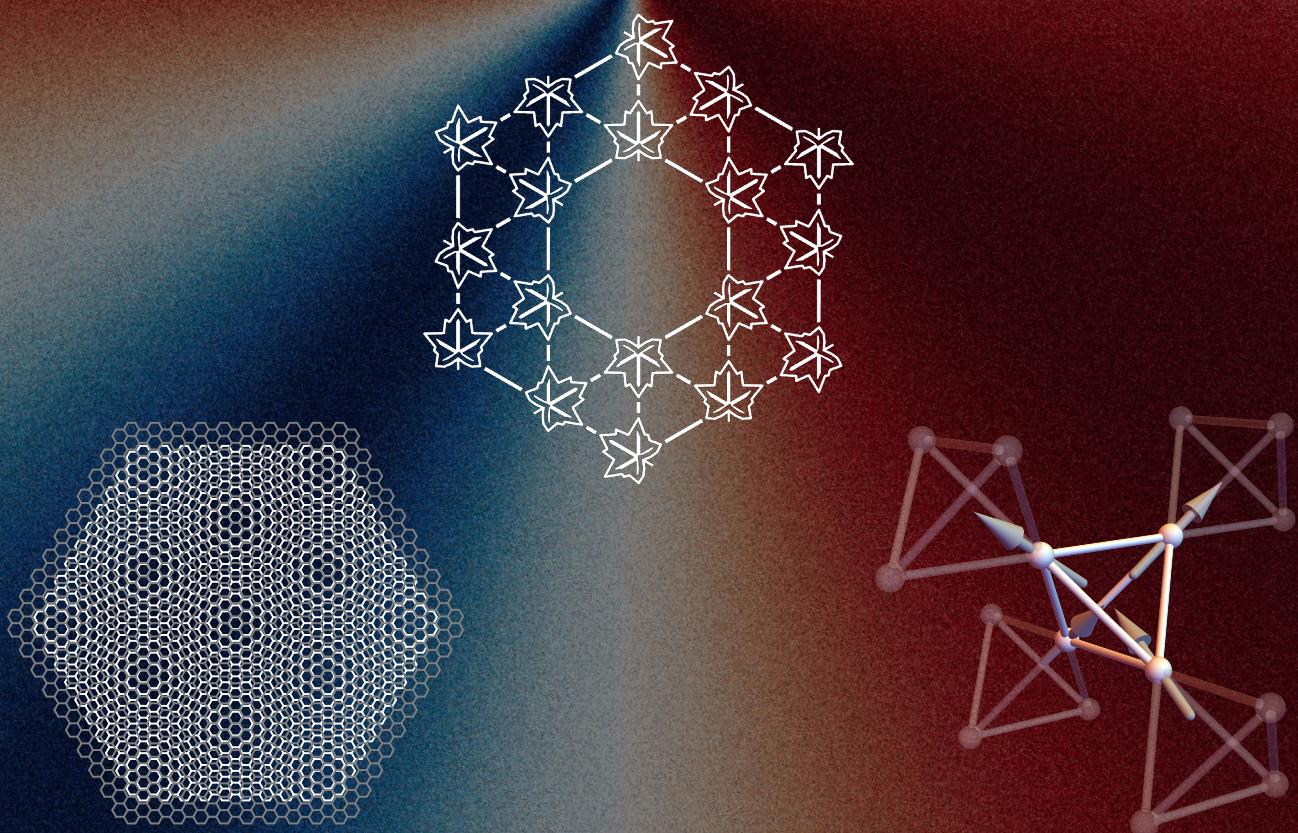


Exploring
Quantum Magnetism
in
Moiré, Maple-leaf, and
Pyrochlore Materials
via Functional Renormalization



Lasse Gresista

Dissertation

Exploring Quantum Magnetism in Moiré, Maple-Leaf, and Pyrochlore Materials via Functional Renormalization

Inaugural-Dissertation

zur

Erlangung des Doktorgrades

der Mathematisch-Naturwissenschaftlichen Fakultät

der Universität zu Köln

vorgelegt von

Lasse Gresista

aus Troisdorf



Angenommen im Jahr 2025

Abstract

Magnets have fascinated people for centuries and strongly shaped society—from compass needles guiding explorers to electric motors powering modern technology. The best-known examples are ferromagnets, where aligned electron spins produce macroscopic magnetization. Depending on the ions and crystal structure, however, many other magnetically ordered states can emerge, such as antiferromagnets with alternating spin alignment. This thesis focuses on even more exotic cases in which competing interactions in combination with quantum fluctuations suppress such order entirely—even at zero temperature. In such systems, spins instead remain in a quantum superposition of many configurations with strong long-range entanglement, forming what is known as a quantum spin liquid (QSL).

QSLs are a phase of matter beyond the Landau paradigm, as—by definition—they do not break any conventional symmetries. Moreover, their entangled nature produces fractionalized excitations—quasiparticles carrying only fractions of the electron’s quantum numbers—which in two dimensions can even obey anyonic statistics and have no counterpart in fundamental particles. These properties make QSLs a fascinating phase of matter, in principle readily accessible through tabletop experiments. Furthermore, initial interest in QSLs was motivated by proposed links to high-temperature superconductivity and topological quantum computing, highlighting their potential practical relevance. Yet since Anderson’s original proposal in 1973, no material has been unambiguously identified to host a QSL.

In this thesis, we numerically study quantum spin models for three promising classes of candidate systems: pyrochlore rare-earth oxides, maple-leaf magnets, and moiré materials. Each hosts mechanisms that are believed to amplify the effects of quantum fluctuations and thus may stabilize exotic quantum ground states without classical counterparts.

In pyrochlore rare-earth oxides, strong spin–orbit coupling gives rise to highly anisotropic interactions that can produce strong frustration. Prime examples are the classical spin-ice materials, which realize a classical spin liquid state described by an emergent gauge theory. For a more general Hamiltonian relevant to many of these compounds, a broad family of classical spin liquids has been identified, characterized by exotic higher-rank emergent gauge theories. We investigate how quantum fluctuations affect the classical spin-liquid phases and discuss the corresponding implications for rare-earth compounds.

Motivated by recent material realizations, we also investigate spin models on the maple-leaf lattice—a two-dimensional network of triangles and hexagons closely related to the triangular and highly frustrated kagome lattice. Spin models on this lattice have been shown to host several valence-bond solid states close to antiferromagnetic phases, which may enable exotic phase transitions involving deconfined criticality or intermediate spin-liquid regimes. We map out the as of yet relatively unexplored phase diagram of the nearest-neighbor model and explore the effect of longer-range interactions, which on related lattices are known to stabilize chiral spin liquids.

The last class we study are moiré materials, where stacking two-dimensional layers with a twist or lattice mismatch induces strong correlations in otherwise weakly interacting systems. Their tunability has revealed a wide range of correlated phases, including Mott insulating states and superconductivity, most prominently in twisted bilayer graphene and related materials. In the strong-coupling regime relevant to the insulating phases, spin becomes entangled with an emergent valley degree of freedom, giving rise to local moments transforming under an enlarged

SU(4) symmetry, that is believed to amplify quantum fluctuations. We investigate such a model for trilayer graphene aligned with hexagonal boron nitride (TG/h-BN), identifying both novel ordered states and putative spin–valley entangled quantum spin liquids.

All of these systems are governed by frustrated quantum spin models, which remain notoriously difficult to analyze even with state-of-the-art many-body methods. The challenges are amplified by the three-dimensional geometry and anisotropic interactions of pyrochlores, the large unit cell of the maple-leaf lattice, and the enlarged local Hilbert space of SU(4) spin–valley models. A central part of this thesis is therefore devoted to extending the pseudo-fermion functional renormalization group (pf-FRG) to these settings. The pf-FRG is by now a well-established method for distinguishing ordered from quantum paramagnetic phases across a broad range of frustrated systems where most other approaches fail. In addition, we develop and apply complementary methods—including classical and mean-field approaches—to support the pf-FRG analysis. Most notably, we introduce a semiclassical Monte Carlo method tailored to SU(4) spin models, which allows us to investigate exotic ordered states of spin–valley and spin–orbit entangled materials.

Contents

| | | |
|----------------|--|-----------|
| Outline | 1 | |
| 1 | Introduction | 3 |
| 1.1 | Classical spins, frustration and classical spin liquids | 6 |
| 1.2 | Quantum fluctuations, valence bond solids and quantum spin liquids | 9 |
| 1.3 | Numerical methods for quantum spin models | 13 |
| 2 | Modeling Mott insulating materials by quantum spin models | 17 |
| 2.1 | Magnetic moments and angular momentum | 17 |
| 2.1.1 | Magnetic moment of a single particle | 18 |
| 2.1.2 | Magnetic moment of an isolated ion | 19 |
| 2.1.3 | Magnetic ions in a crystal field | 20 |
| 2.2 | Localized magnetic moments in solids and the exchange interaction | 22 |
| 2.2.1 | Weak electronic correlations and the tight-binding model | 22 |
| 2.2.2 | Strong electronic correlations and the Hubbard model | 25 |
| 2.2.3 | From the Hubbard to the Heisenberg model and the exchange interaction | 26 |
| 2.2.4 | Spin-orbit entangled Mott insulators and anisotropic exchange | 28 |
| 2.2.5 | Effective SU(4) models for spin-orbit and spin-valley entangled materials | 29 |
| I | Methods | 31 |
| 3 | Classical and mean-field approximations | 33 |
| 3.1 | The Luttinger-Tisza method for classical spin models | 33 |
| 3.1.1 | The classical limit for SU(2) spins | 34 |
| 3.1.2 | The strong and weak constraint | 35 |
| 3.1.3 | Diagonalization in real space | 35 |
| 3.1.4 | Diagonalization in momentum space | 37 |
| 3.1.5 | Examples for collinear, coplanar and noncoplanar states | 40 |
| 3.2 | Semiclassical Monte Carlo for SU(4) spin models | 42 |
| 3.2.1 | The local Hilbert-space of SU(4) spin models | 43 |
| 3.2.2 | The semiclassical limit | 44 |
| 3.2.3 | Monte Carlo implementation | 45 |
| 3.3 | Cluster mean-field theory | 50 |
| 3.3.1 | Cluster mean-field approximation | 50 |
| 3.3.2 | Iterative solution of the self-consistent equations | 52 |
| 3.3.3 | Initial conditions | 53 |
| 3.3.4 | Limitations in the cluster geometry | 53 |
| 3.3.5 | Example: The $J_1 - J_2$ Heisenberg model on the square lattice | 54 |

| | |
|--|------------|
| 4 The pseudo-fermion functional renormalization group | 59 |
| 4.1 The fermionic functional renormalization group | 60 |
| 4.1.1 One-particle irreducible correlation functions | 61 |
| 4.1.2 Flow equations | 62 |
| 4.1.3 Truncation of the flow equations | 63 |
| 4.2 Pseudo-fermion functional renormalization group for SU(2) spin models | 64 |
| 4.2.1 Mapping spins to fermions | 65 |
| 4.2.2 Symmetry constraints on the correlation functions | 66 |
| 4.2.3 Vertex parametrization | 67 |
| 4.2.4 pf-FRG flow equations | 69 |
| 4.2.5 Regulator | 71 |
| 4.2.6 Asymptotic frequency parametrization | 72 |
| 4.2.7 Symmetries in spin space | 73 |
| 4.2.8 Single occupation constraint | 77 |
| 4.3 Generalization to SU(4) spin models | 78 |
| 4.3.1 Symmetry constraints on the correlation functions | 80 |
| 4.3.2 Vertex parametrization | 81 |
| 4.3.3 Flow equations | 83 |
| 4.3.4 Asymptotic frequency parametrization | 84 |
| 4.4 Numerical implementation | 85 |
| 4.4.1 Algorithms in the PFFRGSolver.jl package | 86 |
| 4.4.2 Exploiting real- and spin-space symmetries in spin-orbit coupled materials | 87 |
| 4.5 Calculating phase diagrams | 91 |
| 4.5.1 Distinguishing magnetically ordered and paramagnetic states | 93 |
| 4.5.2 Characterizing ordered states | 94 |
| 4.5.3 Characterizing paramagnetic states | 96 |
| 4.5.4 Comparison with neutron-scattering experiments | 97 |
| 4.6 Discussion and outlook | 97 |
| II Applications | 101 |
| 5 The role of quantum fluctuations in pyrochlore rare-earth oxides | 103 |
| 5.1 Spin Hamiltonian | 106 |
| 5.2 Classical phase diagram | 108 |
| 5.2.1 Deriving the classical phase diagram | 108 |
| 5.2.2 Magnetically ordered phases | 109 |
| 5.2.3 Pinch-line spin liquid | 111 |
| 5.2.4 Spin nematic state | 112 |
| 5.3 Effects of quantum fluctuations in the $S = 1/2$ model | 113 |
| 5.3.1 Quantum phase diagram | 113 |
| 5.3.2 Paramagnetic regime | 114 |
| 5.4 Implications for $\text{Yb}_2\text{Ti}_2\text{O}_7$ | 117 |
| 5.4.1 Phase diagram around parameters from Scheie <i>et al.</i> | 118 |
| 5.4.2 Phase diagram around parameters from Thompson <i>et al.</i> and comparison with nonlinear spin-wave theory | 119 |
| 5.4.3 Challenges of parameter fitting in $\text{Yb}_2\text{Ti}_2\text{O}_7$ | 120 |
| 5.5 Discussion | 121 |

| | |
|---|------------|
| 6 Unconventional quantum states in maple-leaf magnets | 123 |
| 6.1 Ordered and valence-bond solid states in the nearest-neighbor model | 126 |
| 6.1.1 Ordered states | 127 |
| 6.1.2 Valence-bond solids | 127 |
| 6.2 Candidate quantum disordered intermediate phase in the Heisenberg antiferromagnet | 128 |
| 6.2.1 Melting of magnetic order | 129 |
| 6.2.2 Onset of dimer singlet order | 129 |
| 6.2.3 Discussion | 131 |
| 6.3 Quantum states in the ferro–antiferromagnetic Heisenberg model | 132 |
| 6.3.1 Schematic phase diagram | 132 |
| 6.3.2 Classical phase diagram from Luttinger–Tisza | 133 |
| 6.3.3 Cluster mean-field theory | 134 |
| 6.3.4 Pseudo-fermion functional renormalization group | 138 |
| 6.3.5 Discussion | 141 |
| 6.4 Noncoplanar orders and putative chiral quantum spin liquids from cross-plaquette interactions | 142 |
| 6.4.1 Methods for constructing the classical phase diagram | 142 |
| 6.4.2 Phase diagram of the antiferromagnet | 143 |
| 6.4.3 Phase diagram of the ferromagnet | 145 |
| 6.5 Discussion | 147 |
| 7 Spin-valley magnetism in moiré materials | 149 |
| 7.1 Single-layer graphene in a moiré potential | 150 |
| 7.1.1 Derivation of the spin-valley model | 151 |
| 7.1.2 Mean-field phase diagram | 153 |
| 7.1.3 Effects of thermal fluctuations | 155 |
| 7.1.4 Discussion | 158 |
| 7.2 Spin-valley magnetism in TG/h-BN | 159 |
| 7.2.1 Derivation of the spin-valley model | 160 |
| 7.2.2 Ordered states at half-filling | 162 |
| 7.2.3 Key results | 163 |
| 7.2.4 Semiclassical analysis | 165 |
| 7.2.5 Quantum fluctuations | 169 |
| 7.2.6 Discussion | 170 |
| 8 Concluding remarks | 173 |
| A Appendix for Chapter 4 | 175 |
| A.1 Susceptibilities from pf-FRG | 175 |
| B Appendix for Chapter 5 | 177 |
| B.1 Hamiltonian in the global and local frame | 177 |
| B.2 Irrep decomposition and order-parameter fields | 179 |
| B.3 Supplemental data for the pf-FRG calculation | 179 |
| C Appendix for Chapter 6 | 181 |
| C.1 Definition of the maple-leaf lattice | 181 |

| | | |
|---------------------|--|------------|
| C.2 | Evolution of the critical scale for the ferro–antiferromagnetic Heisenberg model | 181 |
| C.3 | Supplemental data: cross-plaquette interactions | 182 |
| D | Appendix for Chapter 7 | 187 |
| D.1 | Vertex parametrization for an $SU(2) \otimes U(1)$ symmetric spin-valley model | 187 |
| D.2 | Thermodynamics in TG/h-BN | 187 |
| D.3 | Supplemental data for the semiclassical and quantum phase diagrams | 188 |
| Bibliography | | 193 |

Outline

This thesis develops and applies several numerical approaches—most notably the pseudo-fermion functional renormalization group (pf-FRG)—to three classes of candidate quantum materials: pyrochlore rare-earth oxides, maple-leaf magnets, and moiré systems. Beyond the two introductory chapters, the thesis is divided into **Part I**, which discusses the methods employed, and **Part II**, which presents the results of applying them. The detailed structure is as follows.

Chapter 1 introduces the fundamental theoretical concepts underlying exotic ground states in quantum spin models. We first discuss how frustration can give rise to classical spin liquids and then how quantum fluctuations may stabilize unconventional quantum paramagnets such as valence-bond solids and quantum spin liquids. We also highlight the types of interactions believed to favor these phases, motivating the material classes studied later. **Chapter 2** explains how magnetic Mott-insulating materials can be described by quantum spin models, focusing on the formation of local magnetic moments, the microscopic mechanisms giving rise to the various exchange interactions studied later in this thesis, and the emergence of SU(4) spin models in spin-valley- or spin-orbital-entangled systems.

In **Part I**, **Chapter 3** presents classical and mean-field methods used to complement pf-FRG. We begin with the Luttinger–Tisza method, an elegant approach to classical ground states. We then introduce a semiclassical limit for SU(4) spin models and a Monte Carlo scheme we developed to study their ground states and finite-temperature behavior. We conclude with a description of cluster mean-field theory (CMFT), which incorporates short-range quantum fluctuations at much lower cost than pf-FRG. **Chapter 4** introduces the concepts and definitions underlying pf-FRG, its extension to spin-orbit coupled models with anisotropic interactions, and to SU(4) spin-valley models. We also describe how phase diagrams are obtained in practice.

Part II begins with **Chapter 5**, where we investigate the role of quantum fluctuations in pyrochlore rare-earth oxides. Particular attention is given to the fate of classical higher-rank spin liquids and the implications for real compounds, with a focus on $\text{Yb}_2\text{Ti}_2\text{O}_7$, whose low-energy behavior is not well captured by pure classical treatments. In **Chapter 6**, we combine the Luttinger–Tisza method, CMFT, and pf-FRG to map the phase diagram of the maple-leaf lattice, both for different nearest-neighbor interactions, where multiple valence-bond solid phases arise, and for longer-range cross-plaquette couplings, which may stabilize chiral spin liquids. **Chapter 7** first uses our semiclassical Monte Carlo approach to study spin-valley orders in a model for single-layer graphene subject to a smooth moiré potential, induced by an insulating substrate layer. We then analyze the strong-coupling limit of a spin-valley model derived for trilayer graphene aligned with hexagonal boron nitride (TG/h-BN), investigating possible spin-valley liquid phases with our SU(4) extension of pf-FRG. Finally, **Chapter 8** offers concluding remarks and an outlook.

Chapter 1

Introduction

Magnetic materials have fascinated people for thousands of years. Ancient texts from Greece, India, and China already describe the curious power of the lodestone: a rock that mysteriously pulls on iron. Even today, I've yet to meet anyone who isn't at least a little delighted when given two permanent magnets to fiddle with. We now understand that a lodestone is actually a naturally magnetized crystal of magnetite—an iron oxide with the formula Fe_3O_4 . Materials like magnetite that can hold on to a magnetic field even without any external influence are called *ferromagnets*. They've been known for a very long time, and their impact on society has been enormous: from guiding early explorers with compasses, to driving electric motors and powering most of the devices we use every day, and—perhaps most importantly—keeping souvenir magnets firmly attached to our refrigerators.

A good understanding of the underlying effect, however, only came in the 20th century, with the discovery of the electron and the development of quantum mechanics. As it turns out, magnetism in matter is an inherently quantum mechanical phenomenon. Its origin lies in the *quantized* angular momentum of elementary particles—mostly of the electrons. This angular momentum has two parts: an orbital contribution, which depends on the electron's state, and an intrinsic part called *spin*, which is an inherent property of every fundamental particle. Together, these give each particle a tiny magnetic moment, similar to how a current circling a loop of wire induces a magnetic field. These magnetic moments can be thought of as small magnets in their own right—although incredibly small, quantum mechanical versions.

Many materials have their electrons arranged so that the angular momenta cancel out, leaving no net magnetic moment. That is why the majority of substances hardly react to external magnetic fields. In ferromagnets, however, some electrons carry uncompensated angular momentum, giving rise to finite microscopic magnetic moments, and crucially, the quantum nature of electrons drives an interaction that favors their alignment. At sufficiently low temperatures, these microscopic moments line up and add together, producing a macroscopic magnetization—a permanent magnet strong enough to stick to a fridge.¹

Although most familiar from daily life, ferromagnetism is only one of many ways in which the local magnetic moments of a material can arrange, or *order*. Magnetism is remarkably rich: both the type and strength of local moments, as well as the effective interactions between them, depend intricately on the chemical elements involved and how they are arranged in the crystal lattice. These interactions can take many forms—they may favor aligned or anti-aligned moments, act over long or short distances, prefer specific spatial directions, or vary from one atomic pair to the next. The result is a whole zoo of possible magnetic orders: from ferromagnets, where all moments align, and antiferromagnets, where neighbors point in opposite directions, to more complex textures such as two-dimensional spin spirals or even three-dimensional, topologically

¹In real materials, the moments do not all align at once. Instead, the crystal divides into domains whose magnetizations point in different directions and largely cancel. An external field can reorient and merge these domains, producing a net magnetization that remains even after the field is removed [1].

nontrivial structures like skyrmions [2] and hopfions [3]. In other cases, the interactions are too weak to cause spontaneous order. The moments then only line up when an external field is applied, and fall back into disorder as soon as it is switched off. This is called *paramagnetism*.

Both ordered and paramagnetic states have something in common: they can be well described within a classical picture. Their local magnetic moments behave like tiny compass needles pointing in different directions in three-dimensional space. Materials of this kind—such as the ferromagnet magnetite—are relatively well understood, both theoretically and experimentally. But not all magnets fit this picture. When competing interactions favor different alignments of the magnetic moments, ground states may emerge with *no* classical analog, with properties that are inherently quantum mechanical. Rather than settling into a single spin configuration, they exist in a vast quantum superposition of many states, bound together by strong, long-range *entanglement*. Crucially, this entanglement cannot be tuned away gradually—returning to an essentially classical state requires crossing a phase transition. Such states therefore form a distinct phase of matter, and in magnetic systems they are known as quantum spin liquids (QSLs) [4].

Unlike ordered magnets or paramagnets, QSLs are notoriously difficult to describe. In fact, there is still no universally agreed-upon definition of what exactly constitutes a QSL, since “entanglement” itself is not easily quantified [4]. What all QSLs share, and what gives them their name, is that despite strong interactions between local moments, the system never settles into an ordered pattern—even at absolute zero. In this sense, they behave like a liquid that never freezes, no matter how much it is cooled (the only known physical example of this is liquid helium). And the mystery doesn’t stop at theory. Since Anderson first proposed the QSL in 1973 [5], no material has yet been unambiguously shown to host such a state. Still, a large and active research community—including the work presented in this thesis—is devoted to understanding QSLs and to finally realizing them in experiment. Which naturally raises the question: why should we care about this elusive state of matter at all?

There are multiple answers to this question, which I will explore in more detail later. Broadly, the motivations fall into two categories. The first is their genuinely exotic nature. QSLs represent a phase of matter outside the traditional Ginzburg–Landau framework [6, 7], as they specifically do *not* spontaneously break the symmetries of the system.² Moreover, when these systems are excited—for instance by heating—the resulting quasiparticles can carry only fractions of the electron’s quantum numbers, with no counterparts among the known fundamental particles [6]. The possibility of studying such deeply fundamental phenomena in a crystal, accessible through tabletop experiments, is fascinating in its own right.

The second category of motivation is more practical, though also more speculative. A subset of QSLs feature topological order—a global property that makes their excitations more robust against local perturbations [6]. This has made topological QSLs attractive candidates for quantum computing, where stability against noise is one of the central challenges [8]. Current efforts, such as Microsoft’s pursuit of a topological quantum computer, rely on hybrid semiconductor–superconductor nanowire devices that are believed to host similar excitations, rather than on intrinsic spin-liquid materials [9]. The discovery of a true topological spin liquid in a magnetic material could be transformative for the development of robust, large-scale quantum devices. Finally, the motivation initially driving research on QSLs was Anderson’s proposal [10] of a

²Some QSLs, such as chiral spin liquids, do break additional symmetries—for example, time-reversal.

In such cases, symmetry breaking is part of their character, but not their only defining feature: what still distinguishes them from conventional magnets is the absence of long-range order, the presence of long-range entanglement and fractionalized excitations.

possible link to high-temperature superconductivity—an effect still poorly understood [11, 12] but of enormous technological potential. Although this connection remains speculative, gaining insight into QSLs may eventually help in the search for superconductors with higher transition temperatures, potentially paving the way for practical applications.

Regardless of the motivation, the main objective of the research that culminated in this thesis is to explore novel classes of magnetic materials in which quantum effects—often referred to collectively as *quantum fluctuations*—are expected to play a significant role. Such systems are promising candidates to host quantum spin liquids as well as other inherently quantum ground states. As theoretical physicists, we don’t study these materials directly in the lab. Instead, we work with effective models that capture their essential physics—whether for materials that already exist or ones that might be realized in the future. In particular, this thesis focuses on materials in which the electrons are strongly localized to the ions—commonly referred to as Mott insulators—which can be effectively described by *quantum spin models* that capture the interactions between the localized magnetic moments within the crystal.

These models are notoriously hard to solve, even with the most advanced computational tools in solid-state physics. A major part of this thesis is therefore devoted to further developing computational methods for spin models that lie beyond the reach of existing approaches. Alongside the use of various classical, semiclassical, and mean-field approaches—which heavily approximate or even neglect quantum fluctuations—the main focus of this work is on extending and applying the *pseudo-fermion functional renormalization group* (pf-FRG) to new classes of materials [13, 14]. By now, pf-FRG is a well-established method, capable of distinguishing between conventional long-range order and unconventional quantum ground states by systematically incorporating quantum fluctuations beyond the mean-field level. Importantly, it can also handle highly frustrated spin models in both two *and* three dimensions—a regime where most other methods break down [14]. In this thesis, we have extended and applied the pf-FRG to three novel material classes: pyrochlore rare-earth oxides, maple-leaf magnets, and moiré systems. All of them feature different types of interactions, crystal structures, and, in the case of moiré materials, even novel types of effective local magnetic moments. The guiding question throughout is: *What role do quantum fluctuations play in these systems, and can they stabilize ground states that are absent in a purely classical analysis?*

The results of our numerical analysis, along with introductions to the physical properties and broader relevance of the studied materials, are presented in the respective chapters of Part II (Applications). Before that, Part I (Methods) is devoted to the computational techniques employed in this work, with particular emphasis on the new developments and implementations made during the course of the thesis. The most significant advances are the extension of our pf-FRG implementation to spin-orbit-coupled systems with highly anisotropic bond-dependent interactions, and to spin models formulated in terms of SU(4) generators rather than the conventional SU(2) spin operators. The remainder of the introduction sets the stage by outlining the common themes that give rise to unconventional states in magnetic materials. We begin with conventionally ordered phases that admit a classical description, and then illustrate how *frustration* and quantum fluctuations can drive the emergence of unconventional quantum states—including, but not limited to, QSLs. This motivates the spin models and material classes studied in this thesis. We then provide a brief overview of standard numerical approaches to quantum many-body systems, concluding with an explanation of why we ultimately employ the pf-FRG. While this first chapter emphasizes theoretical concepts and effective models, Chapter 2 provides an intuitive motivation for why materials can be described by quantum spin models in the first place, and discusses the microscopic origins of the different interactions that can arise.

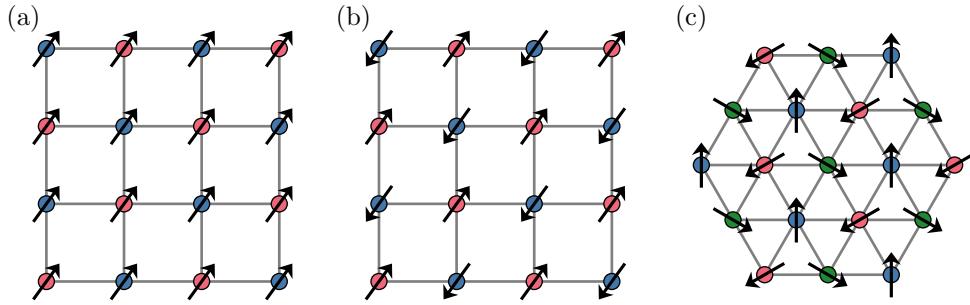


Figure 1.1 – Classically ordered states (a) Ferromagnetic order on the square lattice. (b) Néel order on square lattice. (c) 120° order on the triangular lattice. The colors highlight the bipartite (tripartite) nature of the square (triangular) lattice.

1.1 Classical spins, frustration and classical spin liquids

In this thesis, we focus on magnetic materials in which magnetism arises from electrons that are strongly localized to the magnetic ions of a crystal lattice³. In such systems—commonly referred to as Mott insulators—the electrons are bound to the ions and do not move through the crystal. Their spin and orbital degrees of freedom, and the associated magnetic moments, however, remain free to fluctuate. Although our ultimate aim is to study quantum effects, as already emphasized in the introduction, many magnetic materials can be captured to a good extent by a classical description. Such an approach is particularly suitable at high temperatures, where thermal fluctuations can destroy quantum coherence, and—as we will see later—also for local moments originating from *large* angular momentum J .

Frustration in classical spin models In the classical description of Mott insulators, the local magnetic moments—often simply referred to as spins—are represented by normalized three-dimensional vectors $\mathbf{S}_i = (S_i^x, S_i^y, S_i^z) \in \mathbb{R}^3$ of fixed length $|\mathbf{S}_i|^2 = 1$, localized on site i of a crystal lattice. The simplest model for their interaction is the Heisenberg Hamiltonian

$$H = J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j, \quad (1.1)$$

where the sum goes over nearest neighbors $\langle ij \rangle$. The sign of the exchange constant J determines whether neighboring spins prefer to align parallel (ferromagnetic, $J < 0$) or antiparallel (antiferromagnetic, $J > 0$).

In the ferromagnetic case, energy is minimized by all spins pointing in the same direction, as shown for the square lattice in Fig. 1.1(a). Antiferromagnets, on the other hand, are more subtle, as the ground state depends intricately on the lattice geometry. If the lattice is *bipartite*, meaning that its sites can be divided into two sublattices such that every nearest-neighbor bond connects different sublattices (as highlighted for the square lattice in Fig. 1.1), the ground state is the *Néel state*, with neighboring spins exactly antiparallel [illustrated in Fig. 1.1(b)]. On non-bipartite lattices, however, it is impossible to form a state that fully minimizes all bond energies simultaneously—a phenomenon known as *geometric frustration*. On the triangular lattice, for example, not all three spins in a triangle can be pairwise antiparallel [Fig. 1.2(a)]. Instead, the

³Magnetism can also arise from the collective behavior of itinerant electrons in metals, but such cases will not be considered in this thesis.

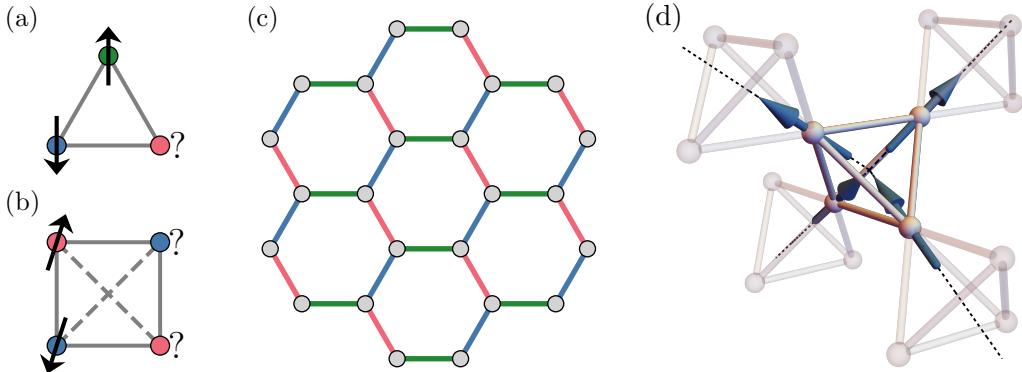


Figure 1.2 – Examples of frustration. (a) Antiferromagnetically coupled Ising spins on a single triangle illustrate the basic case of geometric frustration. (b) On the square lattice, adding next-nearest-neighbor interactions leads to exchange frustration. (c) The Kitaev honeycomb model, where each color represents a coupling $J_a S_i^a S_j^a$ ($a = x, y, z$) between different spin components, provides another example of exchange frustration via anisotropic exchange. (d) Ising spins constrained to the local $\langle 111 \rangle$ axes on the pyrochlore lattice (dotted lines). The blue arrows depict a two-in/two-out configuration characteristic of classical spin ice, a spin-liquid state that emerges from geometric frustration with antiferromagnetic coupling.

minimum-energy configuration is one in which neighboring spins form 120° angles, as shown in Fig. 1.1(c).

Frustration can also arise without special lattice geometry, but instead through competing interactions. A classic example is the square-lattice Heisenberg model with added next-nearest-neighbor couplings [Fig. 1.2(b)], which compete with the nearest-neighbor exchange. Another is the Kitaev honeycomb model [15], defined by the Hamiltonian

$$H = J_x \sum_{(ij) \in x\text{-bonds}} S_i^x S_j^x + J_y \sum_{(ij) \in y\text{-bonds}} S_i^y S_j^y + J_z \sum_{(ij) \in z\text{-bonds}} S_i^z S_j^z, \quad (1.2)$$

where the Ising-like interactions depend on the orientation of the bond, as illustrated in Fig. 1.2(c). Although the honeycomb lattice is bipartite and thus unfrustrated in the geometric sense, the bond-dependent couplings generate a distinct form of exchange frustration: a spin aligned to minimize an x -bond interaction cannot simultaneously minimize the y - and z -bond terms.

Classical spin liquids On the triangular-lattice antiferromagnet, the system resolves its frustration by forming a 120° ordered state. This state is the unique ground state, apart from global spin rotations and a twofold degeneracy associated with the clockwise or counterclockwise winding of the spins on each elementary triangle, known as the chirality. For other models, such as the Heisenberg model on the kagome lattice [made of corner-sharing triangles] or the pyrochlore lattice [made of corner-sharing tetrahedra, Fig. 1.2(d)], the effect of frustration is more severe: for AFM exchange, there is no *unique* ordered ground state. Instead, an *extensive* number of configurations minimize the classical Hamiltonian, the vast majority of which are disordered and show no regular pattern. Such a system retains a finite *residual entropy* even at zero temperature.

If a frustrated model exhibits an extensively degenerate ground-state manifold, two scenarios are possible once temperature effects are taken into account. In some cases, a discrete subset of configurations carries higher entropy than the rest, so thermal fluctuations select these states and induce long-range order at any small but finite temperature. This counterintuitive mechanism is known as *order-by-disorder* [16], and explains, for example, the selection of coplanar states in the classical kagome antiferromagnet [17].

In other cases, such as the classical pyrochlore antiferromagnet, thermal fluctuations are insufficient to lift the degeneracy, and the system remains disordered even at the lowest temperatures [18]. The corresponding low-temperature phase is then referred to as a *classical spin liquid* and, despite being disordered, already displays remarkable emergent properties within a purely classical description.

Emergent gauge theories To make this more transparent, let us stay on the pyrochlore lattice but simplify the problem by considering the Ising limit, where each spin is constrained to point along a single axis \hat{z} . For concreteness, we take \hat{z} to point from the center of each tetrahedron toward its vertices [see Fig. 1.2(d)], so that every spin can only point ‘in’ or ‘out’ of the tetrahedron. In this limit, the Hamiltonian can be rewritten as

$$H = J \sum_{\langle ij \rangle} S_i^{\hat{z}} S_j^{\hat{z}} = \frac{J}{2} \sum_t \left| \sum_{i \in t} S_i^{\hat{z}} \right|^2 + \text{const.}, \quad (1.3)$$

where, in the second step, the sum runs over all elementary tetrahedra of the lattice. Since, in this form, the Hamiltonian contains only positive terms, it is minimized by spin configurations satisfying the local *constraint*

$$\sum_{i \in t} S_i^{\hat{z}} = 0. \quad (1.4)$$

This condition requires that, on each tetrahedron, exactly two spins point inwards and two point outwards. For a single tetrahedron, six configurations satisfy this constraint, and when extended to the infinite lattice in the thermodynamic limit, they give rise to an *extensively* degenerate ground-state manifold. Thermal fluctuations again fail to lift this degeneracy. Remarkably, the residual entropy of this state is exactly the same as for the proton disorder in water ice first predicted by Pauling [19]. The magnetic analogue has therefore been termed *spin ice* [20].

The spin-ice state becomes particularly striking when viewed through its effective low-energy description. We can define an effective field \mathbf{B} that represents the local magnetization and lives on the bonds of the lattice connecting the tetrahedron centers [depicted as blue arrows in Fig. 1.2(d)]. Upon coarse-graining, the spin-ice constraint translates into a Gauss’ law for the effective field

$$\sum_a \partial_a B^a \equiv \nabla \cdot \mathbf{B} = 0, \quad (1.5)$$

directly analogous to classical electromagnetism. This constraint can be resolved by introducing a vector potential \mathbf{A} such that $\mathbf{B} = \nabla \times \mathbf{A}$, with \mathbf{A} defined only up to U(1) gauge transformations. Spin ice is therefore a textbook example of an *emergent gauge theory*, a recurring theme in the broader field of spin liquids.

The simplest excitations above the spin-ice manifold correspond to flipping a single spin. This locally violates the constraint on the two tetrahedra sharing that spin, which in the gauge-theory language corresponds to the appearance of finite monopole charges $\rho = \nabla \cdot \mathbf{B}$. Remarkably, the two monopoles can be separated without any additional energy cost by flipping further spins along a connected path—meaning they are *deconfined*. This provides a particularly simple illustration of *fractionalization*: a single elementary excitation (a spin flip) splits into two independent monopole excitations connected by a Dirac string. In this way, spin ice offers the rare opportunity to study the physics of emergent magnetic monopoles—objects that most likely do not exist as fundamental particles—within a purely condensed-matter setting.

And although all of this sounds very theoretical, the most remarkable fact is that the spin-ice phase is actually realized in materials. The famous dipolar spin-ice compounds $\text{Dy}_2\text{Ti}_2\text{O}_7$ and

$\text{Ho}_2\text{Ti}_2\text{O}_6$ [20] are strongly believed to host the spin-ice state at low temperatures and experimental probes such as neutron scattering have observed signatures consistent with the predicted emergent gauge field and deconfined monopole excitations [21–23]. The dipolar spin-ice compounds belong to the larger material class of pyrochlore rare-earth oxides [24]. It was recently discovered that the phase diagrams of Hamiltonians relevant to many of these compounds host a rich variety of classical spin-liquid phases [25, 26]. Many of them are generalizations of the spin-ice state, where the ground state manifold is also defined by a microscopic constraints that leads to low energy descriptions via generalization of the Gauss' laws and emergent *higher rank gauge theories* [27, 28]. We explore this in greater detail in Chapter 5, where we investigate the role of quantum fluctuations in pyrochlore rare-earth oxides and analyze a concrete example of such an exotic classical spin-liquid phase.

1.2 Quantum fluctuations, valence bond solids and quantum spin liquids

We have seen that frustration can already give rise to unusual phases of matter within a purely classical description of magnetic moments. In real materials, however, the underlying electrons are quantum objects. Their intrinsic quantum nature introduces fluctuations that persist even at zero temperature and, when combined with frustration, can drastically alter the low-temperature physics.

Quantum spins To capture these quantum effects, we must describe the quantized angular momentum of the localized electrons—which we will, for simplicity, continue to refer to collectively as “spin”. Unlike in the classical picture, spins are no longer three-dimensional vectors but quantum operators $\mathbf{S}_i = (S_i^x, S_i^y, S_i^z)$ acting on the local Hilbert space at site i . Mathematically, the components of the spin operator are generators of the Lie algebra $\mathfrak{su}(2)$ of $\text{SU}(2)$, or in other words, of the rotations in the complex spin space. They are defined by the commutation relations

$$[S_i^a, S_j^b] = \delta_{ij} \epsilon_{abc} S_i^c, \quad (1.6)$$

where we set $\hbar = 1$ throughout this thesis. Because different spin components do not commute, they cannot be measured simultaneously.

The size of the local spin is specified by the spin quantum number S , which determines the “magnitude” of the angular momentum. It is defined by the eigenvalue of the square of the spin operator⁴

$$S_i^2 = S(S+1)\mathbb{1}, \quad (1.7)$$

which we assume is the same across all sites i . For a given S , the projection of the spin along a chosen quantization axis (conventionally the z -axis) takes the discrete values $m = -S, -S+1, \dots, S$. The corresponding eigenstates are labeled by $|S, m\rangle$ (or just $|m\rangle$ if S is fixed) and satisfy

$$S^2 |S, m\rangle = S(S+1) |S, m\rangle \quad S^z |S, m\rangle = m |S, m\rangle, \quad (1.8)$$

where we now neglect the site dependence for brevity. These states form a complete basis of the $(2S+1)$ -dimensional local Hilbert space. Since the spin operators do not commute, these basis states cannot be simultaneous eigenstates of S^x and S^y . Their action is best understood by introducing the ladder operators

$$S^+ = S^x + iS^y \quad S^- = S^x - iS^y \quad (1.9)$$

⁴In group theory language, this is the quadratic Casimir element of the Lie-Algebra of $\text{SU}(2)$

which act on the basis states by raising or lowering the magnetic quantum number m as

$$S^+ |S, m\rangle = \begin{cases} |S, m+1\rangle & \text{if } m < S \\ 0 & \text{if } m = S \end{cases} \quad S^- |S, m\rangle = \begin{cases} |S, m-1\rangle & \text{if } m > -S \\ 0 & \text{if } m = -S \end{cases}. \quad (1.10)$$

In this way, the ladder operators introduce transitions between the different basis states of the local Hilbert space. In this thesis, we will mainly focus on the most fundamental case, $S = \frac{1}{2}$. Here the spin projection takes only two values, $m = \pm \frac{1}{2}$, usually referred to as “spin up” $|\uparrow\rangle$ and “spin down” $|\downarrow\rangle$. In this basis the spin operators reduce to simple Pauli matrices

$$\mathbf{S} = \frac{1}{2}(\boldsymbol{\theta}^x, \boldsymbol{\theta}^y, \boldsymbol{\theta}^z)^T. \quad (1.11)$$

The Heisenberg dimer Some consequences of the quantum nature of spins can already be understood by considering the simplest case of two spin- $\frac{1}{2}$ moments coupled via a Heisenberg interaction (sometimes referred to as the Heisenberg dimer). The corresponding Hamiltonian can be written using spin raising and lowering operators as

$$H = J\mathbf{S}_1\mathbf{S}_2 = J(S_1^x S_2^x + S_1^y S_2^y + S_1^z S_2^z) = JS^z S^z + \frac{J}{2}(S_1^+ S_2^- + S_1^- S_2^+). \quad (1.12)$$

From this form it is straightforward to see that the eigenstates, expressed in the product basis $|m_1 m_2\rangle$, group into multiplets characterized by the total spin S_{tot} , defined as the expectation value $S_{\text{tot}}(S_{\text{tot}} + 1)$ of $\mathbf{S}_{\text{tot}}^2 = (\mathbf{S}_1 + \mathbf{S}_2)^2$. There is a singlet state with $S_{\text{tot}} = 0$ and eigenenergy $E_s = -3J/4$,

$$\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle), \quad (1.13)$$

and the three triplet states with $S_{\text{tot}} = 1$ and eigenenergy $E_t = J/4$

$$|\uparrow\uparrow\rangle, \quad \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle), \quad |\downarrow\downarrow\rangle. \quad (1.14)$$

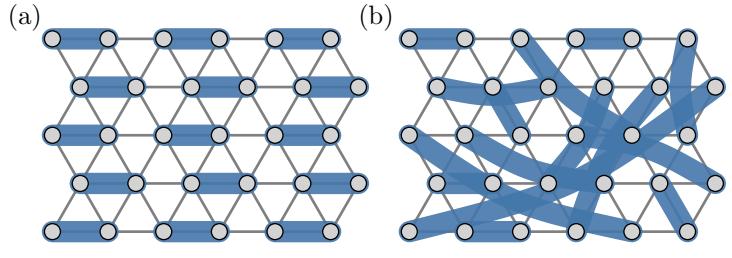
For ferromagnetic exchange ($J < 0$), the triplet manifold forms the ground state. This includes the fully polarized states, which correspond to the classical ferromagnet with maximal local magnetization $|\langle S_i^z \rangle| = 1/2$. For antiferromagnetic exchange ($J > 0$), however, the ground state is the singlet. Unlike the classical Néel state $|\uparrow\downarrow\rangle$, which is *not* an eigenstate of the Hamiltonian, the singlet has vanishing local magnetization, $|\langle S_i \rangle|^2 = 0$, and thus represents a *quantum paramagnet*. This has no classical analog: in the classical picture the local magnetization was fixed to $|\mathbf{S}_i|^2 = 1$. Moreover, the bond energy of the singlet $E_s = -3J/4$, is substantially lower than that of the Néel configuration $E_{\text{Néel}} = -J/4$, highlighting why quantum fluctuations are generally far more pronounced in systems with antiferromagnetic exchange.

Ground states of quantum spin models The two-spin problem illustrates how quantum fluctuations can favor singlets over classical order. The natural next question is what types of quantum ground states can be realized in extended spin systems.

For unfrustrated models, such as ferromagnets, or antiferromagnets on bipartite lattices, the tendency toward conventional long-range order usually prevails: most nearest-neighbor Heisenberg models still develop long-range magnetic order, though quantum fluctuations reduce the local magnetization compared to the classical product state. These ordered states break the global SU(2) spin symmetry and often lattice symmetries as well.

Figure 1.3 –

Valence bond solids on the triangular lattice. Blue bonds indicate pairs of spins forming a singlet state $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$, which can arrange in regular (a) or irregular patterns. A resonating valence bond (RVB) state is the quantum superposition of all possible singlet coverings.



In frustrated systems, however, quantum fluctuations are far more effective at destabilizing classical order, opening the door to unconventional ground states. One natural possibility, directly inspired by the Heisenberg dimer, is that every spin forms a singlet with exactly one other spin, so that the entire lattice is covered by pairwise bonds. Such states are called *valence-bond solids* (VBS), or simply *dimer states* [29]. Two examples on the triangular lattice are shown in Fig. 1.3. Although a VBS has zero local magnetization everywhere—and hence qualifies as a quantum paramagnet—it typically breaks translational or rotational lattice symmetries. There exist many possible valence-bond solids, distinguished by different patterns of singlet coverings, as well as more general singlet phases in which larger clusters of spins form states with $S_{\text{tot}} = 0$ [29]. In Chapter 6, we will explore the possible realization of such states on the relatively unexplored maple-leaf lattice, which is expected to host multiple distinct VBS phases in its phase diagram [30, 31].

An even more intriguing possibility arises when the ground state is not locked into a single dimer pattern but instead forms a coherent superposition of an extensive number of distinct singlet coverings. In such a state, all lattice symmetries are preserved, and the result is a highly entangled phase that breaks no conventional symmetry. This represents the prototype of a quantum spin liquid, first envisioned by Anderson in 1973 under the name *resonating valence-bond* (RVB) state [5]. By now, theorists have proposed a whole zoo of quantum spin liquids [4, 32]. Since they cannot be classified by broken symmetries, they are instead characterized by the nature of their low-energy excitations and the emergent gauge fields that, much like in classical spin ice, mediate their interactions [4, 7]. Excitingly, excitations in QSLs behave as quasiparticles that carry only fractions of the quantum numbers of the original spins—for instance, spinons with spin $S = 1/2$ but no electric charge found in U(1) spin liquids [33]. In gapped *topological* spin liquids in two dimensions, these excitations can even behave as *anyons*, obeying exchange statistics more general than those of fermions or bosons [6].

In practice, determining the true ground state of frustrated spin models is notoriously difficult and even for simple nearest-neighbor antiferromagnetic Heisenberg models, the situation is unresolved on many lattices. A classic example is the triangular lattice antiferromagnet, which was long discussed as a candidate quantum spin liquid but is now believed to realize conventional 120° order [34, 35]. Similarly, the Heisenberg antiferromagnet on the pyrochlore lattice was once thought to host a U(1) quantum spin liquid, but recent numerical studies instead point to a symmetry-breaking valence-bond ground state [36, 37].

These challenges make the rare cases of exactly solvable quantum spin models with established spin-liquid ground states especially celebrated. The most famous such example is the Kitaev honeycomb model [defined in Eq. (1.2) and illustrated in Fig. 1.2(c)] for which Kitaev showed that, by representing the spin operators as real (spinless) Majorana fermions coupled to an emergent \mathbb{Z}_2 gauge field, the model becomes exactly solvable. Its ground state is a long-range entangled quantum spin liquid that is gapless and topologically trivial in the isotropic limit

$J_x = J_y = J_z$. However, when time-reversal symmetry is broken—either explicitly by an applied magnetic field or spontaneously for strong anisotropic couplings (e.g. $J_x \gg J_y, J_z$)—the system becomes gapped and hosts non-Abelian anyonic excitations [15]. This shows that models and materials with highly anisotropic, bond-dependent interactions may also be good candidates to host spin liquid ground states.

Candidate quantum spin liquid materials To summarize, realizing a QSL or other quantum paramagnetic ground state generally requires strong frustration. This can arise in two ways. First, through *geometric frustration*, where isotropic interactions act on highly frustrated lattices such as the triangular, kagome, or pyrochlore lattices. Second, through *exchange frustration*, either from competing nearest- and further-neighbor couplings, or, as motivated by the Kitaev model, from highly anisotropic bond-dependent interactions. Remarkably, both types of interactions are realized in real materials: conventional Heisenberg exchange is present in most Mott insulators [1], while strong anisotropic Kitaev-type interactions can arise in Mott insulators with strong spin-orbit coupling, as first emphasized by Khaliullin and Jackeli [38, 39]. Such interactions are typically encountered in compounds based on heavier 4d and 5d transition metals, or in certain 4f rare-earth magnets [40]. We will return to the microscopic origin of these different interactions in the next chapter. Broadly speaking, candidate quantum spin liquid materials also fall into these two regimes [41]. A paradigmatic example of the geometrically frustrated route is Herbertsmithite $\text{ZnCu}_3(\text{OH})_6\text{Cl}_2$, a $S = 1/2$ kagome antiferromagnet suggested to host a spin-liquid ground state [42, 43]. On the anisotropic side, the most prominent Kitaev material is $\alpha\text{-RuCl}_3$, which is now understood to order magnetically at zero field but is thought to realize a field-induced spin liquid at intermediate fields [44, 45]. There are many more materials exhibiting signatures consistent with a QSL ground state [41], yet unambiguous, smoking-gun evidence for a true QSL phase in any material is still missing.

Enhanced fluctuations in $\text{SU}(4)$ spin models The considerations laid out in this section also guide the choice of materials and spin models studied in this thesis. In Chapter 5, we investigate highly anisotropic nearest-neighbor models on the pyrochlore lattice, while in Chapter 6 we study Heisenberg models with competing interactions on the recently proposed, frustrated maple-leaf lattice. In Chapter 7, by contrast, we explore a different mechanism that enhances the effects of quantum fluctuations.

This mechanism can be realized in systems where an additional bi-valued quantum number beyond spin is present—for example, orbital states in spin-orbital entangled Mott insulators or valley indices in moiré materials built from stacked two-dimensional layers of graphene. The combined spin and orbital/valley degrees of freedom may give rise to local moments described not by the familiar generators of $\text{SU}(2)$ but by generators of the enlarged symmetry group $\text{SU}(4)$ (examples of spin-orbital models include Refs. [46, 47], and for spin-valley models Refs. [48, 49]). We note that models of this form were already explored in the so-called Kugel–Khomskii models in the context of Jahn–Teller physics in transition-metal oxides [50]. Since the Lie algebra of $\text{SU}(N)$ has $N^2 - 1$ generators, $\text{SU}(4)$ provides 15 independent “spin-directions” instead of the three of $\text{SU}(2)$, thereby likely enhancing the potential of quantum fluctuations to destabilize magnetic order. This makes such systems promising candidates for realizing exotic quantum phases, including spin-valley and spin-orbital liquids. For example, on the square lattice the $\text{SU}(2)$ Heisenberg antiferromagnet has long-range Néel order, but in the large- N limit the ground state becomes a “staggered flux spin liquid” [51]. Similarly, while the $\text{SU}(2)$ Heisenberg model on the triangular lattice exhibits magnetic 120° order (see, e.g., [52]), its $\text{SU}(4)$ counterpart is

believed to favor a nonmagnetic valence-bond solid ground state [53]. Aside from the potential for exotic spin-orbital or spin-valley liquid states, interactions that break the $SU(4)$ symmetry can also give rise to novel types of combined spin and valley order. In Chapter 7, we will investigate these possibilities in detail using both the pf-FRG and a semiclassical Monte Carlo approach for $SU(4)$ spin models that we have developed.

1.3 Numerical methods for quantum spin models

We have seen that the spin models most likely to host exotic quantum ground states tend to be the hardest ones to tackle: they live in two or three dimensions, on frustrated lattices, and frequently involve highly anisotropic interactions. These are precisely the regimes where many of the most powerful many-body methods begin to struggle. In the following, I give a brief overview of the standard approaches, their strengths and weaknesses, and why we ultimately turn to the pf-FRG.

Exact diagonalization The most straightforward approach is to brute-force the problem: Restrict to a finite lattice and exactly diagonalize the resulting Hamiltonian matrix numerically. For a system of N spin-1/2 sites the Hilbert space grows exponentially as 2^N , so the method is limited to small clusters. Currently, around $N \sim 30\text{--}40$ sites are possible for typical Heisenberg models (sometimes more, when a lot of symmetries of the spin model can be exploited), so finite-size effects are often severe, especially in three-dimensional models [54]. Nevertheless, in one and two dimensions, ED on these system sizes can sometimes already yield accurate insights into ground states and low-energy excitations. Additionally, within its limits, the results are exact, making exact diagonalization (ED) an invaluable benchmark for more approximate methods.

Quantum Monte Carlo Quantum Monte Carlo (QMC) methods are built on a simple but powerful idea: thermal expectation values of an operator O ,

$$\langle O \rangle = \frac{1}{Z} \text{Tr}[O e^{-\beta H}], \quad (1.15)$$

with partition function $Z = \text{Tr } e^{-\beta H}$, can be rewritten in a path-integral formulation as statistical averages

$$\langle O \rangle = \frac{1}{Z} \int \mathcal{D}\phi O(\phi) p(\phi), \quad (1.16)$$

over configurations ϕ weighted by a probability density $p(\phi)$ that is determined by the Hamiltonian. The integral $\int \mathcal{D}\phi$ ranges over an exponentially large configuration space, but Monte Carlo integration circumvents this by generating a finite set ensemble of configurations $\{\phi_i\}$ sampled according to $p(\phi)$. The infinite integral is then replaced by a finite average,

$$\langle O \rangle \approx \frac{1}{M} \sum_{i=1}^M O(\phi_i), \quad (1.17)$$

which converges to the exact result as the number of samples M increases. Because the most probable configurations are sampled most often, this is remarkably efficient. When QMC is applicable, it essentially provides exact results, up to statistical errors that can be reduced systematically by increasing M [54].

However, for Hamiltonians containing operators with nontrivial anticommutation relations, such as fermions or spins, the weights $p(\phi)$ can become negative or even acquire a complex phase. In this case, the path integral is no longer a sum over positive probabilities but instead involves strong cancellations between contributions. The result is a highly oscillatory integral, for which the statistical error decreases only exponentially slowly with system size or inverse temperature [55]. This difficulty is known as the *sign problem*. With few notable exceptions, most *frustrated* spin models are affected by the sign problem and therefore cannot be efficiently studied with QMC.

Variational approaches A different strategy is to determine the ground state directly by exploiting the variational principle,

$$E(|\psi\rangle) \equiv \frac{\langle\psi|H\psi\rangle}{\langle\psi|\psi\rangle} \geq E_0, \quad (1.18)$$

where E_0 is the ground-state energy. This inequality holds for any state $|\psi\rangle$ in the Hilbert space. The exact ground state can, in principle, be obtained by minimizing $E(|\psi\rangle)$ over the full Hilbert space. For a lattice of N spin-1/2 sites (local Hilbert space dimension 2), a general state can be written as

$$|\psi\rangle = \sum_{m_1, \dots, m_N = \pm 1/2} c_{m_1, \dots, m_N} |m_1 m_2 \dots m_N\rangle, \quad (1.19)$$

requiring 2^N complex coefficients c_{m_1, \dots, m_N} . Just as in ED, this exponential growth restricts exact treatments to very small systems. Instead, variational approaches construct clever Ansätze for the wave function that depend on far fewer parameters, and then minimize the energy within this restricted subspace of the Hilbert space. The reliability of the method hinges on whether the chosen Ansatz is capable of representing the true ground state (or at least a close approximation).

For frustrated magnets, one particularly successful class of Ansätze are *tensor networks* [56]. In one dimension, these take the form of *matrix product states* (MPS), where the wave function is parametrized as

$$|\psi\rangle = \sum_{m_1, \dots, m_N = \pm 1/2} \sum_{a_1, \dots, a_N = 1}^b A_{m_1}^{a_1 a_2} A_{m_2}^{a_2 a_3} \dots A_{m_N}^{a_N a_1} |m_1 m_2 \dots m_N\rangle, \quad (1.20)$$

with A_{m_i} denoting $b \times b$ matrices and b the bond dimension. An MPS thus requires only $N \times 2 \times b^2$ parameters, scaling linearly with system size N . In two dimensions, the natural generalization is to replace matrices A by higher-rank tensors, yielding *projected entangled pair states* (PEPS). The bond dimension b controls how much entanglement can be encoded: for $b = 1$ the state reduces to a simple product state, while larger b allows progressively more entanglement to be captured. It has been shown that all states obeying an area-law in the entanglement entropy can be efficiently represented in this way, i.e., without exponential growth of b [56]. This class is believed to include most *ground states* of gapped local Hamiltonians (including gapped spin-liquid phases). For one-dimensional systems this statement can be proven rigorously [57], and even gapless ground states only weakly violate the area law [58]. Here, the density-matrix renormalization group (DMRG), which provides an efficient algorithm to obtain ground-state MPS, has essentially solved the ground state problem for one-dimensional spin models—even in the presence of frustration [59]. In two dimensions, PEPS are likewise expected to capture most gapped ground states of interest. The practical challenge is that contracting a PEPS is computationally very demanding, with the cost increasing rapidly as the bond dimension is increased. Nevertheless, efficient algorithms

have been developed [60], and variational studies based on PEPS have already been applied with considerable success to a range of two-dimensional frustrated spin models (see, e.g. Refs [61–63]). Three-dimensional systems, however, as well as models with enlarged local Hilbert spaces such as SU(4) spins, are still largely beyond current numerical reach.

Mean-field methods Because many numerical methods fail for frustrated quantum spin models, an alternative route is to simplify the models themselves. The most common such approximation is the mean-field approach. The idea is to replace the complicated many-body interactions a particle experience by an average effective field, effectively reducing the problem to a noninteracting one. In its simplest form for spin systems, one substitutes spin operators by their expectation value as

$$\mathbf{S}_i \mathbf{S}_j \rightarrow \mathbf{S}_i \langle \mathbf{S}_j \rangle + \langle \mathbf{S}_i \rangle \mathbf{S}_j - \langle \mathbf{S}_i \rangle \langle \mathbf{S}_j \rangle \quad (1.21)$$

This, however, can only capture ordered states: for quantum paramagnets $\langle \mathbf{S}_j \rangle = 0$, so the mean-field Hamiltonian vanishes.

A more general strategy is to represent spins in terms of *partons*, which can be bosons or fermions. Here, a common, choice, which is of particular relevance for this thesis, is the Abrikosov fermion representation [64]

$$S_i^a = \frac{1}{2} f_{is'}^\dagger \theta_{s's}^a f_{is} \quad (1.22)$$

where θ^a are Pauli matrices, f_{is}^\dagger creates a fermion at site i with spin $s = \pm 1/2$, and we assume summation over repeated spin indices. As we will discuss in more detail in the context of the pf-FRG in Chapter 4, this mapping faithfully reproduces the spin Hilbert space only if one enforces the single-occupancy constraint

$$f_{is}^\dagger f_{is} = 1, \quad (1.23)$$

which restricts the enlarged fermionic Hilbert space to exactly one fermion per site. In this representation, a general interaction between spins mediated by an exchange matrix \mathbf{J} turns (up to constants) into

$$\begin{aligned} \mathbf{S}_i^T \mathbf{J} \mathbf{S} &= \frac{1}{4} \sum_{ab} \theta_{s'_1 s_1}^a \theta_{s'_2 s_2}^b f_{is'_1}^\dagger f_{is'_2}^\dagger f_{is_2} f_{is_1} \\ &\rightarrow \frac{1}{4} \sum_{ab} \theta_{s'_1 s_1}^a \theta_{s'_2 s_2}^b \frac{1}{2} \left(\langle f_{is'_1}^\dagger f_{is'_2}^\dagger \rangle f_{is_2} f_{is_1} \right) + \left(f_{is'_1}^\dagger f_{is'_2}^\dagger \langle f_{is_2} f_{is_1} \rangle \right) \end{aligned} \quad (1.24)$$

where, in the second step, a mean-field approximation has been applied. The resulting Hamiltonian is quadratic in the fermionic operators and can be readily diagonalized, and the expectation values can be determined self-consistently.

At this stage, however, several choices have already been made: one may decouple in different channels (e.g. grouping creation with annihilation operators, or choosing inter-site versus on-site pairings), and one could equally well represent the spins in terms of bosons rather than fermions. In practice it is usually impossible to explore all possible mean-field Ansätze, so calculations are typically carried out within a particular decoupling scheme [4]. This inevitably introduces a bias, though it can be justified in certain controlled limits—for example, in the $N \rightarrow \infty$ limit of SU(N) models [65]. For frustrated $S = 1/2$ models in two or three dimensions, however, mean-field theory is generally uncontrolled. Nevertheless, such Ansätze have played a central role in developing our understanding of spin liquids: Wen's celebrated classification of quantum spin liquids by their projective symmetry group is based precisely on this type of parton mean-field construction [7].

Pseudo-fermion functional renormalization group. In this thesis we primarily employ the pseudo-fermion functional renormalization group (pf-FRG) [13, 14], which will be described in detail in Chapter 4. The method builds on the Abrikosov fermion representation of spins introduced above, combined with the fermionic functional renormalization group (FRG) [66]. It is therefore not surprising that pf-FRG is closely connected to mean-field theory: in the limit of large spin $S \rightarrow \infty$ [67], where magnetic order is favored, and in the limit $N \rightarrow \infty$ for $SU(N)$ models [68], which tend to stabilize quantum paramagnetic states, pf-FRG reduces exactly to the corresponding mean-field theories. For the physically most relevant case of $S = 1/2$ and $N = 2$, however, pf-FRG goes beyond mean-field by retaining interactions between different decoupling channels. This enables it to distinguish between magnetically ordered and paramagnetic phases without the explicit bias of mean-field approaches. Most importantly, pf-FRG remains applicable in regimes where many other methods break down: frustrated three-dimensional spin systems, models with highly anisotropic exchange [14], and, as we will demonstrate in this thesis, even $SU(4)$ spin models. For the materials and spin models studied in this thesis, pf-FRG thus represents one of the very few—and in some cases the only—unbiased many-body approaches currently available.

Chapter 2

Modeling Mott insulating materials by quantum spin models

In the previous chapter, we explored how quantum fluctuations and frustration can give rise to exotic quantum phases of matter with no classical analogue. We argued that such phases are naturally captured by quantum spin models, particularly in Mott insulating materials where electrons are localized to the ions of a crystal. In this chapter, we take a step back to provide an intuitive motivation for why quantum spin models are an appropriate description of such materials in the first place.

We begin by showing how the angular momentum of a single electron gives rise to a magnetic moment, then discuss how the combined angular momenta of many electrons shape the magnetic properties of an isolated atom, and finally see how the surrounding environment of neighboring atoms can further modify these moments. From there, we'll turn to a central question: why should electrons localize in a solid at all? After all, in a generic crystalline material, electrons are also allowed to move around—the most obvious example being metals where the movement of electrons allows them to conduct electric current. Understanding the mechanisms that lead to their localization is therefore essential in justifying spin-only descriptions. This discussion then naturally also leads to the origin of the effective interactions between the localized moments. We explain how both ferromagnetic and antiferromagnetic Heisenberg interactions may emerge, and then outline how strong spin-orbit coupling can generate highly anisotropic interactions. Finally, we discuss how $SU(4)$ spin models can arise both in spin-orbit entangled Mott insulators and in moiré materials.

This section is intended primarily as a pedagogical introduction, and, in truth, was written as much to clarify my own understanding as to explain it to others. When working in theoretical magnetism, and especially in method development, it is all too easy to take effective spin models for granted without reflecting on their microscopic origin. Therefore, I want to provide a concise and hopefully intuitive overview in this chapter. At the same time, we introduce many of the concepts and notation that will be used later when discussing spin models for real materials in part II.

Most of the material covered here can also be found in standard textbooks. I particularly recommend the chapter on magnetism in matter in *The Feynman lectures on physics - Vol. 2* [69] for great intuitive explanations, and Ref. [1] for a more up-to-date and comprehensive overview, on which this section is strongly based. For derivations in the formalism of second quantization, we closely follow Ref. [70].

2.1 Magnetic moments and angular momentum

In general, the magnetic moment μ of an object, quantifies how strongly it responds to an external magnetic field. To be precise, for the following definitions involving magnetic moments

we really mean the magnetic *dipole* moment. Classically, the magnetic dipole moment determines the torque experienced in a uniform magnetic field \mathbf{B} via

$$\boldsymbol{\tau} = \boldsymbol{\mu} \times \mathbf{B}. \quad (2.1)$$

More generally, in a system where the energy depends on an external magnetic field, we can define the magnetic moment as the coefficient of the term in the Hamiltonian linear in \mathbf{B} . For an isolated magnetic moment, the Hamiltonian would thus read

$$H = -\boldsymbol{\mu} \cdot \mathbf{B}, \quad (2.2)$$

which implies the magnetic moment minimizes its energy by aligning with the external magnetic field. We will now discuss how this magnetic moment is related to angular momentum first in a single particle, then in an isolated ion, and finally in molecules or solid materials.

2.1.1 Magnetic moment of a single particle

To motivate the connection between the a magnetic moment and angular momentum, let us first recap how magnetic dipole moments may arise in a classical picture. In classical electromagnetism, the origin of magnetic moments is always an electric current I . A current loop with area vector \mathbf{A} in a uniform magnetic field \mathbf{B} experiences a torque due to the Lorentz force on the moving charges. By integrating the torque around the loop one finds

$$\boldsymbol{\tau} = IA \times \mathbf{B} \implies \boldsymbol{\mu} = IA, \quad (2.3)$$

so the magnetic moment is proportional to the enclosed current and area. For a single electron ($q = -e$, mass $m = m_e$) moving in a circular orbit of radius r with velocity v , the current is $I = qv/2\pi r$, and the angular momentum is $L = mr v$. Both point along the normal of the orbit, leading to the relation

$$\boldsymbol{\mu}_{\text{orbital}} = \frac{q}{2m} \mathbf{L}. \quad (2.4)$$

We see that the orbital magnetic moment is directly proportional to the angular momentum \mathbf{L} . The natural unit of this orbital moment is the Bohr magneton,

$$\mu_B = \frac{e}{2m_e} \hbar, \quad (2.5)$$

which sets the scale of electronic magnetic moments (from now on, we again set $\hbar = 1$).

The classical picture of an electron orbiting a nucleus is of course not true. And even if it were, the idea that magnetism in matter can be described by the statistical mechanics of microscopic currents turns out to be wrong: the Bohr–van Leeuwen theorem shows that statistical mechanics of microscopic currents alone cannot produce macroscopic magnetism—not even paramagnetism—as their contributions cancel out in thermal equilibrium- [1]. This means we inevitably need to consider a quantum theory. Luckily for us, and maybe somewhat magically, the fact that magnetic moments are proportional to the angular momentum still holds true in quantum mechanics. This angular momentum however, is described by an operator \mathbf{L} whose components obey the same commutation relations as the spin-operators introduced in Eq. (1.6). We call the associated orbital angular momentum quantum number l , which is always an integer, and the associated magnetic quantum number $m_l = -l, \dots, l$.

The magnetic moment of a quantum particle has not only an orbital contribution, which depends on the electron's state, but also an intrinsic one: the *spin* \mathbf{S} . We refer to the spin

quantum numbers by S and m_s . The magnetic moment associated to the spin has a different proportionality constant depending on the particle, which is quantified by the g -factor defined via

$$\boldsymbol{\mu}_{\text{spin}} = g \frac{q}{2m} \boldsymbol{S}. \quad (2.6)$$

The g -factor for the electron, for example, is approximately two ($g_e \approx 2.0023$) [1]. The total magnetic moment is the sum of orbital and spin parts given by

$$\boldsymbol{\mu} = \boldsymbol{\mu}_{\text{orbital}} + \boldsymbol{\mu}_{\text{spin}} = \frac{q}{2m} (\boldsymbol{L} + g\boldsymbol{S}). \quad (2.7)$$

In atoms, orbital and spin angular momenta are coupled by the relativistic *spin-orbit interaction* (SOI) of the form

$$H_{\text{SOI}} = \lambda \boldsymbol{L} \cdot \boldsymbol{S}, \quad (2.8)$$

where λ determines the strength of the SOI and depends on the ions. A simple way to motivate this term is to view the problem in the electron's rest frame: the positively charged nucleus appears to orbit the electron, producing a magnetic field proportional to its orbital motion. This field interacts with the electron's spin magnetic moment, leading to an effective coupling between \boldsymbol{L} and \boldsymbol{S} . This relativistic effect grows rapidly with atomic number, roughly as $\lambda \sim Z^4$, since heavier nuclei bind electrons closer and at higher velocities.

With spin-orbit coupling present, S^z and L^z no longer commute with the Hamiltonian, so m_s and m_l cease to be exact quantum numbers (though for light atoms they remain approximately conserved). Instead, the relevant conserved quantity is the *total* angular momentum

$$\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}, \quad (2.9)$$

with quantum numbers J and m_J . It is therefore convenient to express the magnetic moment in terms of \boldsymbol{J} as

$$\boldsymbol{\mu} = g_J \frac{q}{2m} \boldsymbol{J}, \quad (2.10)$$

where g_J is the Landé g -factor. Its explicit form depends on the relative sizes of \boldsymbol{L} and \boldsymbol{S} and can be found in standard references [1]. For our purposes, the key point is that the magnetic moment is proportional to the total angular momentum, with the proportionality set by a material-specific g -factor. As we will elaborate later, the g -factor of an ion in a real material can also depend on the spatial direction due to an anisotropic environment (other neighboring ions), in which case it is promoted to a 3×3 matrix \boldsymbol{g} called *g-tensor*.

2.1.2 Magnetic moment of an isolated ion

So far our discussion has focused on a single particle. In real Mott insulating materials, however, magnetism originates from electrons bound to ions, made of protons and neutrons that also carry magnetic moments. Fortunately, protons and neutrons are around 2000 times heavier than electrons, and as the magnetic moment is inversely proportional to the mass of the particle we can neglect their contribution for our purposes.

For an isolated atom, Eq. (2.10) still applies, but now \boldsymbol{L} , \boldsymbol{S} , and \boldsymbol{J} denote the total angular momenta of all electrons in the atom. Determining these exactly for the ground state of the ion requires solving the many-electron problem, but very useful intuition comes from the shell and orbital picture of atomic physics. Here, electron-electron interactions are approximated by a mean-field potential, and the resulting single-particle states are grouped into shells labeled by the principal quantum number n and subshells by orbital angular momentum $l = 0, 1, 2, 3, \dots$

(denoted s, p, d, f, \dots) as familiar from the solution of the Hydrogen atom. Each orbital can host two electrons with opposite spin, and the orbital states are “filled up” by the electrons in an order that minimizes the energy of the atom. For many cases, the filling of shells follows the “Aufbau principle” (fill up shells with lower $n + l$ first, or lower n for equal $n + l$), and the filling of the subshells/orbitals follows Hund’s rules. There are, however, many exceptions to both rules and in the end the actual electron configuration for a certain ion needs to be determined by experiment.

A key takeaway is that only *partially filled* shells contribute to the magnetic moment. In filled shells, spin and orbital contributions cancel pairwise as for every $+m_l$ or $+m_s$ electron there is also an electron with $-m_l$ and $-m_s$. In partially filled shells, on the other hand, electrons tend to maximize their total spin S (Hund’s first rule). This can be understood as follows: electrons repel each other via the Coulomb interaction and gain energy when farther apart. Because the Pauli principle allows overlap only for opposite spins, electrons with the same spin tend to avoid each other more, reducing their Coulomb energy. The values for the orbital and total angular momenta L and J are approximately determined by the second Hund’s rule (after maximizing S , maximize L), and third Hund’s rule ($J = |L - S|$ if the shell is less than half full and $J = |L + S|$ if it is more than half full). The magnitude of the magnetic moment μ (defined as the eigenvalue of μ^2) of the ion is then given by the effective J of the partially filled shells

$$\mu = g_J \mu_B \sqrt{J(J+1)}. \quad (2.11)$$

That is why, when discussing materials, we usually refer only to the partially filled orbitals—for example, calling them “3d transition metals” or “4f rare-earth compounds”. This nomenclature indicates that the magnetism arises from electrons in the corresponding partially filled 3d or 4f shells.

From the periodic table alone it follows that most elements have unpaired electrons and thus finite magnetic moments, i.e. they are paramagnetic in isolation. In practice, however, such ions are rarely stable in their neutral state: unpaired electrons are highly reactive and tend to form bonds with neighboring atoms. This leads to the formation of molecules or solids, where the situation becomes more complex.

2.1.3 Magnetic ions in a crystal field

In Mott insulators, valence electrons in partially filled shells remain localized on the ions, giving rise to local magnetic moments. Examples include 3d transition metals such as Fe^{2+} in FeO or Ni^{2+} in NiO , as well as rare-earth ions like Dy^{3+} and Ho^{3+} in the spin-ice compounds $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Ho}_2\text{Ti}_2\text{O}_7$ discussed in the previous chapter. These localized electrons are not truly isolated: they can interact with electrons on neighboring ions (an effect we set aside until the next section), and they also feel the electrostatic potential of the surrounding *nonmagnetic ions* (ligands), such as oxygen.

In the orbital picture, the magnetic ion donates electrons to the ligands, leaving behind a positively charged ion embedded in a negatively charged *crystal field*. A prototypical case is a transition-metal oxide where a 3d ion sits at the center of an octahedron of oxygen ligands, as in NiO or LaMnO_3 and illustrated in Fig. 2.1(a). The five d orbitals, which are degenerate in the isolated atom, are then split by the crystal field into three t_{2g} and two e_g orbitals. Intuitively, this arises because the spatial extend of some orbitals overlap strongly with the ligands and are pushed to higher energy, while others are oriented further away, as is also illustrated in Fig. 2.1. The same principle applies to other ligand geometries, though with different splittings.

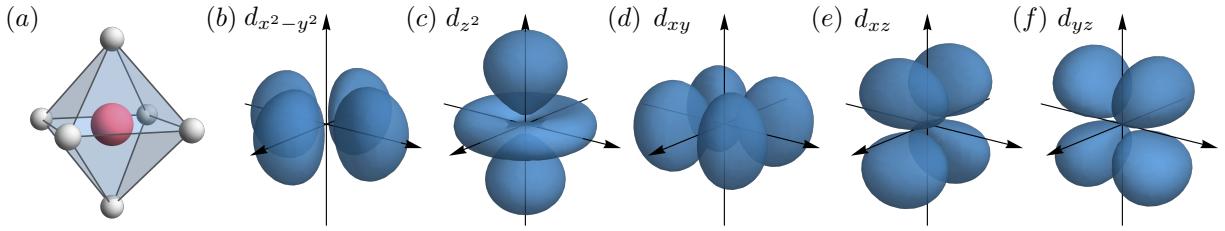


Figure 2.1 – Crystal field splitting of atomic 3d orbitals. (a) Central ion (red) in an octahedral ligand environment (white). (b)–(f) Shapes of the five 3d orbitals, visualized by regions where their wave functions $|\Psi|^2$ exceeds a threshold. The two e_g orbitals (b, c) point directly toward the ligands, while the three t_{2g} orbitals (d-f) lie between them. As a result, the electrostatic crystal field lifts their degeneracy.

The way electrons occupy these orbitals depends on the specific material. Crucially, crystal fields can strongly alter the angular momenta predicted by Hund’s rules. In many 3d compounds, for example, the orbital contribution is quenched ($L = 0$), leaving the spin S as the dominant magnetic degree of freedom. This explains why, in most materials with weak spin–orbit coupling, magnetism is governed primarily by spin.

In contrast, for rare-earth ions the spin–orbit interaction is strong, and local moments with different effective total angular momentum J are realized. The anisotropic crystal field can then cause large anisotropies in the magnetic moments. In spin-ice materials, for example, this results in localized moments that effectively behave like Ising variables constrained to point along local \hat{z} directions. Additionally, the response to a magnetic field—that is, the magnetic moment—is then also anisotropic. This is captured by generalizing the g_J -factor from Eq. (2.10) to a site-dependent 3×3 matrix \mathbf{g} , known as the g -tensor, such that the magnetic moments of a magnetic ion i is given by

$$\boldsymbol{\mu}_i = \mu_B \mathbf{g}_i \cdot \mathbf{J}. \quad (2.12)$$

When the g -tensor is anisotropic, the magnetic moment is no longer parallel to the total angular momentum \mathbf{J} but aligns preferentially along directions set by \mathbf{g} . In a local basis where \mathbf{g} is diagonal, it typically takes the form

$$\mathbf{g}_{\text{local}} = \begin{pmatrix} g_{\perp} & 0 & 0 \\ 0 & g_{\perp} & 0 \\ 0 & 0 & g_{\parallel} \end{pmatrix}. \quad (2.13)$$

Here $g_{\parallel} \gg g_{\perp}$ favors moments aligned along the local \hat{z} axes, corresponding to easy-axis anisotropy as in spin-ice compounds. Conversely, $g_{\parallel} \ll g_{\perp}$ confines the moments to the local $\hat{x}\hat{y}$ plane, giving rise to easy-plane anisotropy.

In summary, predicting the ground state and magnetic moment of an ion in a solid is highly nontrivial. A useful starting point is to identify localized electrons in partially filled shells. Incorporating crystal-field effects and spin–orbit coupling then provides an estimate for the total angular momentum \mathbf{J} of the unpaired electrons, and thus for the effective magnetic degrees of freedom that govern the material’s magnetism. Deriving this fully from first principles is rarely feasible, so the actual behavior must ultimately be determined experimentally. To move beyond the simple paramagnetic response of independent local moments, it is essential to also account for interactions between magnetic ions, which we discuss in the next section.

2.2 Localized magnetic moments in solids and the exchange interaction

We have seen how magnetic moments arise from electrons localized on individual ions and how they are modified by the crystal field of surrounding ligands. So far, however, we have ignored interactions with other magnetic ions and their electrons. In solids this is a crude approximation: when ions bind to form a molecule or a crystal, the atomic-orbital picture breaks down at least to some degree, since electrons are influenced both by the potentials of nearby nuclei and by Coulomb repulsion from other electrons. These effects drive two competing tendencies.

On the one hand, proximity to neighboring nuclei favors delocalization: electrons can lower their energy by spreading over several ions—or even across the entire crystal. This is analogous to a particle in a box, where the (purely kinetic) energy $E \sim 1/L^2$ decreases as the accessible length L increases. On the other hand, minimizing Coulomb repulsion drives electrons to stay as far apart as possible, favoring localization: if one electron is already close to an ion, a second pays an energy penalty for occupying the same region. Strong electronic correlations can therefore stabilize localized electrons and are the main reason local magnetic moments can exist in a solid at all.

In real materials both tendencies are present, and the electronic properties are set by their competition. When correlations are weak, electrons delocalize as in metals or band insulators. When correlations are strong, electrons remain localized, giving rise to Mott insulators. Remarkably, in the latter case the dominant interaction between localized moments—the *exchange interaction*—emerges precisely from the interplay of delocalization, Coulomb repulsion and the exchange statistics of fermions.

In this section we first discuss the limit of weak electronic correlations, described by a non-interacting tight-binding Hamiltonian. We then show how adding strong correlations drives electron localization and simultaneously generates the exchange interaction, thereby motivating the quantum spin models central to this thesis. Afterwards, we examine how different microscopic ingredients can produce distinct types of exchange, such as the anisotropic interactions in Kitaev materials and the rare-earth oxides studied in Chapter 5. Finally, we briefly outline how SU(4) models arise in an analogous way. To connect back to the single-ion orbital picture, we adopt a ground-up derivation in second quantization, closely following Ref. [70].

2.2.1 Weak electronic correlations and the tight-binding model

Let us start in the limit of very weak electronic interactions, but with electrons feeling the full periodic potential V of the ions in the crystal lattice. The corresponding Hamiltonian in second quantized form is

$$H_e = \sum_s \int d^d r c_s^\dagger(\mathbf{r}) \left[\frac{\mathbf{p}^2}{2m} + V(\mathbf{r}) \right] c_s(\mathbf{r}), \quad (2.14)$$

where $c_s^\dagger(\mathbf{r})$ and $c_s(\mathbf{r})$ are fermionic field operators creating/annihilating a fermion with spin $s = \pm 1/2$ perfectly localized at \mathbf{r} . To obtain a lattice model, we label the electronic states by the unit cell index $i = 1, \dots, N$ (with $N \rightarrow \infty$ for an infinite lattice or periodic boundary conditions) and by a multi-index α that encodes spin, orbital, and, if present, sublattice degrees of freedom. In this notation, the fermionic field operators can be expanded in the atomic orbital basis

$$c^\dagger(\mathbf{r})_s = \sum_{i\alpha} \varphi_{i\alpha}(\mathbf{r}) c_{i\alpha}^\dagger, \quad c(\mathbf{r})_s = \sum_{i\alpha} \varphi_{i\alpha}^*(\mathbf{r}) c_{i\alpha}, \quad (2.15)$$

where $\varphi_{i\alpha}$ is the wave function of the orbital localized at the unit cell i and $c_{i\alpha}^\dagger$ creates an electron in this orbital. Plugging this into the noninteracting Hamiltonian transforms it into the familiar

tight-binding form

$$H_{\text{TB}} = \sum_{i,j,\alpha,\beta} t_{ij}^{\alpha\beta} c_{i\alpha}^\dagger c_{j\beta}, \quad (2.16)$$

with the hopping elements

$$t_{ij}^{\alpha\beta} = \langle \varphi_{i\alpha} | H_e | \varphi_{j\beta} \rangle = \int d^d r \varphi_{i\alpha}^*(\mathbf{r}) [\epsilon_\beta + \Delta V_i(\mathbf{r})] \varphi_{j\beta}(\mathbf{r}), \quad (2.17)$$

where ϵ_β is the energy of the atomic orbital β and ΔV_i is the correction to the atomic potential at site i due to neighboring ions. This Hamiltonian describes the hopping of electrons from sites i and orbital α to site j with orbital β , which is mainly controlled by the effective overlap of the corresponding orbitals. The quadratic Hamiltonian can be readily diagonalized by a Fourier transformation

$$c_{i\alpha} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k} \in 1.\text{BZ}} e^{i\mathbf{r}_i \cdot \mathbf{k}} c_{\mathbf{k}\alpha} \quad c_{i\alpha}^\dagger = \frac{1}{\sqrt{N}} \sum_{\mathbf{k} \in 1.\text{BZ}} e^{-i\mathbf{r}_i \cdot \mathbf{k}} c_{\mathbf{k}\alpha}^\dagger, \quad (2.18)$$

where $\sum_{\mathbf{k} \in 1.\text{BZ}}$ goes over the first Brillouin zone (and should be considered as an integral for the infinite lattice). Utilizing the translational invariance of the hopping matrix $t_{ij}^{\alpha\beta} = t_{\mathbf{R}_i - \mathbf{R}_j}^{\alpha\beta}$ this yields the diagonal Hamiltonian

$$H_{\text{TB}} = \sum_{\mathbf{k} \in 1.\text{BZ}} \sum_m \epsilon_m(\mathbf{k}) c_{\mathbf{k}m}^\dagger c_{\mathbf{k}m}, \quad (2.19)$$

where the electronic bands $\epsilon_m(\mathbf{k})$ are the eigenvalues of the hopping matrix in momentum space $t^{\alpha\beta}(\mathbf{k}) = \sum_i t_{R_i}^{\alpha\beta} e^{-i\mathbf{R}_i \cdot \mathbf{k}}$. The eigenstates of this Hamiltonian are $|\psi\rangle_{\mathbf{k}m} = c_{\mathbf{k}m}^\dagger |0\rangle$, which are Bloch states made from linear combination of the atomic orbitals (LCAOs), where the precise combination is derived from the eigenvectors of $t^{\alpha\beta}(\mathbf{k})$. Filling these bands up to the Fermi energy ϵ_F distinguishes metals (partially filled bands) from band insulators (all bands full or empty).

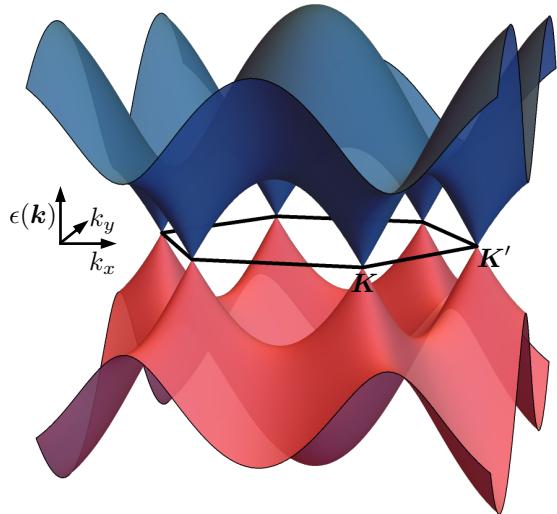
The number of bands m equals the number of orbitals included in the calculation, which in principle is infinite. A tight-binding description in terms of atomic orbitals is therefore only practical if most orbitals have very weak overlap, so that only a small subset contributes appreciably near the Fermi energy. In practice, one typically treats only the valence electrons of the outermost shells explicitly, since the core orbitals lie at very low energies, are tightly bound to the nucleus, and have negligible overlap with other orbitals. Their influence on the valence electrons can then be incorporated effectively through a screened nuclear potential.

As a concrete example—and one that will be relevant again when we discuss moiré materials later—consider graphene, a single atomic layer of graphite in which carbon atoms form a honeycomb lattice [as visualized in Fig. 1.2(c)]. A neutral carbon atom has six electrons with configuration $1s^2 2s^2 2p^2$. Only the valence $2s^2$ and $2p^2$ electrons contribute significantly to bonding. Of these, three sp^2 -hybridized orbitals (combinations of s and p orbitals) form fully filled σ -bands, while the remaining $2p^z$ orbital gives rise to two partially filled π and π^* bands. Since the π orbitals have little overlap with the filled σ bands, the low-energy physics of graphene is well captured by a single-orbital tight-binding model [71] with nearest-neighbor hopping of the form

$$H = -t \sum_{\langle ij \rangle s} (c_{is}^\dagger c_{js} + h.c.) = -t \sum_{\mathbf{k}} \begin{pmatrix} c_{kA}^\dagger & c_{kB}^\dagger \end{pmatrix} \begin{pmatrix} 0 & f(\mathbf{k}) \\ f^*(\mathbf{k}) & 0 \end{pmatrix} \begin{pmatrix} c_{kA} \\ c_{kB} \end{pmatrix} \quad f(\mathbf{k}) = \sum_{j=1}^3 e^{i\mathbf{k} \cdot \boldsymbol{\delta}_j}, \quad (2.20)$$

Figure 2.2 –

Band structure of the nearest-neighbor tight-binding model on the honeycomb lattice. The plot shows the electronic dispersion $\epsilon(\mathbf{k})$ obtained from the nearest-neighbor tight-binding Hamiltonian on the honeycomb lattice as relevant for graphene. The hexagon outlines the first Brillouin zone. At its corners—the inequivalent \mathbf{K} and \mathbf{K}' points, in this context often referred to as “Dirac points”—the two bands touch and the dispersion becomes linear. This linear crossing underlies many of the remarkable electronic properties of graphene at low energies.



where $h.c.$ denotes the Hermitian conjugate and the sum goes over the honeycomb lattice sites, with the nearest-neighbor distance a . In the second step we applied the Fourier-transform, A/B indices label the sublattice, and δ_j are the three nearest-neighbor vectors of sublattice A. Diagonalization leads to the band structure

$$\epsilon_{\pm}(\mathbf{k}) = \pm|f(\mathbf{k})| = \pm t \sqrt{3 + 2 \cos(\sqrt{3}k_x a) + 4 \cos\left(\frac{\sqrt{3}}{2}k_x a\right) \cos\left(\frac{3}{2}k_y a\right)} \quad (2.21)$$

shown in Fig. 2.2. While this model does not reproduce graphene’s full realistic band structure, it captures the essential linear dispersion at the inequivalent \mathbf{K} and \mathbf{K}' points of the Brillouin zone, in this context also called Dirac points, which govern most of its fascinating properties at low energies. Descriptions of graphene thus often involve an expansion of the Hamiltonian around those points. Because the states near \mathbf{K} and \mathbf{K}' are degenerate, this expansion introduces an additional bi-valued quantum number—the *valley degree of freedom*—which labels whether an electron resides near \mathbf{K} or \mathbf{K}' . As mentioned earlier, this degree of freedom plays a central role in strong-coupling descriptions of moiré materials by $SU(4)$ spin models, which we will revisit at the end of this chapter.

Another key takeaway from this calculations is that the band width W is proportional to the hopping matrix element $W \sim t$, which is a general feature. *Flat electronic bands* are therefore a good indicator that the energy gain from delocalization is rather weak, which makes the effects of electron-electron interactions much more important.

For a more realistic reproduction of band structures within the simplest possible tight-binding models, atomic orbitals are often not the optimal basis. Strong overlap may require many orbitals and further-neighbor hoppings, and electronic correlations can complicate matters further. Instead, first-principles or *ab initio* methods (such as density functional theory) are commonly used to compute approximate single-particle Bloch states, which already include electron-electron interactions at a mean-field level [72]. From these, one can construct Wannier functions via

$$|\psi_{im}\rangle = \frac{1}{\sqrt{N}} \sum_{\mathbf{k} \in \text{B.Z.}} e^{-i\mathbf{k} \cdot \mathbf{R}_i} |\psi_{\mathbf{k}m}\rangle \quad , \quad (2.22)$$

where a gauge freedom in the Bloch states is tuned to yield *maximally localized Wannier functions* fitted to the bands of interest [73]. These Wannier functions are strongly localized around the

ions and may or may not resemble atomic or molecular orbitals, depending on the strength of orbital overlap.¹ They provide a natural bridge between first-principles calculations and effective low-energy Hamiltonians, and thus can serve as the starting point for deriving interacting models.

2.2.2 Strong electronic correlations and the Hubbard model

The main strength of the tight-binding description is that it yields simple, exactly solvable models providing a qualitative descriptions of metals and band insulators. The magnetic materials of interest here, however, are *Mott insulators*, which lie beyond this picture: although their band structure suggests metallic behavior due to partially filled bands, strong electronic correlations open a gap and yield insulating states with localized unpaired electrons on the magnetic ions.

To describe this, the tight-binding picture remains a natural starting point. Consider again a single-orbital system (as motivated for graphene), but now include electron–electron interactions. In second quantization the Coulomb interaction reads [70]

$$H_{ee} = \frac{1}{2} \sum_{s,s'} \int d^d r \int d^d r' V_{ee}(\mathbf{r} - \mathbf{r}') c_s^\dagger(\mathbf{r}) c_{s'}^\dagger(\mathbf{r}') c_{s'}(\mathbf{r}) c_s(\mathbf{r}). \quad (2.23)$$

In its simplest form $V_{ee}(\mathbf{r} - \mathbf{r}')$ is the bare coulomb potential $V(\mathbf{r} - \mathbf{r}') \sim \frac{e^2}{|\mathbf{r} - \mathbf{r}'|}$, but more accurately one uses a screened version from first principle calculations [74] that takes into account the partial blocking of the potential due to electrons and orbitals that are not specifically treated in the effective (in this case one-band) model.

To obtain a lattice model, we expand the electronic field operators in the atomic-orbital (or Wannier) basis defined in Eq. 2.15, which yields

$$H_{ee} = \sum_{ss'} \sum_{ii'jj'} U_{ii'jj'} c_{is}^\dagger c_{i's'}^\dagger c_{j's'} c_{js} \quad (2.24)$$

with

$$U_{ii'jj'} = \frac{1}{2} \int d^d r \int d^d r' \varphi_i^*(\mathbf{r}) \varphi_j(\mathbf{r}) V_{ee}(\mathbf{r} - \mathbf{r}') \varphi_{i'}^*(\mathbf{r}') \varphi_{j'}(\mathbf{r}'). \quad (2.25)$$

In the limit where the overlap of neighboring orbitals is weak, the most important term is the on-site interaction $U_{iiii} \equiv U/2$. In this case, the electronic behavior is well described by the (one-band) Hubbard model

$$H = H_{TB} + H_U = -t \sum_{\langle ij \rangle s} (c_{is}^\dagger c_{js} + h.c.) + \sum_i U n_{i\uparrow} n_{i\downarrow}, \quad (2.26)$$

where $n_{is} = c_{is}^\dagger c_{is}$ is the density operator counting how many electrons of spin s are at site i .

In the limit of large electronic correlations $U \gg t$, the ground state is obtained by minimizing $n_{i\uparrow} n_{i\downarrow}$ which counts the amount of doubly occupied sites. At half-filling (on *average* one electron per site), this results in exactly one electron localized at each site. Any hopping creates a doubly occupied site at an energy cost U , becoming favorable only if $t \gtrsim U$. The large- U limit thus realizes a Mott insulating ground state. Whether a material falls into this regime is set by the ratio U/t , which is controlled primarily by how localized the relevant orbitals are. Most conventional Mott insulators are therefore found among $3d$ transition-metal compounds—including the cuprates, famous for high-temperature superconductivity—and among $4f$ rare-earth oxides, both characterized by strongly localized orbitals.

¹The existence of exponentially localized Wannier states is guaranteed for time-reversal symmetric insulators. If time-reversal symmetry is broken, bands may acquire a finite Chern number, in which case such states do not exist and no conventional tight-binding description is possible [73].

2.2.3 From the Hubbard to the Heisenberg model and the exchange interaction

We have seen how electronic correlations can give rise to localized moments. To explain magnetic effects beyond paramagnetism, however, we also need to consider the effective interaction between those moments. A natural first guess for the origin of such interactions is the magnetic dipole–dipole coupling known from classical electrodynamics: A magnetic moment μ_i produces a magnetic field that interacts with local magnetic moments μ_j in its vicinity. Indeed, such an interaction can be observed in magnetic materials, but it is way to weak to account for many magnetic phenomena: if two magnetic dipoles are separated by the vector \mathbf{r} the corresponding potential energy is [1]

$$E = \frac{\mu_0}{4\pi r^3} \left[\mu_i \cdot \mu_j - \frac{3}{r^2} (\mu_i \cdot \mathbf{r})(\mu_j \cdot \mathbf{r}) \right], \quad (2.27)$$

where μ_0 is the vacuum permeability. With an ion's local magnetic moment at the order of $\mu \sim \mu_B$ and a typical distance between neighboring ions of $r \approx 1 \text{ \AA}$, the energy-scale of the dipolar interaction is around 0.6 meV, which via $E = k_B T$ corresponds to a temperature of around 1 K. There are, however, many materials that magnetically order at much larger temperatures (otherwise there would be no permanent magnets at room temperature). Thus, the dominant interaction cannot always be dipolar. Instead, in most Mott insulators it arises from the competition between strong electronic correlations and the kinetic energy gained through delocalization in the ionic potentials—captured by the parameters U and t in the Hubbard model—together with the exchange statistics of the fermionic electrons.

The resulting *exchange interaction* can, in its simplest form, be derived directly from the Hubbard model in the large- U/t limit by expanding to second order in t/U , which we will now show. We mainly follow an approach via conventional perturbation theory described in Ref. [75]. An alternative approach based on a Schrieffer–Wolff transformation is presented in Ref. [70].

We again consider the single-orbital Hubbard model at half-filling. In the limit $U \gg t$, we treat H_{TB} as a perturbation to H_U , whose eigenenergies depend only on the number of doubly occupied sites n . Each sector with fixed n is highly degenerate. We label its states by $|nk\rangle$ with energies $E_n = nU$, where k indexes the particular configuration of spins and occupations across the lattice. The ground-state subspace consists of all singly occupied configurations $\{|0k\rangle\}$, and our aim is to derive an effective Hamiltonian H_{eff} within this subspace, accurate to second order in t . The corresponding matrix element can be obtained from degenerate perturbation theory as²

$$\langle 0k | H_{\text{eff}} | 0k' \rangle = \langle 0k | H_0 | 0k' \rangle + \langle 0k | H_{\text{TB}} | 0k' \rangle + \sum_{p>0,l} \frac{\langle 0k | H_{\text{TB}} | pl \rangle \langle pl | H_{\text{TB}} | 0k' \rangle}{E_0 - E_p} + O(t^3). \quad (2.28)$$

The first term is zero because $E_0 = 0$, the second term is zero because $H_{\text{TB}} |0k'\rangle$ will always produce a state with one doubly occupied site orthogonal to $|0k\rangle$. The same reasoning shows that only $p = 1$ terms contribute to the sum in the third term and we can therefore replace $E_p = E_1 = U$ and remove the restriction $p > 0$ from the sum. Identifying $\mathbf{1} = \sum_{pl} |pl\rangle \langle pl|$, the effective Hamiltonian becomes

$$H_{\text{eff}} = -\frac{1}{U} P_0 H_{\text{TB}}^2 P_0 = -\frac{t^2}{U} P_0 \sum_{\langle ij \rangle, s} \sum_{\langle i'j' \rangle s'} P_0 c_{is}^\dagger c_{js} c_{i's'}^\dagger c_{j's'} P_0 \quad (2.29)$$

²To connect the effective Hamiltonian framework to conventional (degenerate) perturbation theory learned in quantum mechanics courses see e.g. Ref. [76]

where $P_0 = \sum_k |0k\rangle \langle 0k|$ projects onto the singly occupied subspace. This Hamiltonian describes virtual hopping processes from site j' to i' and j to i . Since the resulting state again has to lie in the singly occupied subspace, the electrons have to hop to a neighboring site and back again, implying $i' = j$ and $j' = i$. Restricting to these terms and reordering the operators gives

$$H_{\text{eff}} = -\frac{t^2}{U} \sum_{\langle ij \rangle} \sum_{ss'} P_0 \left(c_{is}^\dagger c_{is'} \delta_{ss'} - c_{is}^\dagger c_{is'} c_{js'}^\dagger c_{js} \right) P_0 \quad (2.30)$$

Although not obvious at first sight, this effective Hamiltonian is—up to an additive constant—nothing but the Heisenberg model introduced in Chapter 1. To see this, we can express fermion bilinears in terms of spin operators using the Abrikosov fermion representation from Eq. (1.22). Since the density operator $n_i = \sum_s c_{is}^\dagger c_{is}$ is fixed to one within the singly occupied subspace, it acts as a constant proportional to the identity. Because Pauli matrices together with the identity span the space of 2×2 Hermitian matrices, any local bilinear $O_i = \sum_{ss'} A_{ss'} c_{is}^\dagger c_{is}$ (with \mathbf{A} Hermitian) can be written as a combination of spin and density operators. Products $O_i O_j$ can thus be re-expressed in terms of spin–spin interactions. In the present case, this mapping leads directly to the Heisenberg Hamiltonian (as can be verified explicitly)

$$H_{\text{eff}} = \frac{2t^2}{U} \sum_{\langle ij \rangle} (\mathbf{S}_i \cdot \mathbf{S}_j - \frac{1}{4}) = J \sum_{\langle ij \rangle} (\mathbf{S}_i \cdot \mathbf{S}_j) + \text{const.}, \quad (2.31)$$

with strictly positive (antiferromagnetic) interaction $J \sim \frac{t^2}{U}$.

Although the derivation may appear technical, it reveals a simple physical picture of the exchange interaction: Electrons can reduce their energy by the virtual process of hopping to a different site and back. Due to the Pauli principle, only electrons with opposite spins can contribute to this process, and therefore have an energetic advantage—this favors an antiferromagnetic alignment. This mechanism is often referred to as *kinetic exchange*. In many materials, the hopping matrix element t_{ij} is not a direct overlap between the magnetic orbitals, but is mediated by intermediate nonmagnetic ions—for instance, oxygen ligands in transition-metal or rare-earth oxides. In this case, the interaction is known as *superexchange*. In Mott insulators with tightly bound orbitals and large U , superexchange typically dominates, leading to predominantly antiferromagnetic couplings, as indeed observed in most $3d$ transition-metal compounds.

If there is a significant orbital overlap between neighboring sites, then longer-range terms of the coulomb integral $U_{ii'jj'}$ need to be taken into account. If we only consider nearest neighbor sites i and j , these include the term U_{ijij} , which leads to interactions of the form $\sum_{i \neq j} V_{ij} n_i n_j$. This couples electron densities at neighboring sites, which in principle can lead to order in the charge distribution called *charge density waves*. We will always consider the strong U limit, where one electron is fixed to each site, for which this term is a constant.

The second nearest neighbor term is U_{ijji} . Projecting onto the singly-occupied subspace and using the representation of spin operators via fermions as above yields

$$H_F = \sum_{\langle ij \rangle} J_{ij}^F \mathbf{S}_i \mathbf{S}_j + \text{const.}, \quad (2.32)$$

where $J_F \sim U_{ijji}$ is usually *positive*. This interaction is called *direct exchange* and typically induces a ferromagnetic coupling between neighboring sites that is weaker than the kinetic exchange.

2.2.4 Spin-orbit entangled Mott insulators and anisotropic exchange

We showed how the simple one-orbital Hubbard model can explain how Heisenberg type interactions arise in Mott insulating materials. These interactions are $SU(2)$ symmetric—multiplication $\mathbf{S}_i \rightarrow \mathbf{U}^\dagger \mathbf{S}_i \mathbf{U}$ by a matrix $\mathbf{U} = \exp \sum_a \lambda^a \mathbf{S}^a \in SU(2)$ leaves them invariant. As already mentioned in the introductory chapters, in many of the spin liquid candidate materials, interactions that explicitly break the $SU(2)$ symmetry and that depend on the spatial direction of the bond are important.

Such interactions arise if we take into account spin-orbit coupling, in which case the spin couples to the orbital angular momentum, which in turn couples directly to the crystal field of the neighboring ions. The crystal field is inherently anisotropic and explicitly breaks rotational symmetry. As discussed in Sec. 2.1, this lifts the degeneracy between orbitals with different spatial orientation. When spin-orbit coupling is strong and the orbital angular momentum is not quenched (hence $L \neq 0$), spin and orbital degrees of freedom combine into a total angular momentum \mathbf{J} . The resulting magnetic moments and their mutual interactions can then be highly anisotropic.

Already in 1960 Toru Moriya, based on the work of Igor Dzyaloshinskii, showed that including a spin-orbit coupling term $\sim \lambda \mathbf{S} \cdot \mathbf{L}$ in the Hamiltonian and derivation of the super-exchange interaction leads to asymmetric interactions $\sim D(S_i^x S_j^y - S_i^y S_j^x)$ later dubbed Dzyaloshinskii-Moriya (DM) interactions. These may exist even for weak spin-orbit coupling. For materials with strong spin-orbit coupling—such as heavier $4d$ and $5d$ transition-metal compounds (like many Iridates or Ruthenates) and $4f$ rare-earth systems (like the spin ice compounds $\text{Ho}_2\text{Ti}_2\text{O}_7$ and $\text{Dy}_2\text{Ti}_2\text{O}_7$)—an even wider variety of interactions can arise.

To derive the super-exchange interaction for these systems, one has to take into account hopping processes between different orbitals $t_{ij}^{\alpha\beta}$ mediated by the Ligand environment, on-site direct exchange interactions between different orbitals (also called Hund's interaction) and the local spin-orbit energy $\sim \lambda \mathbf{S} \cdot \mathbf{L}$. In the strong-coupling limit, the resulting interactions can then be described by Hamiltonians of the form

$$H = \sum_{ij} \mathbf{S}_i^T \mathbf{J}_{ij} \mathbf{S}_j \quad (2.33)$$

with a fully nondiagonal 3×3 coupling matrix of the form

$$\mathbf{J}_{ij} = \begin{pmatrix} J + K & D + \Gamma' & D + \Gamma' \\ -D + \Gamma' & J & \Gamma \\ -D + \Gamma' & \Gamma & J \end{pmatrix}, \quad (2.34)$$

where J denotes the Heisenberg, K the Kitaev, D the antisymmetric DM, and Γ, Γ' symmetric off-diagonal couplings. Among those, however, only terms that are compatible with the crystal symmetry are present. For example, if the crystal is symmetric under an inversion along the bond between sites i and j , this implies $\mathbf{J}_{ij} = \mathbf{J}_{ji}$ which rules out antisymmetric interactions. The sign and magnitude of the couplings also differs from compound to compound.

Most commonly, the Heisenberg interaction still remains dominant. Remarkable exceptions, however, are the so-called Kitaev materials [40]. Jackeli and Khaliullin [38, 39] showed that in transition-metal ions with a d^5 electronic configuration and an edge-sharing octahedral ligand environment, the effective local moments are spin-orbit entangled $j = 1/2$ states that couple predominantly via the anisotropic Kitaev interaction K . In this specific geometry, different super-exchange hopping paths destructively interfere, strongly suppressing the usual Heisenberg exchange. Another example are the pyrochlore rare-earth oxides discussed in Chapter 5, where all exchange couplings can be of comparable magnitude [24].

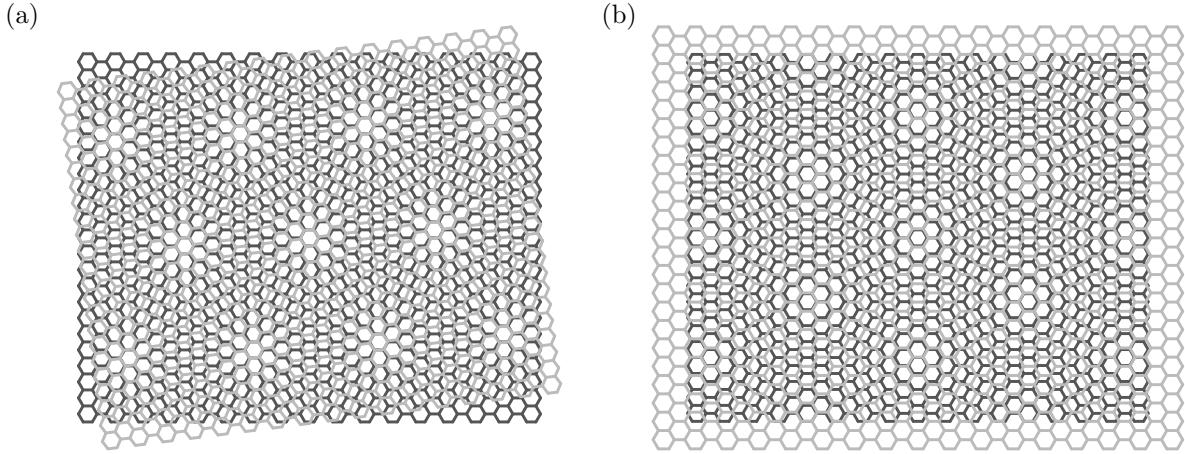


Figure 2.3 – Moiré patterns in stacked honeycomb lattices. (a) Moiré pattern arising from a small twist angle between two equivalent honeycomb layers, as realized in twisted bilayer graphene. (b) Moiré pattern formed when two honeycomb layers with slightly different lattice constants are stacked, as in heterostructures where graphene is aligned with hexagonal boron nitride (h-BN), whose lattice constant is slightly larger than that of graphene. In both cases the moiré pattern forms a triangular lattice.

2.2.5 Effective SU(4) models for spin-orbit and spin-valley entangled materials

Interestingly, including orbital degrees of freedom can also change the very nature of the effective local moments. In the context of the Jahn–Teller effect in transition-metal oxides, Kugel and Khomskii [50] showed that adding a twofold orbital degeneracy to the Hubbard model (introducing four rather than two fermion flavors) leads, in the strong-coupling limit, to an effective spin–orbital model where the physical spin couples to an orbital pseudospin τ_i (also of spin- $\frac{1}{2}$ type). More recently, for the concrete material α -YbCl₃, which exhibits strong spin–orbit coupling, analogous derivations suggest that the low-energy manifold, formed by spin–orbit entangled $j = 3/2$ moments, is well described by

$$H = J \sum_{\langle ij \rangle} (1 + \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j) (1 + \boldsymbol{\tau}_i \boldsymbol{\tau}_j \mathbf{1}) , \quad (2.35)$$

including two coupled pseudospins $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$. Here, the products of these operators on the same site should be interpreted as combined spin-orbital operators

$$\sigma_i^\mu \tau_i^\nu \equiv P_0 c_{isl}^\dagger \theta_{ss'}^\mu \theta_{ll'}^\nu c_{is'l'} P_0 , \quad (2.36)$$

where $\mu = 0, 1, 2, 3$, and we define $\boldsymbol{\theta}^0 = \mathbf{1}$. P_0 again projects two a subspace of fixed filling, in this case one electron per site. These operators are traceless and Hermitian, and thus form a basis of the Lie algebra of SU(4). The Hamiltonian above corresponds to the fully SU(4)-symmetric case, characterized by equal coupling between all generators—directly analogous to the SU(2) Heisenberg model. In real materials, however, SU(4)-breaking terms are typically also present.

Remarkably, a completely different class of materials—the so-called moiré materials—can also be described by such effective models. Moiré materials are built by stacking layers of different two-dimensional crystals, most prominently graphene. Rotating the layers even by a small angle misaligns their atomic lattices, creating a slowly varying pattern of overlapping regions, known as a moiré pattern illustrated in Fig. 2.3(a). The most famous example material realization is twisted bilayer graphene (TBG). Moiré patterns can also arise from stacking two layers with

slightly different lattice constants, as illustrated in Fig. 2.3. This is the case in trilayer graphene aligned with hexagonal boron nitride—a system we will study in detail in Chapter 7.

On its own, graphene is a semi-metal with weak electronic correlations, and thus far from a Mott insulator [71]. In moiré systems, however, a special mechanism can still drive strongly correlated behavior. This was first pointed out by Bistritzer and MacDonald in 2011 for TBG [77]. Starting from the effective low-energy tight-binding description near the Dirac points of graphene (discussed in Sec. 2.2.1 and shown in Fig. 2.2), they included interlayer tunneling between rotated layers. This coupling produces extremely narrow “flat bands” in the reduced Brillouin zone of the moiré lattice. The degree of flatness is tunable with the twist angle, and at certain so-called *magic angles* the bands become nearly dispersionless. As discussed earlier, a narrow band corresponds to a small hopping amplitude t in a Hubbard description. Even a modest electron interactions U can therefore drive strong correlation effects. Indeed, experiments on TBG and related moiré systems have revealed a wealth of strongly correlated phases, including correlated insulators and superconductivity [78].

Here, we are mainly interested in the Mott insulating states, for which the strong-coupling limit provides a natural starting point. For both TBG and TG/h-BN, extended Hubbard models have been derived that include fermions carrying both spin and valley degrees of freedom defined on the moiré lattice [48, 49]. In the limit $U/t \rightarrow \infty$, these models reduce to effective SU(4) spin–valley models. Unlike the fully symmetric SU(4) case discussed earlier, however, the interactions in moiré systems typically break this symmetry. For instance, in Chapter 7 we study an explicit model for TG/h-BN derived in Ref. [49], which contains terms of the form

$$(1 + \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j)(\tau_i^x \tau_j^x + \tau_i^y \tau_j^y), \quad (1 + \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j)(\tau_i^x \tau_j^y - \tau_i^y \tau_j^x), \quad (2.37)$$

explicitly breaking SU(4) down to $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$. Aside from a quantum paramagnetic region, we find that such interactions also induce quite novel order in the combined spin and valley degrees of freedom.

Part I

Methods

Chapter 3

Classical and mean-field approximations

The exact solution of quantum spin models quickly becomes intractable once we move beyond just a handful of interacting spins, due to the exponential growth of the Hilbert space and the complexity introduced by long-range entanglement. As argued in the previous chapters, highly frustrated systems or highly symmetric $SU(4)$ models are particularly prone to such strongly entangled states, where quantum fluctuations dominate. By contrast, in spin models intended to capture real materials, quantum paramagnetic phases are actually rare: the presence of additional interactions usually drives the system away from perfect frustration and into a magnetically ordered state.

Qualitatively, magnetically ordered states are often simpler to analyze using approaches that neglect, or strongly approximate, the underlying entanglement structure. For $SU(2)$ models, the most common strategy is to work with their classical counterparts, already introduced in Chapter 1. We begin this chapter with Sec. 3.1, which discusses the Luttinger–Tisza method—a particularly elegant and computationally inexpensive approach for finding the ground state of classical $SU(2)$ spin models, and one that we will employ repeatedly throughout this thesis. Before that, however, we briefly explain how the classical limit emerges from the quantum model by considering the large-spin limit $S \rightarrow \infty$.

The situation is more subtle for $SU(4)$ models. A naive replacement of $SU(4)$ spin operators with real classical vectors generally fails to capture the structure of the local Hilbert space, even when inter-site entanglement is neglected. In Sec. 3.2, we introduce an alternative route to define a meaningful semiclassical limit. To study $SU(4)$ models within this framework, we have developed a dedicated Monte Carlo implementation made available in the Julia package `SemiClassicalMC.jl` [C1], whose foundations we describe in detail.

Finally, another way to simplify the quantum problem is to retain entanglement only within finite clusters rather than neglecting it altogether. In Sec. 3.3 we introduce such an approach known cluster mean-field theory (CMFT). Here, a finite cluster of quantum spins is solved exactly using ED, while inter-cluster couplings are treated at the mean-field level. CMFT thus strikes a balance between the overly simplistic conventional mean-field theory and the intractable full quantum problem. For $SU(2)$ Heisenberg models, we have implemented this method in the Julia package `ClusterMeanFieldTheory.jl` [C2].

3.1 The Luttinger–Tisza method for classical spin models

A powerful method for finding the ground state of classical spin models is the Luttinger–Tisza (LT) method, sometimes called the Luttinger–Tisza–Lyons–Kaplan method. Introduced by Luttinger and Tisza [79] for dipolar interactions on an eight-site cluster, it was soon generalized to Ising models on Bravais lattices [80] and to Heisenberg models on non-Bravais lattices by Lyons and Kaplan [81]. Later extensions include treatments of certain three-dimensional noncoplanar

states by Schmidt and Richter. Comprehensive accounts can be found in Refs. [82–84], which form the basis of this section.

The LT method is most commonly formulated in momentum space, where translational invariance makes it possible to work directly with infinite lattices and to capture incommensurate ground states. Also useful, however, is the real-space formulation applied to a finite lattice. Although less frequently discussed, the real-space version can be more straightforward: when it applies, it directly yields explicit spin configurations in real space, including certain noncoplanar states that would require a multi- \mathbf{q} Ansatz in momentum space¹.

In what follows, we set the stage by clarifying how classical spins emerge from quantum spins in the limit $S \rightarrow \infty$, providing a more rigorous justification for classical spin models. We then introduce the real-space version of the LT method, followed by the more conventional momentum-space formulation. Finally, we illustrate the method with three examples, showing how collinear, coplanar, and even noncoplanar states can be captured within both approaches.

3.1.1 The classical limit for $SU(2)$ spins

Quantum spin operators are defined by their commutation relations in Eq. (1.6). Their noncommutativity underlies the quantum nature of spins, which differ from classical angular momenta in two important ways. First, two components of the quantum spins cannot be known at the same time. For example, consider the $S = 1/2$ spin-up eigenstate of the S^z operator with eigenvalue $m = 1/2$. Written in the S^x eigenbasis, this state is

$$|\uparrow_z\rangle = \frac{1}{\sqrt{2}}(|\uparrow_x\rangle + |\downarrow_x\rangle). \quad (3.1)$$

An S^z measurement yielding $m = 1/2$ projects the state exactly onto $|\uparrow_z\rangle$, but a consecutive S^x measurement will therefore give $\pm 1/2$ with equal probability—the outcome is completely random. Second, for a fixed spin quantum number S an S^z measurement can only yield the $2S+1$ discrete outcomes $m = -S, -S+1, \dots, S-1, S$. By contrast, a classical angular momentum is represented by a three-dimensional vector with continuous components, all of which can be specified simultaneously.

Interestingly, the behavior of a classical angular momentum can be recovered from quantum spin operators by taking the limit $S \rightarrow \infty$, as we will now show. In order to retain a finite total angular momentum in this limit, the spin operators need to be rescaled as

$$\mathbf{s} = \frac{1}{S} \mathbf{S}. \quad (3.2)$$

The commutators then become

$$[s^a, s^b] = \frac{1}{S^2} i \epsilon_{abc} s^c \xrightarrow{S \rightarrow \infty} 0, \quad (3.3)$$

so the components of \mathbf{s} commute and the eigenvalues of s^z take the continuous values

$$\left\{ \frac{m}{S} \mid m = -S, -S+1, \dots, S-1, S \right\} \xrightarrow{S \rightarrow \infty} [-1, 1]. \quad (3.4)$$

The operators \mathbf{s} therefore effectively behave like three dimensional vectors normalized to

$$\mathbf{s}^2 = \frac{S(S+1)}{S^2} \mathbf{1} \xrightarrow{S \rightarrow \infty} \mathbf{1}. \quad (3.5)$$

¹Mathematically, the two formulations are of course equivalent when restricted to a finite lattice.

These observations motivate the replacement of spin operators \mathbf{S} by real-valued three-dimensional vectors \mathbf{s} normalized to $|\mathbf{s}|^2 = 1$. Although the arguments above were heuristic, it was indeed rigorously shown that the partition function of a quantum spin model converges to the partition function of the corresponding classical spin model in the (rescaled) limit of $S \rightarrow \infty$ [85, 86]. Throughout the rest of this thesis, we use \mathbf{S} for both classical and quantum spins, with the distinction evident from context.

3.1.2 The strong and weak constraint

The goal of the LT method is to find the ground-state spin configuration that minimizes the energy of the classical Heisenberg model²

$$H = \sum_{ij} J_{ij} \mathbf{S}_i \mathbf{S}_j, \quad (3.6)$$

where \mathbf{S}_i are classical spins. The minimization is made considerably harder due to the restriction that classical spins must have fixed length,

$$\mathbf{S}_i^2 = 1 \quad \forall i, \quad (3.7)$$

which in the context of the LT method is called the *strong constraint*. The strong constraint is inherently nonlinear: if two spin configurations each satisfy the constraint, their superposition will, in general, not. The central idea of the LT method is to relax this requirement and instead minimize the Hamiltonian under the *weak constraint*

$$\sum_i |\mathbf{S}_i|^2 = N. \quad (3.8)$$

This condition directly follows from the strong constraint, but only implies a global normalization which significantly simplifies the problem. Every configuration satisfying the strong constraint also satisfies the weak constraint. Thus, if a minimization under the weak constraint happens to yield a configuration that also fulfills the strong constraint, we have found the true classical ground state. If, on the other hand, we only find spin-configurations that fulfill the weak constraint, their energy still provides a lower bound for the ground-state energy.

Interestingly, it has even been argued that enforcing only the weak constraint may better approximate the *quantum* problem, where the local magnetizations $|\langle \mathbf{S} \rangle|^2$ don't have a fixed magnitude [87]. While this is only a heuristic point, we indeed find remarkable agreement between weak-constraint LT results and those obtained from the pseudo-fermion functional renormalization group, as will be discussed in Part II.

3.1.3 Diagonalization in real space

Let us now explain how the minimization of the Hamiltonian under the weak constraint can be carried out in practice. The weak constraint can be enforced by introducing a *single* lagrange multiplier λ into the Hamiltonian

$$H_\lambda = H - \lambda \sum_i (\mathbf{S}_i^2 - 1) \quad (3.9)$$

²Applications to spin models with interactions beyond Heisenberg type are possible. In this thesis, we only apply the LT method to pure Heisenberg models and thus restrict ourselves also in this section.

and searching for the stationary state of this function with respect to both λ and the spin configuration $\{\mathbf{S}_i\}$. The condition $\partial_\lambda H_\lambda = 0$ simply recovers the weak constraint (3.8), while $\partial_{\mathbf{S}_i} H_\lambda = 0$ leads to the eigenvalue equation

$$\sum_j J_{ij} \mathbf{S}_j = \lambda \mathbf{S}_i . \quad (3.10)$$

Stationary states are therefore configurations $\{\mathbf{S}_{\lambda i}\}$ where each spin-component is an eigenvector of J_{ij} with the eigenvalue λ , and which in addition satisfy the weak constraint. The energy of such a configuration is

$$H(\{\mathbf{S}_{\lambda i}\}) = \sum_{ij} \mathbf{S}_{\lambda i} J_{ij} \mathbf{S}_{\lambda j} = N\lambda \geq N\lambda_{\text{LT}} , \quad (3.11)$$

which is minimized by choosing eigenvectors with the lowest eigenvalue, which we denote by λ^{LT} .

We next show how to explicitly obtain such a spin-configuration $\{\tilde{\mathbf{S}}_i\} \equiv \{\mathbf{S}_{\lambda^{\text{LT}} i}\}$. Since J_{ij} is real and symmetric, it can always be diagonalized by a set of orthogonal eigenvectors. We denote the set of such eigenvectors with eigenvalue λ_{LT} as $\{\psi_i^\gamma\}$, with $\gamma = 1, \dots, \Gamma$ (for the nondegenerate case $\Gamma = 1$). Since the eigenvalue equation is valid for all three spin components independently, a spin configuration can then be constructed by superimposing these eigenvectors with any set of three-dimensional coefficients $\{\mathbf{c}^\gamma\}$ as

$$\tilde{\mathbf{S}}_i(\{\mathbf{c}^\gamma\}) = \sum_{\gamma=1}^{\Gamma} \psi_i^\gamma \mathbf{c}^\gamma . \quad (3.12)$$

The weak constraint can be enforced by a suitable normalization of ψ_i^γ and \mathbf{c}^γ . A convenient choice is

$$\sum_i (\psi_i^\gamma)^2 = N/\Gamma . \quad (3.13)$$

Inserting $\tilde{\mathbf{S}}_i(\{\mathbf{c}^\gamma\})$ with this normalization into the weak constraint (3.8) implies that

$$|\mathbf{c}^\gamma|^2 = 1 , \quad (3.14)$$

where we used that the orthogonality relation

$$\sum_i \psi_{\lambda,i}^\gamma \psi_{\lambda,i}^{\gamma'} = \frac{N}{\Gamma} \delta_{\gamma,\gamma'} . \quad (3.15)$$

Thus, for any choice unit vectors $\{\mathbf{c}^\gamma\}$, we obtain a spin configuration $\tilde{\mathbf{S}}_i(\{\mathbf{c}^\gamma\})$ that minimizes the energy under the weak constraint. In general, however, such a configuration does not satisfy the strong constraint. To numerically check whether a particular choice of \mathbf{c}^γ exists where it does, we can minimize the deviation from the strong constraint (SC) defined as

$$\chi_{\text{SC}}^2(\{\mathbf{c}^\gamma\}) = \frac{1}{N} \sum_i \left(\tilde{\mathbf{S}}_i^2(\{\mathbf{c}^\gamma\}) - 1 \right)^2 . \quad (3.16)$$

If this minimization yields a set of $\{\mathbf{c}^\gamma\}$ such that $\chi_{\text{SC}}^2(\{\mathbf{c}^\gamma\}) = 0$, then the corresponding spin configuration is a true ground-state spin configuration on the finite lattice. If no such set exists, then the real-space LT method on a finite lattice cannot capture the ground state.

For the cases we studied, if there are fewer than three degenerate eigenvectors, i.e. $\Gamma \leq 3$, the strong constraint was always satisfied by choosing the set \mathbf{c}^γ to be orthogonal. The simplest choice in this case is to take Cartesian unit vectors, $\mathbf{c}^\gamma = \mathbf{e}^\gamma$, for which $\tilde{S}_i^a = \psi_i^a$. We have, however, not attempted to prove that this holds in general. For $\Gamma > 3$, or when this simple choice fails, one can instead numerically minimize χ_{SC}^2 .

Finally, to approximate the ground state of the infinite lattice, the finite lattice used in the LT calculation must be chosen carefully to avoid finite-size effects. In practice, this can be checked by varying the lattice size and verifying whether the spin configuration or the ground-state energy per site changes.

3.1.4 Diagonalization in momentum space

The more common approach to the Luttinger-Tisza method is to diagonalize the coupling matrix J_{ij} in momentum space for an infinite lattice. This can be achieved by exploiting the translational invariance of the spin model. To this end, we split the lattice index into two parts, $i = (n, \alpha)$, where n labels the unit cell, and $\alpha = 1, \dots, M$ the basis site in the unit cell (i.e., the sublattice). For pure Bravais lattices, $M = 1$. Translational invariance then implies

$$J_{ij} \equiv J_{m\alpha, n\beta} = J_{\alpha\beta}(\mathbf{R}_n - \mathbf{R}_m), \quad (3.17)$$

where \mathbf{R}_n is the position of the unit cell n . Defining the Fourier transform of the spin operator as

$$\mathbf{S}_{m\alpha} = \sum_{\mathbf{q} \in 1.BZ} e^{i\mathbf{q}\mathbf{R}_m} \mathbf{S}_\alpha(\mathbf{q}) \quad \mathbf{S}_\alpha(\mathbf{q}) = \frac{M}{N} \sum_i e^{-i\mathbf{q}\mathbf{R}_m} \mathbf{S}_{m\alpha}, \quad (3.18)$$

we can rewrite all relations from the real-space formulation in momentum space. The eigenvalue equation (3.10) then becomes

$$\sum_\beta J_{\alpha\beta}(\mathbf{q}) \mathbf{S}_\beta(\mathbf{q}) = \lambda \mathbf{S}_\alpha(\mathbf{q}), \quad (3.19)$$

with the Fourier-transformed coupling matrix

$$J_{\alpha\beta}(\mathbf{q}) = \sum_n e^{i\mathbf{q}\mathbf{R}_n} J_{\alpha\beta}(\mathbf{R}_n), \quad (3.20)$$

and where \mathbf{q} is restricted to the first Brillouin zone. This matrix is Hermitian and satisfies

$$J_{\alpha\beta}(-\mathbf{q}) = J_{\alpha\beta}(\mathbf{q})^*, \quad (3.21)$$

implying that all eigenvalues are real, the eigenvalues at $\pm\mathbf{q}$ are identical and the corresponding eigenvectors are complex conjugates.

Instead of diagonalizing the full $N \times N$ matrix J_{ij} , in the momentum space approach we only need to diagonalize the $M \times M$ matrix $J_{\alpha\beta}(\mathbf{q})$ for every \mathbf{q} in the first Brillouin zone.

The weak constraint now reads

$$\sum_\alpha \sum_{\mathbf{q} \in 1.BZ} |\mathbf{S}_\alpha(\mathbf{q})|^2 = M, \quad (3.22)$$

where we used $\sum_n e^{i(\mathbf{q}+\mathbf{q}')} = N/M \delta_{\mathbf{q}, \mathbf{q}'}$ and $\mathbf{S}_\alpha(-\mathbf{q}) = \mathbf{S}_\alpha(\mathbf{q})^*$. Similarly, the classical energy can be expressed as

$$H = \frac{N}{M} \sum_{\alpha\beta} J_{\alpha\beta}(\mathbf{q}) \mathbf{S}_\alpha(-\mathbf{q}) \mathbf{S}_\beta(\mathbf{q}). \quad (3.23)$$

In order to see how to obtain minimal energy configurations, we express the spin operators in terms of the M eigenvectors $\psi_\alpha^\beta(\mathbf{q})$ of $J_{\alpha\beta}(\mathbf{q})$ associated to the eigenvalues $\lambda^\beta(\mathbf{q})$ as

$$\mathbf{S}_\alpha(\mathbf{q}) = \sum_\gamma \mathbf{c}^\beta(\mathbf{q}) \psi_\alpha^\beta(\mathbf{q}), \quad (3.24)$$

where the coefficients $\mathbf{c}^\beta(\mathbf{q}) \in \mathbb{C}^3$ are *complex* three-component vectors. Choosing the normalization

$$\sum_\alpha |\psi_\alpha^\beta(\mathbf{q})|^2 = M, \quad (3.25)$$

the weak constraint reduces to

$$\sum_{\mathbf{q}} \sum_\beta |\mathbf{c}^\beta(\mathbf{q})|^2 = 1, \quad (3.26)$$

which means only finite weight can be attributed to the different eigenvalues $\lambda^\beta(\mathbf{q})$ and wave-vectors \mathbf{q} . For the energy of a stationary state this implies

$$H(\{\mathbf{c}^\beta(\mathbf{q})\}) = N \sum_{\mathbf{q}} \sum_\beta |\mathbf{c}^\beta(\mathbf{q})|^2 \lambda^\beta(\mathbf{q}) \geq N \lambda^{\text{LT}} \quad (3.27)$$

where λ^{LT} is again the smallest eigenvalue of $J_{\alpha\beta}(\mathbf{q})$ for any \mathbf{q} . This is consistent with the real-space result discussed above. Thus, a state with minimal energy $N\lambda^{\text{LT}}$ is therefore obtained by only including wave-vectors \mathbf{q}^{LT} and their eigenvectors with minimal eigenvalues $\lambda = \lambda^{\text{LT}}$ and by setting $\mathbf{c}^\gamma(\mathbf{q})$ to zero for all other momenta and eigenvectors.

The set of minimizing wave-vectors $\{\mathbf{q}^{\text{LT}}\}$ already provides important information about the ground state. A discrete set of minimal \mathbf{q}^{LT} vectors generally corresponds to an ordered state. If \mathbf{q}^{LT} vectors lie at incommensurate momenta, they indicate possible incommensurate phases—although in practice the corresponding real-space configurations often fail to satisfy the strong constraint. Finally, if the \mathbf{q}^{LT} vectors form a continuous manifold, this signals an extensive degeneracy which may imply classical spin liquid behavior [27].

Having found minimal \mathbf{q}^{LT} vectors and the corresponding eigenvectors, we still need to specify how to construct a real-space spin configuration. Let us assume there are Γ eigenvectors $\psi_\alpha^\gamma(\mathbf{q}^{\text{LT}})$ with minimal eigenvalue, for each minimal \mathbf{q}^{LT} -vector. The most general Ansatz for the minimal-energy spin configuration in momentum space is then the multi- \mathbf{q} state

$$\tilde{\mathbf{S}}_\alpha(\mathbf{q}) = \sum_{\mathbf{q}^{\text{LT}}} \delta_{\mathbf{q}, \mathbf{q}^{\text{LT}}} \sum_{\gamma=1}^{\Gamma} \mathbf{c}^\gamma(\mathbf{q}^{\text{LT}}) \psi_\alpha^\gamma(\mathbf{q}^{\text{LT}}), \quad (3.28)$$

where the coefficients have to satisfy $\mathbf{c}^\beta(-\mathbf{q}) = \mathbf{c}^\beta(\mathbf{q})^*$ so that $\mathbf{S}_\alpha(-\mathbf{q}) = \mathbf{S}_\alpha(\mathbf{q})^*$, insuring real spin configurations. The weak constraint can be satisfied by suitable normalization of $\mathbf{c}^\gamma(\mathbf{q}^{\text{LT}})$. An inverse Fourier transform yields the real-space spin configuration

$$\begin{aligned} \mathbf{S}_{m\alpha} &= \sum_{\mathbf{q}^{\text{LT}}} \sum_\gamma e^{i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \mathbf{c}^\gamma(\mathbf{q}^{\text{LT}}) \psi_\alpha^\gamma(\mathbf{q}^{\text{LT}}) \\ &= \sum_{\mathbf{q}^{\text{LT}} > 0} \sum_\gamma \left(e^{i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \mathbf{c}^\gamma(\mathbf{q}^{\text{LT}}) \psi_\alpha^\gamma(\mathbf{q}^{\text{LT}}) + e^{-i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \mathbf{c}^\gamma(\mathbf{q}^{\text{LT}})^* \psi_\alpha^\gamma(\mathbf{q}^{\text{LT}})^* \right) \\ &= \sum_{\mathbf{q}^{\text{LT}} > 0} \sum_\gamma 2 \operatorname{Re} \left(e^{i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \mathbf{c}^\gamma(\mathbf{q}^{\text{LT}}) \psi_\alpha^\gamma(\mathbf{q}^{\text{LT}}) \right), \end{aligned} \quad (3.29)$$

where in the first step we used $\mathbf{c}^\beta(-\mathbf{q}) = \mathbf{c}^\beta(\mathbf{q})^*$ and $\psi_\alpha^\gamma(-\mathbf{q}) = \psi_\alpha^\gamma(-\mathbf{q})^*$. The sum $\mathbf{q}^{\text{LT}} > 0$ indicates that only one member of each degenerate pair $\pm\mathbf{q}^{\text{LT}}$ is included, and $\text{Re}(\dots)$ denotes taking only the real part. To obtain physical spin configurations, we can again minimize the deviation from the strong constraint χ_{SC}^2 (defined in Eq. (3.11)) numerically. In practice, however, this becomes increasingly challenging when many degenerate \mathbf{q}^{LT} vectors are present.

That is why, typically, the Ansatz is simplified to a single- \mathbf{q} state, where only one eigenvector associated to a single momentum pair $\pm\mathbf{q}^{\text{LT}}$ is used to construct the spin configuration. In the case of multiple degenerate eigenvectors at the same \mathbf{q}^{LT} , we may freely choose any vector ψ_α within the degenerate subspace. The single- \mathbf{q} state is then takes the form

$$\mathbf{S}_{m\alpha} = e^{i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \psi_\alpha \mathbf{c} + e^{-i\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} \psi_\alpha^* \mathbf{c}^*. \quad (3.30)$$

We now have to make a suitable choice for the complex coefficients \mathbf{c} . The weak constraint (3.26) already implies $|\mathbf{c}|^2 = 1/2$. Because the Heisenberg Hamiltonian is rotationally invariant, for single- \mathbf{q} states we can restrict the \mathbf{c} to lie in a plane. To ensure $\mathbf{S}_{m\alpha}^2$ is independent of the unit cell m , a convenient choice is

$$\mathbf{c} = \frac{1}{2\sqrt{2}} (1, e^{\pm i\pi/2}, 0)^T, \quad (3.31)$$

which yields the spiral spin configuration

$$\mathbf{S}_{m\alpha} = |\psi_\alpha| \begin{pmatrix} \cos(\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m + \phi_\alpha) \\ \pm \sin(\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m + \phi_\alpha) \\ 0 \end{pmatrix}, \quad (3.32)$$

where ϕ_α is phase of $\psi_\alpha = |\psi_\alpha| e^{i\phi_\alpha}$. We are free to choose also any plane other than the xy -plane and can add an additional global phase, which would simply rotate the spiral configuration as a whole.

The strong constraint for the single- \mathbf{q} spiral becomes

$$\mathbf{S}_{m\alpha}^2 = |\psi_\alpha|^2 = 1 \quad \forall \alpha = 1, \dots, M. \quad (3.33)$$

Thus, whenever an eigenvector of $J_{\alpha\beta}(\mathbf{q}^{\text{LT}})$ has constant magnitude $|\psi_\alpha|$ across all basis sites, a single- \mathbf{q} spiral provides a valid ground-state configuration satisfying the strong constraint. This is trivially fulfilled for all Bravais lattices, since there is only one basis site. This proves the spiral theorem for Bravais lattices: the ground state of a classical Heisenberg model with interactions that respect the translational symmetry of the lattice is always a coplanar spiral [81, 82]. We note that this result has additionally been proven for crystal lattices with two symmetry equivalent basis sites per unit cell (see, e.g., Ref. [88]).

Let us summarize the practical steps we use to apply the momentum-space Luttinger-Tisza approach. The first step is to implement a function that calculates the coupling matrix in momentum space $J_{\alpha\beta}(\mathbf{q})$ (for simple cases this can be done analytically). Next, we need to identify all wave vectors \mathbf{q}^{LT} in the first Brillouin zone for which $J_{\alpha\beta}(\mathbf{q}^{\text{LT}})$ has the minimal eigenvalue λ^{LT} . In practice, we define a function $\lambda_{\text{min}}(\mathbf{q}, J_{\alpha\beta})$ that calculates the lowest eigenvalue of the coupling matrix at wave-vector \mathbf{q} , and then perform a global minimization over the Brillouin zone. To ensure that all global minima are identified, we start the minimization from a large set of \mathbf{q} -vectors. A practical approach is to first use the set of physically allowed \mathbf{q} -vectors of a finite but large periodic lattice, then gradually increase the lattice size until the set of minimizing vectors \mathbf{q}^{LT} converges. In cases where \mathbf{q}^{LT} forms a continuous manifold, we instead increase the size until the desired resolution in \mathbf{q} -space is reached. For each vector in $\{\mathbf{q}^{\text{LT}}\}$ we then compute the

Table 3.1 – Collinear, coplanar, and noncoplanar spin configurations from Luttinger–Tisza.

Shown are results for nearest-neighbor AFMs on the honeycomb, triangular, and maple-leaf lattices. For the maple-leaf lattice, the three inequivalent nearest-neighbor couplings are chosen such that the system lies in the noncoplanar phase VI defined in Chapter 6. The first column depicts the lattices, with the magnetic unit cell indicated by a dashed outline and inequivalent spin directions marked by different colors. The second column illustrates the corresponding spin orientations on the unit sphere. The third column shows the minimizing \mathbf{q}^{LT} vectors in momentum space, and the fourth column gives the explicit formula for the real-space spin configuration within unit cell m and sublattice α , where ψ_α denotes the LT eigenvectors in momentum space.

| Lattice | Spins | \mathbf{q}^{LT} vectors | Real-space configuration |
|---------|-------|----------------------------------|---|
| | | | $S_{m\alpha} = -1^\alpha \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ |
| | | | $S_m = \begin{pmatrix} \cos(\mathbf{K} \cdot \mathbf{R}_m) \\ \sin(\mathbf{K} \cdot \mathbf{R}_m) \\ 0 \end{pmatrix}$ |
| | | | $S_{m\alpha} = \begin{pmatrix} \psi_\alpha(\mathbf{q}_1^{\text{LT}}) \cos(\mathbf{q}_1^{\text{LT}} \cdot \mathbf{R}_m) \\ \psi_\alpha(\mathbf{q}_2^{\text{LT}}) \cos(\mathbf{q}_2^{\text{LT}} \cdot \mathbf{R}_m) \\ \psi_\alpha(\mathbf{q}_3^{\text{LT}}) \cos(\mathbf{q}_3^{\text{LT}} \cdot \mathbf{R}_m) \end{pmatrix}$ |

(possibly degenerate) eigenvectors ψ_α^γ with eigenvalue λ^{LT} and check whether any superposition $\psi_\alpha = \sum_\gamma c^\gamma \psi_\alpha^\gamma$ of these eigenvectors has constant magnitude $|\psi_\alpha|$ across all basis sites α . If this condition is satisfied, we have found a spiral spin configuration fulfilling the strong constraint, which can be constructed in real space using Eq. (3.32). If no such superposition exists, one can instead employ the multi- \mathbf{q} Ansatz of Eq. (3.28) and attempt to numerically find coefficients $\{c^\gamma(\mathbf{q}^{\text{LT}})\}$ for which the strong constraint is satisfied. In practice, however, this is only feasible when the number of $\{\mathbf{q}^{\text{LT}}\}$ is small. If all $\{\mathbf{q}^{\text{LT}}\}$ are at commensurate wave-vectors, resorting to the real-space LT approach is usually the simpler choice.

3.1.5 Examples for collinear, coplanar and noncoplanar states

To conclude this section, let us briefly outline a few examples where the Luttinger-Tisza approach correctly captures collinear, coplanar, and even noncoplanar ground states of different classical spin models. A summary that includes illustrations of the considered lattices, the minimizing \mathbf{q}^{LT} vectors, and real-space ground states is provided in Table 3.1.

Collinear state on the Honeycomb lattice. For the AFM nearest-neighbor Heisenberg model on the honeycomb lattice, the ground-state is a simple collinear Néel state, with spins pointing in opposite directions on the two sublattices. In real-space, diagonalizing J_{ij} for any periodic honeycomb lattice of size $N_1 \times N_2$ (where the unit cell is repeated N_i times along translation vector a_i) yields a single eigenvector $\psi_i = (1, -1, 1, -1, \dots)$ of constant magnitude but alternating sign between sublattices. Choosing, for example, $\mathbf{c} = (1, 0, 0)$ in the real-space Ansatz (3.12) directly produces the correct Néel ground state with spins along the x -axis.

In momentum space, diagonalization yields a single minimizing wave-vector $\mathbf{q}^{\text{LT}} = \mathbf{\Gamma} = (0, 0)^T$ with the eigenvector $\psi_\alpha = (1, -1)$, which fulfills the strong constraint as it has constant magnitude. The spiral Ansatz in Eq. (3.32) then reduces to

$$\mathbf{S}_{m\alpha} = -1^\alpha (1, 0, 0)^T \quad (3.34)$$

again giving the correct Néel order.

Coplanar state on the triangular lattice. For the AFM Heisenberg model on the triangular lattice, the ground state exhibits the well-known 120° coplanar order, in which neighboring spins enclose angles of 120° .

In real space, finite lattices of size $3N_1 \times 3N_2$ are commensurate with this order. Diagonalization of J_{ij} yields two eigenvectors $\psi^1 = (-1, 1/2, 1/2, \dots)^T$ and $\psi^2 = (0, \sqrt{3}/2, -\sqrt{3}/2, \dots)^T$ (where we only state the first components that lie on a representative triangle). Choosing orthogonal unit vectors $\mathbf{c}^1 = (1, 0, 0)$ and $\mathbf{c}^2 = (0, 1, 0)$ produces the correct 120° configuration satisfying the strong constraint.

In momentum space, the minimizing wave vectors occur at the Brillouin-zone corners $\mathbf{q}^{\text{LT}} = \mathbf{K}, \mathbf{K}'$, where $\mathbf{K} = 2\mathbf{G}_1/3 + \mathbf{G}_2/3$ and $\mathbf{K}' = \mathbf{G}_1/3 + 2\mathbf{G}_2/3$, with the reciprocal lattice vector \mathbf{G}_i of the triangular lattice. Because the triangular lattice is a Bravais lattice, the eigenvector is trivial $\psi^i = 1$ and the single- \mathbf{q} spiral Ansatz for both \mathbf{K} and \mathbf{K}' gives

$$\mathbf{S}_m = \begin{pmatrix} \cos(\mathbf{K} \cdot \mathbf{R}_m) \\ \pm \sin(\mathbf{K} \cdot \mathbf{R}_m) \\ 0 \end{pmatrix}, \quad (3.35)$$

which is the correct 120° order.

noncoplanar state on the maple-leaf lattice Finally, we show an example where genuinely three-dimensional, noncoplanar ground states can also be captured by the LT approach. We identified two such phases in an nearest-neighbor Heisenberg model on the maple-leaf lattice, which has three symmetry inequivalent interactions. A definition of the model and a detailed discussion of the states will be given in Chapter 6. Here, we focus on couplings in phase VI, where the classical ground state has a 24 site magnetic unit cell whose spins point in twelve different direction that cannot be confined to a single plane. A visualization of the maple-leaf lattice and the ground state is given in the last row of Table 3.1.

In real space, for periodic lattices of size $(2N_1 \times 2N_2)$, there are three eigenvectors ψ_i^γ . For any choice of three orthogonal unit vectors $\{\mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3\}$ the strong constraint is satisfied and the correct noncoplanar ground state is obtained. The simplest choice is again $\mathbf{S}_i^a = \psi_i^a$.

In momentum space, the situation is slightly more complicated, since the single- \mathbf{q} Ansatz cannot capture noncoplanar states. Diagonalization yields three \mathbf{q}^{LT} vectors,

$$\mathbf{q}_1^{\text{LT}} = \frac{\mathbf{G}_1}{2}, \quad \mathbf{q}_2^{\text{LT}} = \frac{\mathbf{G}_1 + \mathbf{G}_2}{2}, \quad \mathbf{q}_3^{\text{LT}} = \frac{\mathbf{G}_2}{2}, \quad (3.36)$$

where \mathbf{G}_i are reciprocal lattice vectors of the maple-leaf lattice. At each of those wave vectors, the coupling matrix has a single eigenvector $\psi_\alpha(\mathbf{q}^{\text{LT}})$ with minimal eigenvalue λ^{LT} . None of them individually fulfill the strong constraint under the single- \mathbf{q} Ansatz, as $|\psi_\alpha|$ varies across the unit cell. This rules out a simple spiral state and necessitates the multi- \mathbf{q} Ansatz of Eq. (3.28).

Numerical minimization of χ_{SC}^2 indeed produces a set of $\{\mathbf{c}(\mathbf{q}^{\text{LT}})\}$ that correctly reproduces the correct noncoplanar ground state satisfying the strong constraint. Guided by the real-space construction, however, we can verify that simply taking $\mathbf{c}(\mathbf{q}_i^{\text{LT}})$ as the three Cartesian unit vectors also produces the correct state. Substituting this into the multi- \mathbf{q} Ansatz yields

$$\mathbf{S}_{m\alpha} = \begin{pmatrix} \psi_\alpha(\mathbf{q}_1^{\text{LT}}) \cos(\mathbf{q}_1^{\text{LT}} \cdot \mathbf{R}_m) \\ \psi_\alpha(\mathbf{q}_2^{\text{LT}}) \cos(\mathbf{q}_2^{\text{LT}} \cdot \mathbf{R}_m) \\ \psi_\alpha(\mathbf{q}_3^{\text{LT}}) \cos(\mathbf{q}_3^{\text{LT}} \cdot \mathbf{R}_m) \end{pmatrix}, \quad (3.37)$$

where the fact that all three \mathbf{q}^{LT} vectors are half reciprocal lattice vectors ensures that $e^{\mathbf{q}^{\text{LT}} \cdot \mathbf{R}_m} = \pm 1$ are real, rendering all eigenvectors real as well. The result can be viewed as a superposition of three spiral-like states which, remarkably, combine such that they satisfy the strong constraint.

3.2 Semiclassical Monte Carlo for SU(4) spin models

In the introductory chapters, we showed that in moiré materials and certain spin-orbit entangled Mott insulators, the local moments may not be conventional spin-operators—which are generators of SU(2)—but spin-orbit or spin-valley degrees of freedom described by generators of SU(4). Our initial motivation for studying such systems is that enlarging the symmetry from SU(2) to SU(4) enhances the effect of quantum fluctuations, similar to frustration, and may thus increase the propensity toward novel quantum-disordered ground states. In realistic models for moiré materials such as twisted bilayer graphene [89, 90], trilayer graphene aligned with h-BN [49, 91, 92], or in transition metals with an orbital degeneracy [50], however, SU(4) symmetry is often broken by (super)-exchange interactions, and ordered states are natural competitors.

To analyze such states in SU(2) spin models, a common approach is to employ the classical $S \rightarrow \infty$ limit discussed in the previous section. A naive extension of this approach to SU(4) spin models might be to decouple the spin and valley degrees of freedom, and then take the $S \rightarrow \infty$ limit for both of them. We begin this section, however, by showing that this approach fails to capture the correct structure of the local Hilbert space, as it neglects important contributions from the local entanglement between spin and valley—especially for a filling of two fermions per side. We then introduce a more appropriate semiclassical limit that preserves the local structure of the SU(4) Hilbert space while neglecting only inter-site entanglement. Finally, we describe how, in this limit, we can explore both the semiclassical ground state and the effects of thermal fluctuations efficiently using conventional Metropolis Monte Carlo and simulated annealing algorithms, which we have implemented in the publicly available Julia package `SemiClassicalMC.jl` [C2].

The method we present was originally developed in Ref. [93] to study spin-nematic correlations in $S = 1$ SU(2) spin models, and later applied to the Bose-Hubbard model in the strong-coupling limit [94]. The semiclassical limit was first extended to SU(4) models to analyze the ground state of the SU(4)-symmetric Heisenberg model in the six-dimensional representation [53]. In our own work [P1, P2], we built on this approach and generalized it to models with broken SU(4) symmetry and finite temperatures. Our presentation here is strongly based on these references.

3.2.1 The local Hilbert-space of SU(4) spin models

Before deriving the semiclassical limit, we first specify the structure of the local Hilbert space of SU(4) models in their representation in terms of fermions with both a spin and valley (or orbital) degree of freedom, as motivated in Chapter 1 and Sec. 2.2.5. For brevity, we adopt the spin–valley terminology relevant to Chapter 7 from now on.

Generic SU(4) spin models are of them the form

$$H = \sum_{ij} \sum_{a,b=1}^{15} T_i^a J_{ij}^{ab} T_j^b, \quad (3.38)$$

where the T_i^a form a basis of the 15-dimensional Lie algebra of SU(4). In principle, any basis can be chosen. In the context of spin–valley entangled Mott insulators, however, a particularly natural choice is provided by spin–valley operators (for brevity, we adopt the spin–valley terminology relevant to Chapter 7 from now on), defined as³

$$T^a = \sigma_i^\mu \tau_i^\kappa = P_n \left(f_{isl}^\dagger \theta_{ss'}^\mu \theta_{ll'}^\kappa f_{is'l'} \right) P_n. \quad (3.39)$$

Here, θ^μ are Pauli matrices (with $\theta^0 = \mathbb{1}$ and $\mu, \kappa = 0, 1, 2, 3 = d, x, y, z$), f_{isl}^\dagger creates a fermion with spin $s = (\uparrow, \downarrow)$, and valley $l = (+, -)$, and we assume summation over repeated spin and valley indices. The spin–valley operators form a basis of the Lie algebra of SU(4) only when including the projector P_n onto the subspace of n fermions per site, effectively enforcing the filling

$$f_{isl}^\dagger f_{isl} = n \quad (3.40)$$

The dimension of the local Hilbert space is determined by the number of ways to place n fermions on a single site. With four possible fermion flavors, this gives $\binom{4}{n}$ states. The relevant cases are:

1. Quarter filling ($n = 1$): one fermion per site, four-dimensional Hilbertspace
2. Half filling ($n = 2$): two fermions per site, six-dimensional Hilbertspace
3. Three-quarter filling ($n = 3$): equivalent to $n = 1$ by particle-hole symmetry.

The trivial cases $n = 0, 4$ correspond to completely empty/filled bands without dynamics. We now discuss the Hilbert space of quarter and then half-filling in detail.

Quarter filling For one fermion per site, a basis of the four-dimensional local Hilbert space is

$$\mathcal{B}_1 = \{|\uparrow+\rangle, |\uparrow-\rangle, |\downarrow+\rangle, |\downarrow-\rangle\}, \quad (3.41)$$

with $|sl\rangle \equiv f_{sl}^\dagger |0\rangle$. It can be equivalently expressed in the factorized form

$$\mathcal{B}_1 = \{|s\rangle \otimes |l\rangle |s = (\uparrow, \downarrow), l = (+, -)\}, \quad (3.42)$$

where $|s\rangle$ and $|l\rangle$ are the eigenstates of conventional SU(2) operators. In this basis, the spin–valley operators are then also tensor products of SU(2) operators

$$\sigma_i^\mu \tau_i^\nu = \sigma_i^\mu \otimes \tau_i^\nu. \quad (3.43)$$

³To only consider the fifteen generators of SU(4) we need to exclude $\sigma^0 \tau^0 = \mathbb{1}_4$ which anyway only constitutes a constant in the spin model. We leave this in for ease of notation.

which, in the standard basis, can be explicitly represented as Kronecker products of Pauli matrices (up to possible factors of 1/2).

In this case, a commonly chosen classical limit is to mean-field decouple the spin and valley degrees of freedom in, e.g., the valley sector as

$$\sigma_i^\mu \otimes \tau_i^\nu \approx \sigma_i^\mu \langle \tau_i^\nu \rangle, \quad (3.44)$$

yielding an effective SU(2) Hamiltonian that can be treated in the conventional classical limit. Going one step further, one may instead approximate the tensor product with an outer product

$$\sigma_i^\mu \otimes \tau_i^\nu \approx \sigma_i^\mu \times \tau_i^\nu \quad (3.45)$$

and then take the $S \rightarrow \infty$ limit in both σ and τ . This results in an effective classical SU(2) spin model, but with two spins per site. Compared to the mean-field approach, this has the advantage that thermal fluctuations around the mean-field values can be taken into account in both the spin and valley sector.

Both approximations, however, share an important limitation: while they naturally discard inter-site entanglement—as any classical limit does—they also ignore intra-site “entanglement” between spin and valley degrees of freedom. In other words, they only capture states that factorize into a direct product $|\psi_s\rangle \otimes |\psi_l\rangle$, and therefore miss superpositions that mix spin and valley. For example, the state

$$\frac{1}{\sqrt{2}} (|\uparrow+\rangle + |\downarrow-\rangle) \quad (3.46)$$

is entirely excluded. As we will see next, this issue becomes even more pronounced at half-filling.

Half-filling At half-filling there are two fermions per site, giving a local Hilbert space of dimension $\binom{4}{2} = 6$. An explicit basis is

$$\{|\uparrow+, \downarrow+\rangle, |\uparrow+, \downarrow-\rangle, |\downarrow+, \uparrow-\rangle, |\uparrow+, \uparrow-\rangle, |\downarrow+, \downarrow-\rangle, |\uparrow-, \downarrow-\rangle\}, \quad (3.47)$$

with $|s_1 l_1, s_2 l_2\rangle \equiv f_{s_1 l_1}^\dagger f_{s_2 l_2}^\dagger |0\rangle$. At this filling, spin and valley become much more tightly intertwined than for quarter filling, as the Pauli principle enforces correlations between them. For instance, if both fermions have $s = \uparrow$, then they must occupy opposite valley states. Locally, this means $\langle \sigma^z \rangle = 1$ automatically implies $\langle \tau^z \rangle = 0$ regardless of the considered Hamiltonian. As a result, spin and valley can no longer be meaningfully decoupled, and there is no consistent way to take the $S \rightarrow \infty$ limit in the spin or valley degree of freedom.

3.2.2 The semiclassical limit

Fortunately, a semiclassical limit that incorporates the local entanglement between spin and valley can be defined in a straightforward way. The essential idea is to forbid entanglement *between sites*, but to keep the full local Hilbert space intact, without taking any $S \rightarrow \infty$ limit.

In the first step, we restrict the Hilbert space to product states of the form

$$|\psi\rangle = \otimes_i |\psi_i\rangle, \quad (3.48)$$

where $|\psi_i\rangle$ is an arbitrary state in the d -dimensional local Hilbert space. Choosing a basis of this space, in our case \mathcal{B}_1 or \mathcal{B}_2 defined above, we can parametrize each $|\psi_i\rangle$ by d complex amplitudes c_i^γ as

$$|\psi_i(\mathbf{c}_i)\rangle = \sum_{\gamma=1}^d c_i^\gamma |\gamma\rangle, \quad (3.49)$$

where $\{|\gamma\rangle\}$ label the $\gamma = 1, \dots, d$ basis states. Normalization of the state demands that the d -dimensional complex vector $\mathbf{c}_i = (c_i^1, \dots, c_i^d)^T$ satisfies the constraint

$$|\mathbf{c}_i|^2 = \sum_{\gamma=1}^d |c_i^\gamma|^2 = 1. \quad (3.50)$$

Thus, a system of N sites is described by N normalized d -dimensional complex vectors \mathbf{c}_i , which corresponds to $N(d - 2)$ real degrees of freedom once normalization and an arbitrary overall phase are removed. To indicate the dependence on these vectors, we will simply write $|\psi(c)\rangle \equiv |\psi(\{\mathbf{c}_i\})\rangle$.

In the next step, we define the semiclassical energy as a function of c as the expectation value of the corresponding product state as

$$H_{\text{sc}}(c) = \langle \psi(c) | H | \psi(c) \rangle. \quad (3.51)$$

Minimizing $H_{\text{sc}}(c)$ over \mathbf{c}_i then yields the semiclassical approximation to the ground state. Conceptually, this approach is equivalent to a mean-field theory where interaction between spin-valley operators are decoupled as

$$\sigma_i^\mu \tau_i^\nu \sigma_j^\kappa \tau_j^\lambda \approx \langle \sigma_i^\mu \tau_i^\nu \rangle \sigma_j^\kappa \tau_j^\lambda + \sigma_i^\mu \tau_i^\nu \langle \sigma_j^\kappa \tau_j^\lambda \rangle - \langle \sigma_i^\mu \tau_i^\nu \rangle \langle \sigma_j^\kappa \tau_j^\lambda \rangle \quad (3.52)$$

and the expectation values are determined self-consistently.

We can go one step further, however, and extend this framework to finite temperatures by including thermal fluctuations around the mean-field solution. This allows us to more accurately calculate expectation values at finite temperatures $T = 1/\beta$. Following Refs. [93, 94], we approximate the semiclassical partition function as

$$Z_{\text{sc}} = \int \prod_i d\mathbf{c}_i \langle \psi(c) | e^{-\beta H} | \psi(c) \rangle \approx \int \prod_i d\mathbf{c}_i e^{-\beta \langle \psi(c) | H | \psi(c) \rangle} = \int \prod_i d\mathbf{c}_i e^{-\beta H_{\text{sc}}(\{\mathbf{c}_i\})}, \quad (3.53)$$

where $\prod_i d\mathbf{c}_i \sim \prod_{i,\gamma} d\text{Re}(c_i^\gamma) d\text{Im}(c_i^\gamma) \delta(|\mathbf{c}_i|^2 - 1)$ includes the integration over all properly normalized sets of $\{\mathbf{c}_i\}$. This corresponds to truncating a cumulant expansion at first order, which becomes exact in the limit $T \rightarrow \infty$ and approaches the mean-field approach for $T \rightarrow 0$. We therefore assume that it captures the essential physics across the full temperature range.

Thermal expectation values of arbitrary operators O can then simply be computed as

$$\langle O \rangle = \int \prod_i d\mathbf{c}_i \langle \psi(c) | O | \psi(c) \rangle e^{-\beta H_{\text{sc}}(c)} \equiv \int \prod_i d\mathbf{c}_i O(c) e^{-\beta H_{\text{sc}}(c)} \quad (3.54)$$

Such integrals can be efficiently evaluated using Markov Chain Monte Carlo, which we describe in the following.

3.2.3 Monte Carlo implementation

To calculate thermal expectation values in the semiclassical limit, we use a Monte Carlo routine based on the standard Metropolis algorithm with local updates. The same framework also allows for simulated annealing to minimize the semiclassical energy and obtain the semiclassical ground state, which we further refine using stochastic gradient descent to more precisely converge to the minimal-energy configuration. We have implemented these algorithms in the publicly available Julia Package SemiClassicalMC.jl [C2]. Since the basics of Markov Chain Monte Carlo and the Metropolis algorithm are well covered in the literature, we only summarize the aspects relevant to our implementation here, and refer the reader to Ref. [95] for a more general introduction.

Markov Chain Monte Carlo integration Let us begin by briefly explaining the general idea behind applying such MCMC methods to calculate expectation value of the form shown in Eq. (3.54). The problem we are facing is to compute integrals over the all possible sets $\{\mathbf{c}_i\}$. As stated before, these can be parametrized by $d - 2$ real numbers per site. The dimension of the integration space therefore grows exponentially with the number of sites according to $d_{\text{int}} \sim (d - 2)^N$. Even for only few sites, this becomes impossible to compute using standard numerical integration techniques. The idea of applying MCMC to compute expectation values is to instead only generate only a finite number of configurations $\{\mathbf{c}_i\}$ according to the Boltzmann distribution $p(c) \sim \exp(-\beta H_{\text{sc}}(c))$. If one has obtained a set of M such configurations $c^{(m)}$, expectation values of operators can be approximated by

$$\langle O \rangle \approx \bar{O} = \frac{1}{M} \sum_{m=1}^M O(c^{(m)}). \quad (3.55)$$

The estimator \bar{O} clearly converges to exact expectation value $\lim_{M \rightarrow \infty} \bar{O} = \langle O \rangle$. For finite M , however, it is itself a statistical variable, with a finite statistical error. This error can be shown to scale as

$$\Delta \bar{O} = \sqrt{\text{var}(\bar{O})} \sim 1/\sqrt{M}, \quad (3.56)$$

which, crucially, is independent of the dimension of the integration regime d_{int} and therefore avoid the *curse of dimensionality* that other integration methods suffer.

Metropolis algorithm To perform the Monte Carlo integration we need to efficiently generate configurations $c^{(m)}$ according to the distribution $p(c) \sim \exp(-\beta H_{\text{sc}}(c))$. To this end, we utilize the conventional Metropolis algorithm with local updates.

From a configuration $c^{(m)}$, a trial configuration c^{new} is proposed and accepted with probability

$$p = \min \left(1, \exp[-\beta(H_{\text{sc}}(c^{\text{new}}) - H_{\text{sc}}(c^{(m)}))] \right). \quad (3.57)$$

If c^{new} is not accepted, instead we set $c^{(m+1)} = c^{(m)}$. The fraction of newly proposed configurations that are accepted over the total number of proposed configurations is called the acceptance rate R .

The generated configurations follow the correct Boltzmann distribution but are not statistically independent, as successive samples exhibit autocorrelations. While this does not bias expectation values, it complicates error estimation. To address this, observables are “measured” only at larger intervals rather than after every update, and statistical errors are extracted using binning analysis (via the Julia package `BinningAnalysis.jl` [96]), which systematically includes autocorrelation effects in the error estimation.

Local Metropolis updates To generate the trial configuration c^{new} , we perform a *local update*. This means we randomly select one of the N sites, say site i , and then update only the vector \mathbf{c}_i in c^{new} , while all other vectors $\mathbf{c}_{i \neq j}$ remain the same as in the previous $c^{(m)}$. A sequence of N such updates is called a *sweep*.

For classical SU(2) spins, which are normalized, real, three-dimensional vectors, a new trial spin can be obtained by sampling a random point on the unit 2-sphere, where one has to be careful to sample the full sphere uniformly. In our case, the situation seems quite different, yet is ultimately very similar: we need to uniformly sample the space of normalized, complex, d -dimensional vectors. This space, however, can also be interpreted as a real $(d - 1)$ -dimensional

hypersphere parameterized by the real and imaginary part of each vector component. A simple way to sample uniformly on an n -dimensional hypersphere (n -sphere) is to draw $n + 1$ random numbers from a *normal distribution* (note that using a uniform distribution does *not* give a uniform sampling on the sphere) and then normalize the resulting vector [97]. For complex vectors, this means drawing normally distributed random numbers for the real and imaginary parts of each component, forming a random complex vector Γ , and then setting

$$\mathbf{c}_i^{\text{new}} = \frac{\mathbf{\Gamma}}{|\mathbf{\Gamma}|}. \quad (3.58)$$

This procedure samples the hypersphere correctly, but at low temperatures it leads to very low acceptance rates, where most proposed moves are rejected and convergence becomes slow. To overcome this, we adopt the *Gaussian trial move* introduced for classical SU(2) spins in Ref. [98]. Instead of fully replacing \mathbf{c}_i with Γ , we generate a new state in the vicinity of \mathbf{c}_i as

$$\mathbf{c}_i^{\text{new}} = \frac{\mathbf{c}_i^{(m)} + \sigma_g \mathbf{\Gamma}}{|\mathbf{c}_i^{(m)} + \sigma_g \mathbf{\Gamma}|}. \quad (3.59)$$

Here, the parameter σ_g sets the average distance between $\mathbf{c}_i^{\text{new}}$ and $\mathbf{c}_i^{(m)}$. This is also an unbiased way of sampling the local Hilbert space, but has the advantage that the acceptance rate can be adjusted dynamically by controlling the value of σ_g : Starting from a large value, we update σ_g every tenth sweep as

$$\sigma_g \rightarrow \frac{0.5}{1 - R_{10}} \sigma_g, \quad (3.60)$$

where R_{10} is the acceptance rate over the last ten sweeps [98]. In practice, this procedure quickly stabilizes the acceptance rate near 50% and significantly accelerates convergence at low temperatures.

Thermalization The metropolis algorithm yields a Markov Chain that generates configurations $\{\mathbf{c}^{(m)}\}$ according to the desired probability distribution—but only once the chain has reached its stationary state. This is guaranteed only in the limit $m \rightarrow \infty$, but in practice occurs much earlier to any reasonable desired accuracy. It is, however, crucial to wait a certain number of *thermalization sweeps* N_t before taking measurements of observables, giving the system time to reach the steady state.

We refer to this stage as *thermalization*, because we typically start simulations from a completely random configuration, which corresponds to an infinite-temperature state. When performing simulations at low temperatures, the system must gradually relax into the region of configuration space relevant for that temperature, determined by its free energy functional. This process can, e.g., be observed in the energy, which decreases steadily until the steady state is reached, where it only fluctuates around its mean value. The time it takes to reach this steady state varies, and therefore the number of thermalization sweeps N_t has to be adjusted for each system and each temperature.

To accelerate thermalization and avoid trapping in local minima, we perform *simulated annealing* during the first $3/4N_t$ thermalization sweeps. The idea of simulated annealing is to start the simulation at a high temperature, and then only gradually lower it to the desired value. Let's say we want to perform simulations at the temperature T . We then start the thermalization at a much higher initial temperature, typically of order $T_i = 2|J|$ (where $|J|$ is some energy scale

given by the interactions), and a large value of $\sigma_g = 60$. We then perform Monte Carlo sweeps using the procedure explained above, but lower the temperature after each sweep as

$$T^{(m+1)} = T^{(m)} \times (T/T_i)^{\frac{4}{3N_t}} \quad (3.61)$$

until the target temperature T is reached. For the remaining $\frac{1}{4}N_t$ sweeps, the temperature is kept constant. Only after the thermalization stage, we begin the measurement sweeps during which the temperature and σ_g are kept constant.

Ground-state minimization by simulated annealing and gradient descent Beyond finite-temperature properties, we are also interested in ground-state phase diagrams, which requires finding the product state (or equivalently the configuration c), that minimizes the semiclassical energy.

Similarly to the challenge in calculating thermal expectation values, this problem also suffers from the high-dimensional configuration space in which the minimization has to be performed. Often, the classical energy will have many local minima and it is difficult to confirm that a minimization has converged to the true global minimum. Fortunately, simulated annealing is particularly suited for this task—it has even been shown to always converge to a global minimum given infinite time (which of course does not strictly guarantee anything for the practical case of finite time runtimes) [99]. At low temperatures, acceptance rates (or equivalently the step width σ_g) become very small, and convergence slows dramatically. In practice, we therefore terminate simulated annealing once R and σ_g drop below predefined thresholds, assuming the system is then sufficiently close to the global minimum. To refine the solution, we subsequently apply stochastic gradient descent, which converges much faster to the nearest minimum.

Unlike the fixed cooling schedule used during thermalization, here we allow for a variable number of sweeps per temperature to ensure proper equilibration. Concretely, the temperature is lowered only after $n_{\text{acc}}N$ updates have been accepted, which leads to frequent temperature decreases at high temperatures, where equilibration is fast, while still allowing for sufficient sampling at low temperatures. To avoid excessively long runs, we additionally impose an upper bound N_T on the number of sweeps per temperature. In practice, typical values we choose are $N_T \approx 4000$ and $n_{\text{acc}} \approx 400$. At each temperature step, we calculate the acceptance rate R , adjust σ_g according to Eq. (3.60) and then lower the temperature by 2 %. Once σ_g has reached the minimal value of $\sigma_g = 0.05$, we keep it fixed. We continue lowering the temperature until the acceptance rate is below $R_{\text{min}} = 0.001\%$, at which point we terminate the calculation.

Starting from the state obtained from simulated annealing, we perform *N_o optimization sweeps* based on stochastic gradient descent. The idea is analogous to local Monte Carlo updates: at each step, we randomly pick a site i , and then update the vector \mathbf{c}_i . Unlike in a Metropolis step, however, we don't choose a random new vector \mathbf{c}^{new} , but the one that locally minimizes the energy (while all other $\mathbf{c}_{j \neq i}$ are fixed). An optimization step therefore always results in a lower energy. The corresponding local energy to be minimizes with respect to \mathbf{c}_i is

$$H_{\text{sc}}^i(c) = \left\langle \psi(c) \left| \sum_{j,ab} J_{ij}^{ab} T_i^a T_j^b \right| \psi(c) \right\rangle = \sum_a \langle \psi_i(\mathbf{c}_i) | T_i^a | \psi_i(\mathbf{c}_i) \rangle \sum_{j,b} J_{ij}^{ab} \langle \psi_j(\mathbf{c}_j) | T_j^b | \psi_j(\mathbf{c}_j) \rangle. \quad (3.62)$$

For conventional SU(2) spin models without on-site interactions, this minimization can be formulated as a linear equation that is solved by setting the spin at site i antiparallel to the mean field of the coupled spins (e.g. $\mathbf{h}_i = \sum_j J_{ij} \mathbf{S}_j$ for simple Heisenberg models). In this context, the method is then often simply referred to as “iterative minimization” [100, 101]. In our case, however, H_{sc}^i is a nonlinear function of \mathbf{c}_i , so we resort to numerical minimization using gradient descent.

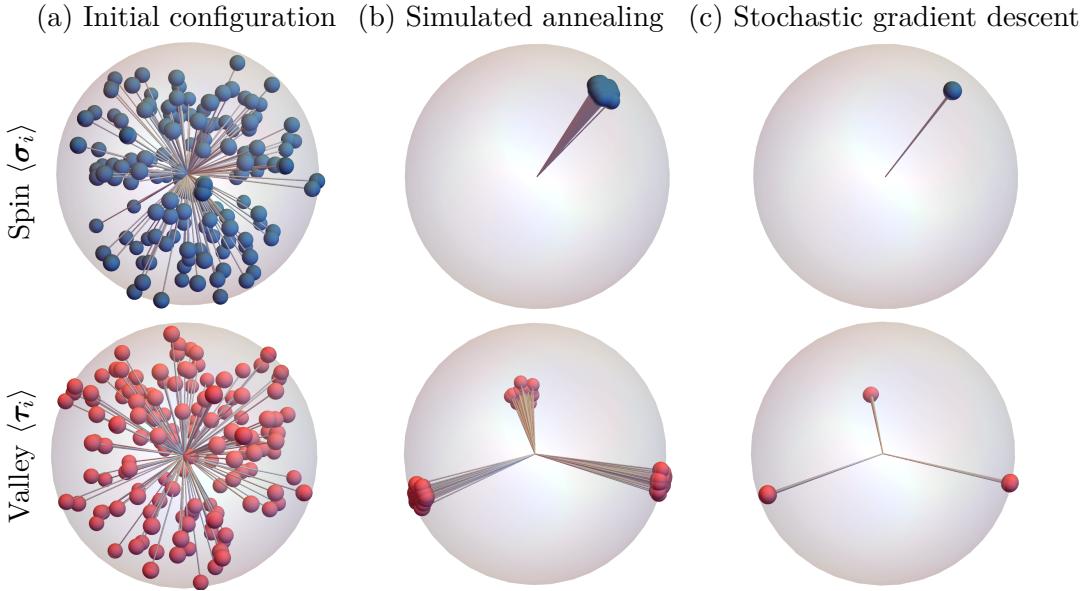


Figure 3.1 – Common-origin plots at different stages of the minimization procedure. Spin (top) and valley (bottom) expectation values for all $N = 144$ sites a $SU(2) \otimes SU(2)$ symmetric spin–valley model. Shown are (a) the initial random configuration, (b) the state after simulated annealing down to $T \approx 0.005|J|$, and (c) the result after ten optimization sweeps with stochastic gradient descent.

In gradient descent, \mathbf{c}_i is iteratively updated in the direction of the negative gradient of H_{sc}^i as

$$\mathbf{c}_i^{\text{new}} = \mathbf{c}_i - \lambda \nabla_{\mathbf{c}_i} H_{\text{sc}}^i(c), \quad (3.63)$$

where λ is a tunable step size. To preserve the normalization $|\mathbf{c}_i| = 1$, however, the minimization needs to be performed on the 11-sphere spanned by the real and imaginary part of \mathbf{c}_i . We implement this using the `Manifolds.jl` package [102] to compute gradients on the sphere (via finite differences) and `Manopt.jl` [103] to carry out the optimization. The gradient descent at site i is terminated once the gradient norm falls below $0.001|J|$.

One *optimization sweep* consists of performing this update on N randomly chosen sites. We repeat the optimization sweeps until the total energy converges to the desired precision. For all systems studied, the energy typically stabilizes well before $N_o \approx 100$ sweeps to an accuracy of $10^{-8}|J|$.

Lets us give a specific example that highlights the efficiency of the combined simulated annealing and stochastic gradient descent minimization. To this end, we consider the $SU(2) \otimes SU(2)$ symmetric spin–valley model on the triangular lattice (already studied, e.g., in Ref. [104]) defined by the Hamiltonian

$$H = \sum_{\langle ij \rangle} \left(J \sum_{a,b=1}^3 (\sigma_i^a \tau_j^b)(\sigma_j^a \tau_j^b) + J_s \sum_{a=1}^3 \sigma_i^a \sigma_j^a + J_v \sum_{a=1}^3 \tau_i^a \tau_j^a \right). \quad (3.64)$$

We choose $J > 0$, large $J_s = -2J$ and $J_v = 2J$, and quarter filling. In this regime, the semiclassical ground state is expected to exhibit FM order in the spin sector and 120° order in the valley sector. Performing the minimization procedure described above, we can confirm this by creating common origin plots of the spin and the valley expectation values. These are shown in Fig. 3.1 after different stages of the minimization procedure for a lattice of 12×12 sites. After simulated annealing, terminated once the acceptance rate dropped below $R = 0.001\%$

(at $T \approx 0.005J$), the correct ordering pattern is already visible, although the vectors remain somewhat spread out. After only $N_o = 10$ additional optimization sweeps, however, both spin and valley order expectation values align perfectly to the ground-state order. This demonstrates the efficiency of stochastic gradient descent for refining ground-state configurations once simulated annealing has identified the correct ordering pattern.

3.3 Cluster mean-field theory

In the classical and semiclassical approximations discussed earlier, the complexity of the quantum many-body problem was drastically reduced by forbidding entanglement between spins on different sites. This simplification fails whenever entanglement or unconventional order plays a central role, as in quantum spin liquids or dimerized phases. While our primary tool for analyzing such phases is the pf-FRG which we will describe in Chapter 4, we complement it with a conceptually and computationally simpler approach called cluster mean-field theory (CMFT), which can likewise distinguish between magnetic and paramagnetic ground states, although with some limitations.

Unlike the semiclassical limits who eliminating entanglement entirely, CMFT solves the full quantum problem on small spin clusters using exact diagonalization (ED), while approximating inter-cluster interactions via a mean-field decoupling. Compared to conventional ED, CMFT has the advantage that spontaneous symmetry breaking can be observed explicitly due to the symmetry breaking induced by the mean-fields. The price, however, is that the $SU(2)$ symmetry of the model is also fully broken in the ED calculation. For highly symmetric models, this significantly reduces the accessible cluster sizes compared to pure ED, which can utilize the symmetries of the model. Despite these limitations, CMFT has been successfully applied, for example, to determine phase boundaries between ordered and paramagnetic phases in the J_1 – J_2 Heisenberg model on the square lattice [105], or to capture different ordered states in the XXZ model on the triangular lattice [106]. In both cases, it yields qualitatively reliable results even with relatively modest cluster sizes.

In the following, we describe the practical steps of implementing the CMFT. We first outline how the approximation leads to a set of self-consistent equations for the cluster magnetizations, which we solve using an iterative procedure. We then discuss the role of initial conditions and cluster geometry, both of which crucially affect convergence and the phases that can be detected. To make the method accessible, we have implemented it in a Julia package `ClusterMeanFieldTheory.jl` [C2]. As an illustrative example and benchmark, we apply CMFT to the J_1 – J_2 Heisenberg model on the square lattice. The description of the method itself is largely based on Ref. [U1].

3.3.1 Cluster mean-field approximation

The CMFT approach begins by partitioning the full lattice on which the spin model is defined into small spin clusters \mathcal{C} , each containing N_C sites. Interactions between spins within a cluster are treated exactly. Interactions between spins in different clusters are approximated by the standard mean-field decoupling scheme. Formally, this is achieved by expanding the spin operators in small perturbations around their expectation values

$$\delta \mathbf{S}_i \equiv \mathbf{S}_i - \langle \mathbf{S}_i \rangle \tag{3.65}$$

and keeping terms up to first order, which yields

$$\begin{aligned}\mathbf{S}_i \cdot \mathbf{S}_j &= (\langle \mathbf{S}_i \rangle + \delta \mathbf{S}_i) \cdot (\langle \mathbf{S}_j \rangle + \delta \mathbf{S}_j) \\ &= \langle \mathbf{S}_i \rangle \cdot \mathbf{S}_j + \mathbf{S}_i \cdot \langle \mathbf{S}_j \rangle - \langle \mathbf{S}_i \rangle \cdot \langle \mathbf{S}_j \rangle + O(\delta \mathbf{S}^2).\end{aligned}\quad (3.66)$$

This is, of course, only a valid approximation if fluctuations away from the expectation value are indeed small. If we assume this is the case, a general Heisenberg Hamiltonian

$$H = \frac{1}{2} \sum_{ij} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \quad (3.67)$$

can be recast as a sum over single-cluster Hamiltonians $H_{\mathcal{C}}$

$$\begin{aligned}H &= \sum_{\mathcal{C}} \left(\sum_{i,j \in \mathcal{C}} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j + \sum_{\mathcal{C}' \neq \mathcal{C}} \sum_{i \in \mathcal{C}, j \in \mathcal{C}'} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \right) \\ &\approx \sum_{\mathcal{C}} \left(\sum_{i,j \in \mathcal{C}} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j + \sum_{i \in \mathcal{C}} \mathbf{h}_i \cdot \mathbf{S}_i + C_{\mathcal{C}} \right) \\ &=: \sum_{\mathcal{C}} H_{\mathcal{C}}.\end{aligned}\quad (3.68)$$

where the mean influence of neighboring clusters is captured within effective fields⁴

$$\mathbf{h}_i^{\mathcal{C}} = \frac{1}{2} \sum_{\mathcal{C}' \neq \mathcal{C}} \sum_{j \in \mathcal{C}'} (J_{ij} + J_{ji}) \langle \mathbf{S}_j \rangle \quad (3.69)$$

and the energy shift

$$C^{\mathcal{C}} = \frac{1}{2} \sum_{\mathcal{C}' \neq \mathcal{C}} \sum_{i \in \mathcal{C}, j \in \mathcal{C}'} (J_{ij} + J_{ji}) \langle \mathbf{S}_i \rangle \langle \mathbf{S}_j \rangle. \quad (3.70)$$

These both still depend on the expectation values of the spins on different clusters $\langle \mathbf{S}_j \rangle$. For the single-cluster Hamiltonians to fully decouple, we impose *periodic boundary conditions* for the magnetizations by assuming that the magnetization patterns repeat identically across all clusters. To this end, we split the site index into $i = (\mathcal{C}, \alpha)$, where \mathcal{C} denotes the cluster and $\alpha = 1, \dots, N_{\mathcal{C}}$ the site within this cluster. Periodic boundary conditions then imply

$$\langle \mathbf{S}_{\mathcal{C}\alpha} \rangle = \langle \mathbf{S}_{\mathcal{C}'\alpha} \rangle \equiv \mathbf{m}_{\alpha}, \quad (3.71)$$

to hold for all clusters $\mathcal{C}, \mathcal{C}'$. This defines the $N_{\mathcal{C}}$ cluster-independent magnetizations \mathbf{m}_{α} , from which the now cluster-independent effective fields \mathbf{h}_{α} and the energy shift C can be calculated using the formulas above.

Restricting the problem to $H_{\mathcal{C}}$ instead of the full lattice Hamiltonian constitutes the cluster mean-field approximation. In the limit $N_{\mathcal{C}} \rightarrow \infty$ this becomes exact, while for $N_{\mathcal{C}} \rightarrow 1$ it reduces to standard mean-field theory. For intermediate cluster sizes, finite size effects remain, but even for moderate cluster sizes qualitatively correct results have been obtained [105, 106].

⁴The site-exchange term J_{ji} is typically accounted for in Heisenberg models by defining $J_{ij} = J_{ji}$ and summing only over $i < j$. For more general spin models, e.g. with DM interactions, only $J_{ij} = J_{ji}^T$ holds and this has to be accounted for to obtain a Hermitian Hamiltonian matrix. To keep the discussion general, we include the term here.

3.3.2 Iterative solution of the self-consistent equations

Although finite temperature CMFT implementations are possible, our interest here is in approximating the ground state. This amounts to solving

$$|\psi_0\rangle = \underset{\psi}{\operatorname{argmin}} \langle \psi | H_C | \psi \rangle . \quad (3.72)$$

If the cluster is small enough, this can simply be achieved by numerically diagonalizing H_C and determining the eigenstates with the lowest eigenvalue. In practice, we do this using the Lanczos algorithm [107] implemented in the `KrylovKit.jl` Julia package [108]. For a thorough explanation of how to set up the Hamiltonian matrix numerically, we highly recommend Ref. [54].

Because we are interested in the ground state, we are, by definition, considering the system at zero temperature. Expectation values, therefore, must be taken with respect to the ground state, $\langle \dots \rangle = \langle \psi_0 | \dots | \psi_0 \rangle$. In particular, the magnetizations \mathbf{m}_α depend directly on $|\psi_0\rangle$, while $|\psi_0\rangle$ itself depends on the Hamiltonian H_C , which again depends on \mathbf{m}_α . The definition of the magnetization (3.71) therefore constitutes a self-consistent equation, implying we need to self-consistently determine the magnetization in addition to the ground state.

To achieve this, we perform a damped fixed-point iteration. Starting with some initial magnetizations \mathbf{m}_α^0 , at each iteration step n , an updated set of magnetizations \mathbf{m}_α^{n+1} is computed based on the values from the previous step \mathbf{m}_α^n as follows:

1. Calculate the effective fields \mathbf{h}_α from \mathbf{m}_α^n using Eq. (3.69).
2. Determine the ground state of the resulting single-cluster Hamiltonian H_C via the Lanczos algorithm.
3. Calculate the magnetizations

$$\mathbf{m}_\alpha^{\text{new}} = \langle \mathbf{S}_\alpha \rangle \quad (3.73)$$

in this ground state.

4. Update the magnetizations according to

$$\mathbf{m}_\alpha^{n+1} = (1 - \lambda) \mathbf{m}_\alpha^{\text{new}} + \lambda \mathbf{m}_\alpha^n ,$$

where $\lambda \in (0, 1]$ is a damping parameter introduced to stabilize convergence.

5. Check for convergence: Stop the iteration if

$$\sum_\alpha |\mathbf{m}_\alpha^{n+1} - \mathbf{m}_\alpha^n| < \epsilon_{\text{abs}} ,$$

for a desired absolute tolerance ϵ_{abs} , otherwise continue with the next step.

In practice, we have found that setting the damping parameter to $\lambda = 0.5$ achieves a good compromise between fast convergence and stability of the method. Without damping (i.e. $\lambda = 1$), we found that the iterations often fail to converge at all. For the absolute tolerance we typically set $\epsilon_{\text{abs}} = 10^{-8}$. Once convergence is reached, we calculate the final ground state using the magnetizations of the last iteration step. From this ground state, various observables can then be straight-forwardly computed.

3.3.3 Initial conditions

To reduce the risk of converging to a local rather than global minimum, we repeat the CMFT calculation with a variety of multiple distinct initial magnetizations \mathbf{m}_α^0 . Among the resulting solutions, we then use the ground state with the lowest energy as our final CMFT estimate for the actual ground state. To test whether CMFT predicts a magnetically ordered or paramagnetic ground state, we always include a calculation with paramagnetic initial conditions $\mathbf{m}_\alpha^0 = 0$ (for which no iteration is required as all fields are always zero).

When the ground state of the corresponding classical spin model is known, we also initialize the magnetization pattern according to this classical state. To account for the reduction of the ordered moment expected from quantum fluctuations, we normalize the initial values to $|\mathbf{m}_\alpha^0| = 1/4$.

In cases where the classical ground state is unknown, or where for finite size clusters the ground state might differ from the ground state in the thermodynamic limit, we instead construct a physically motivated initial state from the real-space Luttinger-Tisza approach as described in Sec. 3.1. To be as close to the true classical ground state as possible, we choose a configuration that minimizes the deviation from the strong constraint by minimizing Eq. (3.16) numerically. Implementing this for general clusters also has the advantage that the classical ground states don't have to be implemented "by-hand" every time a new region of the coupling space or a different model is considered (so long as the ground-state is captured by Luttinger-Tisza). Other, numerically more involved approaches are to obtain the classical ground state by Monte Carlo methods like simulated annealing [95] or iterative minimization [100].

Finally, we also include several runs with completely random initial conditions. Although these typically converge more slowly, they provide an additional check that the obtained solution indeed has the lowest accessible energy.

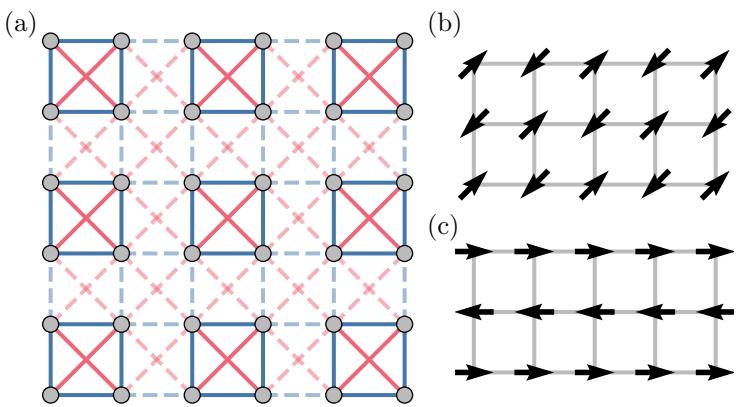
3.3.4 Limitations in the cluster geometry

The effective fields in the single-cluster Hamiltonian completely break the global $SU(2)$ spin symmetry, necessitating the diagonalization of the full $2^{N_c} \times 2^{N_c}$ Hamiltonian matrix in each iteration. As a result, CMFT is more computationally intensive than pure exact diagonalization studies and restricts us to smaller clusters. In our implementation, we find the largest feasible cluster size to be around $N_c = 18$.

Cluster choice is further constrained by the requirement of periodic boundary conditions, which restricts us to symmetry equivalent clusters that tile the full lattice. For magnetically ordered ground states, clusters must additionally be commensurate with the ordering pattern. For paramagnetic states, by contrast, all effective fields vanish ($\mathbf{h}_i = 0$), and CMFT reduces to ED with open boundaries, yielding ground states that respect all spin and cluster symmetries. To access nonmagnetic states that break lattice symmetries—such as dimer or nematic phases—one must explicitly break the symmetry, either through the choice of cluster geometry or by introducing small biases into the couplings J_{ij} . This, however, requires prior knowledge of the candidate symmetry-breaking patterns.

In general, CMFT tends to favor magnetically ordered states, as a paramagnetic state with vanishing magnetizations has zero effective fields and thus lacks a mechanism to lower the energy in the mean-field interactions. One possible remedy is to perform finite-size scaling by extrapolating results with increasing N_c towards the thermodynamic limit [106, 109]. In practice, however, the restrictions on cluster geometries often make systematic finite-size scaling difficult: only a limited set of clusters is allowed, and the maximal accessible cluster size remains relatively small.

Figure 3.2 –
 $J_1 - J_2$ model on the square lattice (a) Square lattice with nearest-neighbor bonds shown in blue (J_1) and next-nearest-neighbor bonds in red (J_2). For an $L = 2 \times 2$ cluster in CMFT, solid lines denote bonds treated exactly within the cluster, while dashed lines indicate bonds approximated at the mean-field level. (b) Néel order and (c) stripe order, which are the classical ground states for $J_2/J_1 < 0.5$ and $J_2/J_1 > 0.5$, respectively.



As a result, finite-size effects are expected to be significant, and fully quantitative predictions are often out of reach. Nevertheless, by comparing results across different cluster geometries and focusing on phases that appear consistently, CMFT may still provide a robust qualitative picture of the phase diagram.

3.3.5 Example: The $J_1 - J_2$ Heisenberg model on the square lattice

As an instructive example, we present CMFT results obtained with our `ClusterMeanFieldTheory.jl` package for the $J_1 - J_2$ Heisenberg model on the square lattice [as illustrated in Fig. 3.4], a system previously analyzed with CMFT in Ref. [105], which provides a useful benchmark for comparison.

The Hamiltonian we consider is

$$H = \sum_{\langle ij \rangle_1} J_1 \mathbf{S}_i \mathbf{S}_j + \sum_{\langle ij \rangle_2} J_2 \mathbf{S}_i \mathbf{S}_j, \quad (3.74)$$

where the sums run over nearest neighbors $\langle ij \rangle_1$ and next-nearest neighbors $\langle ij \rangle_2$. We consider AFM $J_1, J_2 > 0$ and calculate the phase diagram as a function of J_2/J_1 .

Classical phases To obtain suitable initial conditions, we first determine the classical phase diagram. Since the square lattice is a Bravais lattice, this can be done analytically using Luttinger-Tisza. The Fourier-transformed coupling matrix from Eq. (3.20) reduces to the scalar function

$$J(\mathbf{q}) = J_1 [\cos q_x + \cos q_y] + J_2 [\cos(q_x + q_y) + \cos(q_x - q_y)], \quad (3.75)$$

which needs to be minimized in the first Brillouin zone $[0, 2\pi]^2$ (with the lattice constant set to unity). This yields two possible symmetry inequivalent extrema $\mathbf{q}_{\pi,\pi} = (\pi, \pi)^T$ and $\mathbf{q}_{\pi,0} = (\pi, 0)^T$, corresponding to Néel and stripe order shown in Fig. 3.2(b,c), respectively. Their energy per site $\epsilon = E/N$ is equal to $J(\mathbf{q})$ evaluated at these wave-vectors and reads

$$\epsilon_{\text{Néel}} = J(\mathbf{q}_{\pi,\pi}) = -2J_1 + 2J_2, \quad \epsilon_{\text{stripe}} = J(\mathbf{q}_{\pi,0}) = -2J_2, \quad (3.76)$$

implying that the classical ground state exhibits Néel order for $J_2/J_1 < 0.5$ and stripe order for $J_2/J_1 > 0.5$.

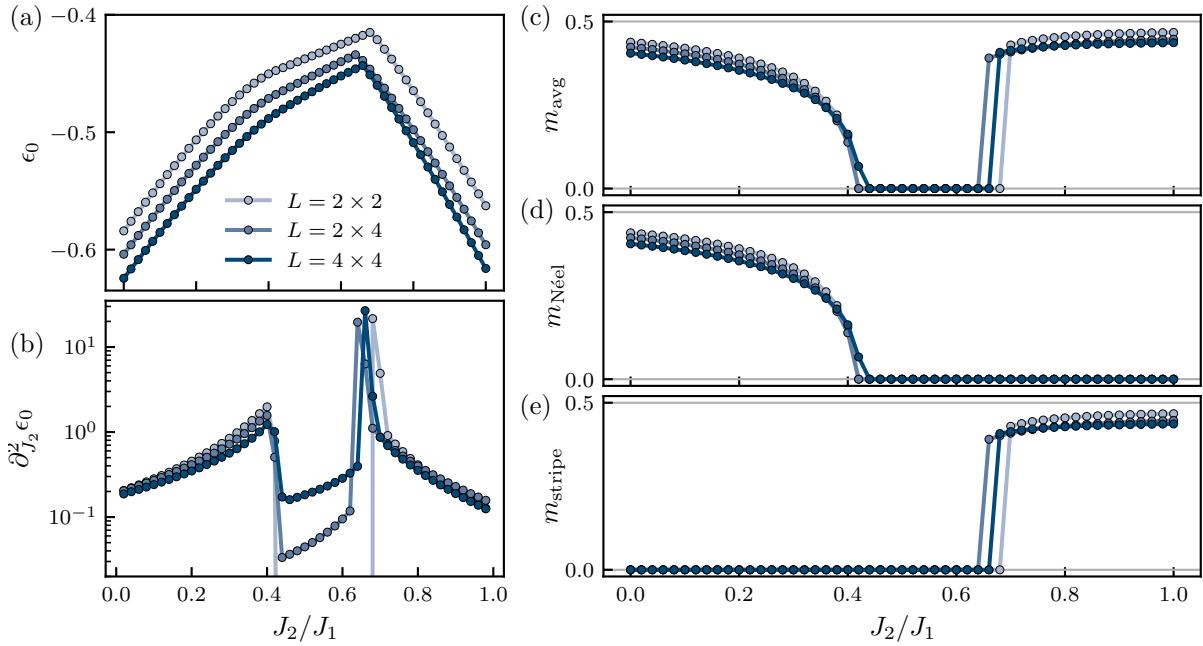


Figure 3.3 – CMFT results for different cluster sizes. (a) Ground-state energy per site ϵ_0 . (b) Second derivative of ϵ_0 with respect to J_2 , showing peaks at $J_2^{c1} \approx 0.41 J_1$ and $J_2^{c2} \approx 0.67 J_1$, which signal phase transitions. (c) Average site magnetization, m_{avg} , revealing a paramagnetic regime for $J_2^{c1} < J_2 < J_2^{c2}$. (d, e) Néel (stripe) order parameter, corresponding to the AFM phases shown in Fig. 3.2(b,c), which dominate at small (large) J_2/J_1 .

CMFT phase diagram We now apply CMFT for clusters of size $L = 2 \times 2, 2 \times 4, 4 \times 4$. We only consider even cluster sizes in both lattice directions so that the clusters are commensurate with the ordered phases. For each cluster, we perform calculations starting from Néel, stripe, paramagnetic, and random initial conditions, and select the converged solution with the lowest energy as the final ground state.

To distinguish between magnetically ordered and paramagnetic states, we compute the average magnetization

$$m_{\text{avg}} = \frac{1}{N_C} \sum_{i \in \mathcal{C}} |\langle \mathbf{S}_i \rangle|, \quad (3.77)$$

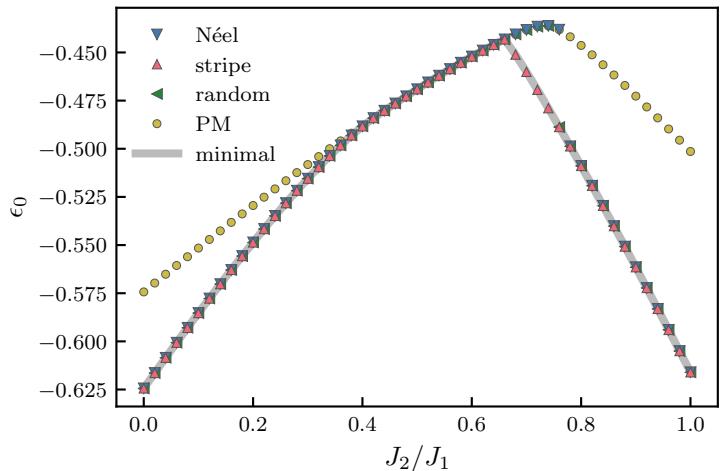
which vanishes only if all local magnetizations are zero, and saturates at its maximum value $m_{\text{avg}} = 1/2$ for a pure product state. While the precise magnitude of the magnetizations can vary slightly from site to site (spins on the boundary of the cluster often show larger magnetizations as they are mean-field coupled to other spins), in practice we always observe the same qualitative behavior—either all magnetizations vanish, or all remain finite.

To further characterize the ordered phases, we define order parameters for the Néel and stripe states as

$$m_{\text{Néel}} = \frac{1}{N_C} \left| \sum_{i \in \mathcal{C}} \cos(\mathbf{q}_{\pi, \pi} \cdot \mathbf{R}_i) \langle \mathbf{S}_i \rangle \right|, \quad m_{\text{stripe}} = \frac{1}{N_C} \left| \sum_{i \in \mathcal{C}} \cos(\mathbf{q}_{\pi, 0} \cdot \mathbf{R}_i) \langle \mathbf{S}_i \rangle \right|, \quad (3.78)$$

where the factors $\cos(\mathbf{q} \cdot \mathbf{R}_i) = \pm 1$ capture the staggered sign of the magnetization. We also calculate the ground-state energy per site $\epsilon_0 = H/N_C$ and its second derivative with respect to J_2/J_1 as an additional indicator of possible phase transitions.

Figure 3.4 –
Ground-state energy for different initial conditions. Converged CMFT results for the $L = 4 \times 4$ cluster using various initial magnetizations \mathbf{m}_α^0 in the fixed-point iteration. Markers indicate individual runs, while the grey line traces the lowest-energy solution. Close to the first order transition from the paramagnet to stripe order, only stripe initial conditions correctly capture the ground state.



The results for all cluster sizes are shown in Fig. 3.3. For small J_2/J_1 , CMFT predicts Néel order, and for large J_2/J_1 it finds stripe order, in agreement with the classical phase diagram. At intermediate J_2/J_1 , however, the average magnetization vanishes, indicating a quantum paramagnetic regime. For the $L = 4 \times 4$ cluster, this occurs between $J_2^{c1} \approx 0.41J_1$ and $J_2^{c2} \approx 0.67J_1$, consistent with results from CMFT in Ref. [105]. The precise location of these critical couplings has been intensely studied, yet significant variation remains across methods. A recent state-of-the-art DMRG study finds a much narrower paramagnetic window, $J_2^{c1} = 0.535(3)J_1$ to $J_2^{c2} = 0.610(5)J_1$ [110], while other approaches such as PEPS [111], DMRG [112], and ED [113] report values closer to our CMFT results. Another open question concerns the character of the quantum paramagnetic regime itself: while some of these studies identify both a QSL and a VBS phase, others find only a VBS. Within CMFT, however, we cannot reliably distinguish between these states.

We can, however, make estimates on the nature of the phase transitions by analyzing both the magnetization and the energy as a function of J_2/J_1 shown in Fig. 3.3. At the Néel to PM transition, the energy exhibits a smooth hump, which shows up as a peak in the second derivative $\partial_{J_2}^2 \epsilon_0$, while the magnetization continuously drops to zero. This behavior indicates a continuous, or a weakly first-order transition. In contrast, at the PM to stripe transition, the energy displays a sharp kink, the second derivative shows a pronounced peak, and the magnetization jumps discontinuously to a finite value—all indicators for a first-order transition. These findings are consistent with previous studies [110–113].

In addition, we observe that the magnitude of the magnetization in the ordered phases decreases systematically with increasing cluster size. This reflects the natural enhancement of quantum fluctuations as the system approaches the thermodynamic limit.

To emphasize our point that choosing the correct initial conditions is crucial in CMFT, Fig. 3.4 shows *converged* solutions obtained with different initial magnetizations \mathbf{m}_α^0 . Especially in the stripe phase close to the phase transition to the paramagnetic phase, results converge to the correct stationary state only if the initial conditions already reflect the stripe ordering pattern. This sensitivity provides further evidence for phase coexistence at the transition and thus supports the interpretation of a first-order transition.

To conclude, although CMFT does not yield quantitatively precise phase boundaries, it is remarkable that such a numerically inexpensive method (typically requiring well below a single core-hour per data point on a modern computer) can produce phase diagrams of frustrated magnetic models that qualitatively agree with results from far more computationally demanding

approaches.

Chapter 4

The pseudo-fermion functional renormalization group

In the previous chapter we introduced several methods to analyze frustrated spin models. All of these approaches involved approximations to the quantum many-body problem that restricted the amount of long-range entanglement taken into account. In principle, however, we exactly want to study phases where quantum fluctuations lead to such long-range entanglement and associated exotic quantum phases this can produce.

With this aim in mind, in this chapter we present the pseudo-fermion functional renormalization group (pf-FRG), first introduced by Reuther and Wölfle in 2010 [13]. Since then, it has become an established method for exploring the competition between magnetically ordered states and nonmagnetic quantum phases in a broad range of highly frustrated spin models and materials, both in two and three dimensions [14]. The pf-FRG rests on two key steps: (i) representing spin operators in terms of complex Abrikosov fermions [64], as already discussed in Eq. (1.22), and (ii) applying the fermionic functional renormalization group (FRG) to the resulting pseudo-fermion Hamiltonian. The central idea of the FRG is to avoid treating all energy scales simultaneously. Instead, one starts from a known high-energy limit and gradually incorporates fluctuations at lower and lower energy scales. Technically, this is achieved by deriving exact differential equations for correlation functions, which are then integrated from high to low energies. This framework enables new approximation schemes that remain effective even for spin models where most conventional methods break down.

In this chapter, we begin by introducing the general framework of the fermionic FRG in Sec. 4.1, before turning to its specialization for spin models in Sec. 4.2. Because the literature already covers many technical details, we aim to keep the exposition here concise. For background on the general FRG, we recommend the textbook by Kopietz *et al.* [114]. For the pf-FRG, a short but comprehensive summary can be found in a recent review [14]. For more details on the many involved calculations, we recommend several excellent PhD thesis [84, 115, 116]. Much of this chapter is largely based on those works.

We continue the chapter by focusing on the methodological developments contributed during my PhD (and partly already in my master’s thesis). The first such development, discussed in Sec. 4.3, is the generalization of the pf-FRG from SU(2) to SU(4) models, relevant for spin-valley or spin-orbital systems as motivated in the introductory chapters and based on Ref. [P3]¹. The second development is the extension of our pf-FRG implementation to spin-orbit coupled materials with anisotropic, nondiagonal interactions. In this setting, the systematic exploitation of combined spin and real-space symmetries becomes crucial. While related implementations already exist [117] and such models have been studied with pf-FRG [118–122], our version is, to

¹Although Ref. [P3] was published during my PhD, much of the conceptual work had already been carried out during my master’s thesis, where a similar analysis is presented.

the best of our knowledge, the most general algorithm for essentially arbitrary crystal lattices that is publicly available—realized within the `PFFRGSolver.jl` Julia package [C3], to which I contributed during this thesis. In Sec. 4.4, we present this algorithm in detail. In Sec. 4.5, we then turn to the more practical aspects of how to actually calculate phase diagrams from the output of the pf-FRG. We close the chapter with an assessment of pf-FRG’s strengths and limitations, combined with outlook on future directions for diagrammatic approaches to quantum spin models.

4.1 The fermionic functional renormalization group

We begin by reviewing the FRG for purely fermionic systems. We closely follow the notation of Ref. [123], but skip the more technical steps of the derivation. Far more detailed discussions are given in textbooks such as Ref. [114].

To set the stage, we consider the general fermionic Hamiltonian $H = H_0 + H_{\text{int}}$ with the kinetic contribution

$$H_0 = \sum_{x'x} h_{x',x} f_{x'}^\dagger f_x, \quad (4.1)$$

and a quartic interaction of the form

$$H_{\text{int}} = \frac{1}{4} \sum_{x'_1, x'_2, x_1, x_2} v_{x'_1, x'_2, x_1, x_2} f_{x'_1}^\dagger f_{x'_2}^\dagger f_{x_1}^\dagger f_{x_2}, \quad (4.2)$$

where $f_{x_i}^\dagger$ and f_{x_i} denote fermionic creation and annihilation operators, and x_i is a multi-index capturing all relevant quantum numbers (for example momentum and spin). Our final objective is to calculate n -particle (or equivalently $2n$ -point) correlation functions of the form

$$G_n(x'_1 \tau'_1, \dots, x'_n \tau'_n; x_1 \tau_1, \dots, x_n \tau_n) = (-1)^n \left\langle T_\tau \left[f_{x_1}(\tau_1) \dots f_{x_n}(\tau_n) f_{x'_n}^\dagger(\tau'_n) \dots f_{x'_1}^\dagger(\tau'_1) \right] \right\rangle, \quad (4.3)$$

where

$$f_x(\tau) = e^{\tau H} f_x e^{-\tau H}, \quad f_x^\dagger(\tau) = e^{-\tau H} f_x^\dagger e^{\tau H} \quad (4.4)$$

are the fermionic operators evaluated at imaginary time $\tau \in [0, \beta)$ in the Heisenberg picture, and T_τ is the imaginary-time-ordering operator. The thermal average $\langle \dots \rangle$ is taken with respect to the partition function $Z = \text{Tr} \exp(-\beta H)$ for the temperature $T = 1/\beta$. From such correlation functions, essentially all physical observables can be constructed.

The FRG works in the path integral framework, where the trace over all states can be replaced by an, in principle, infinite dimensional path integral over Grassmann fields $\bar{\psi}_{x_1}$ and ψ_{x_1} . In this framework, the quantum partition function can be rewritten as

$$Z = \text{Tr} e^{-\beta H} = \int \mathcal{D}(\bar{\psi}, \psi) \exp\{-S(\bar{\psi}, \psi)\}, \quad (4.5)$$

where the action $S = S_0 + S_{\text{int}}$ now takes the role of the Hamiltonian in defining the field-theory. Performing a Fourier transform to Matsubara frequencies, the dependence on τ is replaced by a Matsubara index ω . Including this index in the multi-indices x_i , the noninteracting part of the action S_0 is defined as

$$S_0(\bar{\psi}, \psi) = -(\bar{\psi}, G_0^{-1}\psi) = - \sum_{x'_1, x_1} \bar{\psi}_{x'_1} (-i\omega \delta_{x'_1, x_1} + h_{x_1, x'_1}) \psi_{x_1} \delta_{\omega, \omega'}, \quad (4.6)$$

where the notation (\cdot, \cdot) symbolizes the sum over internal indices, including all normalization factors, as e.g. $1/\beta$ for the sum over Matsubara frequencies, and

$$G_0 = (i\omega\delta_{k',k} - h_{k',k})^{-1} \quad (4.7)$$

is the bare propagator. The interaction part of the action is given by

$$S_{\text{int}}(\bar{\psi}, \psi) = \frac{1}{4} \sum_{x'_1, x'_2, x_1, x_2} v_{x'_1, x'_2, x_1, x_2} \bar{\psi}_{x'_1} \bar{\psi}_{x'_2} \psi_{x_1} \psi_{x_2} \delta_{\omega'_1 + \omega_1, \omega'_2 + \omega_2}. \quad (4.8)$$

In this formalism, the n -particle correlation functions can then be calculated as the path integral

$$G_n(x'_1, \dots, x'_n; x_1, \dots, x_2) = \langle \bar{\psi}_{x'_1} \dots \bar{\psi}_{x'_n} \psi_{x_n} \dots \psi_{x_1} \rangle \quad (4.9)$$

$$= \frac{1}{Z} \int \mathcal{D}(\bar{\psi}, \psi) \bar{\psi}_{x'_1} \dots \bar{\psi}_{x'_n} \psi_{x_n} \dots \psi_{x_1} \exp\{-S(\bar{\psi}, \psi)\}. \quad (4.10)$$

The idea of the FRG is now the following: We introduce an infrared cutoff Λ into the theory, which allows us to derive an exact hierarchy of flow equations describing how correlation functions evolve with Λ . The cutoff Λ is chosen such that correlation functions are analytically known in the limit $\Lambda \rightarrow \infty$, while it vanishes and leaves the full interacting theory unchanged as $\Lambda \rightarrow 0$. Integrating the flow equations from the known limit of $\Lambda \rightarrow \infty$ to $\Lambda \rightarrow 0$, we can obtain the full correlation functions. In practice, however, approximations are inevitable to make any real calculation feasible. Moreover, instead of working directly with the n -particle correlation functions, it is more convenient to formulate the flow in terms of one-particle irreducible (1PI) correlation functions, or *vertices*. These vertices serve as the fundamental building blocks from which all correlation functions can be reconstructed. In the following, we first define the 1PI vertices and then briefly sketch the derivation of the general flow equations.

4.1.1 One-particle irreducible correlation functions

To define the 1PI vertices, we first introduce a generating functional for the full (disconnected) correlation functions as

$$W(\bar{\eta}, \eta) = \frac{1}{Z} \int \mathcal{D}(\bar{\psi}, \psi) e^{S(\bar{\psi}, \psi) - (\bar{\psi}, \eta) - (\bar{\eta}, \psi)}, \quad (4.11)$$

where $\bar{\eta}$ and η are external Grassmann source fields. Correlation functions are obtained via functional derivatives as

$$G_n(x'_1, \dots, x'_m; x_1, \dots, x_m) = \frac{\delta^m}{\delta \bar{\eta}_{x'_1} \dots \delta \bar{\eta}_{x'_m}} \frac{\delta^m}{\delta \eta_{x_m} \dots \delta \eta_{x_1}} W(\bar{\eta}, \eta) \Big|_{\bar{\eta}=\eta=0}. \quad (4.12)$$

These functions still contain a lot of redundancy, since they include disconnected diagrams that simply factorize into lower-order correlators. To exclude these, we define the generating functional for only the *connected* correlation functions as

$$W^c(\bar{\eta}, \eta) = \ln [W(\bar{\eta}, \eta)]. \quad (4.13)$$

Even the connected correlation functions contain redundant information in the form of diagrams that are one-particle *reducible*, which means they can be separated into two disconnected parts

by removing a single propagator-line $G \equiv G_1$. The one-particle *irreducible* building blocks are the vertices, generated by the functional Legendre transform

$$\Gamma(\bar{\varphi}, \varphi) = -W^c(\bar{\eta}, \eta) - (\bar{\varphi}, \eta) - (\bar{\eta}, \varphi), \quad (4.14)$$

with conjugate fields defined as

$$\bar{\varphi} = \frac{\delta W^c}{\delta \eta} \quad \text{and} \quad \varphi = \frac{\delta W^c}{\delta \bar{\eta}}. \quad (4.15)$$

The n -particle vertex Γ_n is then obtained by functional derivatives of Γ , just as connected correlators follow from W^c .

All connected correlation functions can be reconstructed from the 1PI vertices by the *tree expansion*, a diagrammatic construction that introduces no loop structures [114]. This means the internal loop structure of the theory is fully encoded in the vertices. For instance, for the one-particle correlation function $G := G_1^{(c)} = G_1$ the tree expansion yields the Dyson equation

$$G := [G_0^{-1} + \Gamma_1]^{-1} \equiv [G_0^{-1} - \Sigma]^{-1}, \quad (4.16)$$

so that $\Gamma_1 = -\Sigma$ is identified with the self-energy. Similarly, the connected two-particle correlator relates to the two-particle vertex as

$$G^c(x'_1, x'_2; x_1, x_2) = - \sum_{x_3 x_4 x_5 x_6} \Gamma_2(x_3, x_4; x_5, x_6) G(x'_1; x'_3) G(x'_2; x'_4) G(x'_5; x'_1) G(x'_6; x'_2). \quad (4.17)$$

The vertices can be understood as effective n -body interactions, generated from the bare interaction of the classical action by systematically incorporating quantum fluctuations. For this reason, the functional Γ is also referred to as the *effective action*.

4.1.2 Flow equations

To derive the flow equations for the vertices, we introduce an infrared cutoff Λ by modifying the bare propagator, $G_0 \rightarrow G_0^\Lambda$, such that it satisfies the boundary conditions

$$G_0^{\Lambda=0} = G_0 \quad G_0^{\Lambda \rightarrow \infty} = 0. \quad (4.18)$$

For the derivation, the precise Λ -dependence is not important, as long as these boundary conditions are fulfilled. Consequently, all correlation functions, including the vertices, acquire a Λ -dependence. Importantly, in the limit $\Lambda \rightarrow \infty$ they become particularly simple, reducing to their bare values which in our case are

$$\begin{aligned} \Gamma_2^{\Lambda \rightarrow \infty}(x'_1, x'_2; x_1, x_2) &= v_{x'_1, x'_2, x_1, x_2}, \\ \Gamma_m^{\Lambda \rightarrow \infty}(x'_1, \dots, x'_m; x_1, \dots, x_m) &= 0 \quad \text{for } m \neq 3 \end{aligned} \quad (4.19)$$

This can be rigorously derived [123], but also intuitively understood if one is familiar with the diagrammatic expansion of correlation functions: if the bare propagator G_0 is zero, then all lines between vertices in all diagrams are zero as well and only the bare vertices remain. The effective action therefore becomes the bare action in this limit.

In the limit of $\Lambda \rightarrow 0$, on the other hand, the full vertices $\Gamma_m^{\Lambda \rightarrow 0} = \Gamma_m$ are recovered. The behavior for intermediate Λ is governed by the flow equations. Starting from the definition of the effective action Γ , one can derive these flow equations by performing Λ -derivatives on both

sites, which yields an *exact* differential equation for the full effective action (whose $\Lambda \rightarrow \infty$ limit is simply the classical action). More useful for practical purposes, is to formulate flow equations for the individual m -particle vertices by expanding the effective action as

$$\Gamma^\Lambda(\bar{\varphi}, \varphi) = \sum_m \frac{(-1)^m}{(m!)^2} \sum_{x'_1, \dots, x'_m} \sum_{x_1, \dots, x_m} \Gamma_m^\Lambda(x'_1, \dots, x'_m; x_1, \dots, x_m) \bar{\varphi}_{x'_1} \dots \bar{\varphi}_{x'_m} \varphi_{x_m} \dots \varphi_{x_1} \quad (4.20)$$

This leads to an infinite hierarchy of ordinary first-order differential equations of the form

$$\frac{d}{d\Lambda} \Gamma_n^\Lambda = \mathcal{F}^\Lambda(\Gamma_1^\Lambda, \dots, \Gamma_{n+1}^\Lambda). \quad (4.21)$$

As each m -particle vertex depends on the $m+1$ particle vertex, any closed solution requires truncations of this hierarchy. Already the three-particle vertex depends on six of the multi-indices x_i , which usually contain a lattice site (or momentum), a frequency (or time), spin and sometimes orbital indices. Without applying drastic approximations, even the memory requirements of storing such a six-dimensional tensor is numerically infeasible for all but the smallest systems. In the context of the pf-FRG, only the one-particle vertex, i.e., the self-energy $\Gamma_1 \equiv \Sigma$, and the two-particle vertex are treated explicitly, while higher-order vertices are included only approximately by different truncation schemes discussed in the next section. From now on, we drop the subscript and denote the two-particle vertex simply as $\Gamma \equiv \Gamma_2$, and refer to it as just the *vertex*.

Explicitly, the general flow equation for the self-energy reads

$$\frac{d}{d\Lambda} \Sigma^\Lambda(x'_1, x_1) = - \sum_{x'_2 x_2} \Gamma_2^\Lambda(x'_1, x'_2; x_1, x_2) S^\Lambda(x_2; x'_2) \quad (4.22)$$

and for the two-particle vertex

$$\begin{aligned} \frac{d}{d\Lambda} \Gamma_2^\Lambda(x'_1, x'_2; x_1, x_2) &= \sum_{x'_3 x_3} \Gamma_3^\Lambda(x'_1, x'_2, x'_3; x_1, x_2, x_3) S^\Lambda(x_3; x'_3) \\ &+ \sum_{x'_3 x_3, x'_4 x_4} [\Gamma_2^\Lambda(x'_1, x'_2; x_3, x_4) \Gamma_2^\Lambda(x'_3, x'_4; x_1, x_2) \\ &- \Gamma_2^\Lambda(x'_1, x'_4; x_1, x_3) \Gamma_2^\Lambda(x'_3, x'_2; x_4, x_2) - (x'_3 \leftrightarrow x'_4, x_3 \leftrightarrow x_4) \\ &+ \Gamma_2^\Lambda(x'_2, x'_4; x_1, x_3) \Gamma_2^\Lambda(x'_3, x'_1; x_4, x_2) + (x'_3 \leftrightarrow x'_4, x_3 \leftrightarrow x_4)] \\ &\times G^\Lambda(x_3; x'_3) S^\Lambda(x_4; x'_4). \end{aligned} \quad (4.23)$$

Here, we defined the single scale operator

$$S^\Lambda \equiv G^\Lambda [\partial_\Lambda [G_0^\Lambda]^{-1}] G^\Lambda = -\partial_\Lambda^\Sigma G^\Lambda, \quad (4.24)$$

where for the second step we used the Dyson equation (4.16) and define $\partial_\Lambda^\Sigma \equiv \partial_\Lambda^{\Sigma=const.}$.

4.1.3 Truncation of the flow equations

As already noted, the explicit treatment of the three-particle vertex is numerically infeasible for all but the simplest models. In practice, therefore, all pf-FRG implementations to date truncate the flow equations such that the two-particle vertex flow (4.23) no longer depends explicitly on the three-particle vertex.

The simplest option is to drop this contribution altogether while leaving the remaining flow equations unchanged. This so-called level-2 (L2) truncation can be justified for weakly interacting fermionic models [124, 125], at least in the $\Lambda \rightarrow \infty$ limit. For spin Hamiltonians in the pseudo-fermion representation, which are strongly interacting, the L2 fails to capture quantum paramagnetic phases such as spin liquids or valence-bond solids and only detects magnetically ordered states [13].

The standard choice in pf-FRG, which we also adopt throughout this thesis, is the *Katanin truncation* [126]. Here, the three-particle vertex is again set to zero, but the single-scale propagator is replaced by the Katanin propagator as

$$S^\Lambda = -\partial_\Lambda^\Sigma G^\Lambda \rightarrow S_{\text{kat}}^\Lambda = -\partial_\Lambda^\Sigma G = S^\Lambda + G(\partial_\Lambda \Sigma) G. \quad (4.25)$$

This modification feeds the self-energy flow $\partial_\Lambda \Sigma$ back into the vertex, effectively capturing certain $O(\Gamma^3)$ diagrams that would otherwise be neglected. This can be shown to improve the fulfillment of Ward identities—exact relations between correlation functions imposed by symmetries—thereby enhancing self-consistency of the approach [126]. Within this truncation, pf-FRG has successfully identified both magnetically ordered and paramagnetic phases [13, 14].

More recently, a more elaborate scheme called *multiloop truncation* has been developed for general fermionic models [127, 128] and later implemented in the pf-FRG [129, 130]. In the multiloop scheme, higher-order vertex terms are obtained by reinserting parts of the flow equations back into it-self in an iterative scheme, systematically including higher-loop diagrams. In principle, this converges to the *parquet approximation* [64, 131, 132] (a set of self-consistent equations for the self-energy and vertex), which exactly satisfies one-particle Ward identities and becomes independent of the cutoff scheme (unlike the Katanin truncation). However, the multiloop approach is both conceptually and numerically far more demanding and there has, as of yet, not been a concrete pf-FRG study showing it qualitatively improves the results from the Katanin truncation. We therefore restrict ourselves to the Katanin truncation in the remainder of this thesis.

Even with the Katanin truncation and finite system sizes, the flow equations still involve a huge number of vertex components. Crucially, however, the vertices preserve all Hamiltonian symmetries throughout the flow. Many self-energy and vertex components are therefore symmetry-equivalent, meaning only a reduced set needs to be computed explicitly. Efficient FRG implementations thus rely on systematically identifying these symmetries and their action on the vertex functions—a task we take up in the following section for quantum spin models.

4.2 Pseudo-fermion functional renormalization group for SU(2) spin models

We now show how the general framework of the fermionic FRG can be applied to SU(2) quantum spin models of the general form

$$H = \frac{1}{2} \sum_{ij} \sum_{ab} S_i^a J_{ij}^{ab} S_j^b, \quad (4.26)$$

where we assume J_{ij}^{ab} is real to obtain a Hermitian Hamiltonian. Since the SU(2) spin operators S_i^a are neither bosonic nor fermionic, but instead obey the nontrivial commutation relations of the SU(2) Lie algebra (1.6), developing an FRG scheme directly in terms of spins is technically

demanding—though recent progress has been made [133]. The pf-FRG circumvents this difficulty by representing spins in terms of auxiliary fermions, which allows one to directly apply the fermionic FRG framework developed in the previous section and leverage the extensive approximation schemes and diagrammatic techniques already established for interacting fermionic systems.

In the following, we outline the essential steps of the pf-FRG: mapping spins to fermions, identifying and exploiting symmetries of the resulting pseudo-fermion Hamiltonian, deriving symmetry-constrained vertex parametrizations, and formulating the corresponding flow equations. We then discuss the choice of regulator and the treatment of the single-occupation constraint that arises from the fermion mapping.

4.2.1 Mapping spins to fermions

To represent spin operators via fermions, we use the mapping already employed in Eq. (1.22), originally introduced by Abrikosov [64]. Since it is central to the pf-FRG approach, we state it again here. The spin operators on site i are represented by complex fermions as

$$S_i^a = \frac{1}{2} f_{is'}^\dagger \theta_{s's}^a f_{is}, \quad (4.27)$$

where $s, s' = \uparrow\downarrow$ are the spin-indices, θ^a with $a = 1, 2, 3$ are Pauli matrices, and we assume summation over repeated spin indices. The local spin Hilbert space is two-dimensional, spanned by the states $|\uparrow\rangle$ and $|\downarrow\rangle$. In contrast, the local fermionic Hilbert space is spanned by the *four* states

$$|0\uparrow, 0\downarrow\rangle, |1\uparrow, 0\downarrow\rangle, |0\uparrow, 1\downarrow\rangle, |1\uparrow, 1\downarrow\rangle, \quad (4.28)$$

labeled by $|n_\uparrow, n_\downarrow\rangle$ where n_s is the number of fermions with spin s . To obtain a faithful mapping, the unphysical vacuum and doubly occupied states must be projected out by enforcing the single-occupation constraint

$$n_i = f_{is}^\dagger f_{is} = 1 \quad (4.29)$$

on every lattice site. Strategies to handle this constraint within the FRG framework are discussed in Sec. 4.2.8 below. Rewriting the general spin Hamiltonian (4.26) in terms of pseudo-fermions gives

$$H = \frac{1}{8} \sum_{ij} \sum_{ab} J_{ij}^{ab} \theta_{s'_1 s_1}^a \theta_{s'_2 s_2}^b f_{is'}^\dagger f_{js'_2}^\dagger f_{js_2} f_{is_1}. \quad (4.30)$$

What sets this Hamiltonian apart from usual fermionic models—such as the Hubbard model in Sec. 2.2.3—is the *complete absence of a kinetic term*.

For the following it is important to note that fermionic representation carries an intrinsic redundancy. This becomes most transparent by defining the matrix

$$\mathbf{F}_i = \begin{pmatrix} f_{i\uparrow} & -f_{i\downarrow}^\dagger \\ f_{i\downarrow} & f_{i\uparrow}^\dagger \end{pmatrix}, \quad (4.31)$$

which allows the spin operators in the pseudo-fermion representation to be written as

$$S_i^a = \frac{1}{2} \text{Tr} \theta^a \mathbf{F}_i \mathbf{F}_i^\dagger. \quad (4.32)$$

In this notation it is apparent that the *gauge transformation* $\mathbf{F} \rightarrow \mathbf{F}\mathbf{U}$ leaves the spin operator invariant for any matrix $\mathbf{U} \in \text{SU}(2)$. This transformation mixes unphysical states $|0\uparrow, 0\downarrow\rangle, |1\uparrow, 1\downarrow\rangle$

with the physical ones $|1_\uparrow, 0_\downarrow\rangle, |0_\uparrow, 1_\downarrow\rangle$. Relevant for us are the transformations that only affect the physical states, which are the local U(1) gauge transformation

$$U = \begin{pmatrix} e^{i\phi} & 0 \\ 0 & e^{-i\phi} \end{pmatrix} \quad (4.33)$$

which simply amounts to multiplying the fermionic operators with a phase ϕ . Additionally relevant is a transformation that exchanges which of the two unphysical states plays the role of the vacuum. This effectively maps fermions to holes and is encoded in the \mathbb{Z}_2 transformation

$$U = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad (4.34)$$

typically referred to as *particle-hole conjugation*. Both transformation are gauge redundancies that leave spin operators, and thus any spin Hamiltonian, invariant. Since expectation values, including all correlation functions, must remain invariant, the gauge redundancies—together with the physical symmetries of the Hamiltonian—impose strong constraints on the structure of the vertex functions, as we will discuss in the next section.

4.2.2 Symmetry constraints on the correlation functions

In addition to the gauge redundancy discussed above, the general spin Hamiltonian, and therefore also its pseudo-fermion representation, has several intrinsic physical symmetries. To state all the relevant ones: The Hamiltonian is Hermitian, time-translation invariant (doesn't explicitly depend on τ) and time-reversal symmetric (with time reversal acting on spins as $\mathbf{S} \rightarrow -\mathbf{S}$). Combined with the gauge redundancy, these symmetries can be exploited to constrain the structure of the one- and two-particle correlation functions, and equivalently their vertices. Such constraints are crucial for an efficient numerical treatment of the flow equations.

The general procedure is straightforward: (i) specify how a given symmetry or gauge transformation acts on the fermionic operators, (ii) insert this into the definition of the correlation functions (4.3), and (iii) demand invariance. In Matsubara frequency space, the one- and two-particle correlators are explicitly given by

$$G(x'_1; x_1) = - \int d\tau'_1 d\tau_1 e^{i(\tau'_1 \omega'_1 - \tau_1 \omega_1)} \langle T_\tau f_{i_1 s_1}(\tau_1) f_{i'_1 s'_1}^\dagger(\tau'_1) \rangle, \quad (4.35)$$

$$G(x'_1 x'_2; x_1, x_2) = - \int d\tau'_1 d\tau'_2 d\tau_1 d\tau_2 e^{i(\tau'_1 \omega'_1 + \tau'_2 \omega'_2 - \tau_1 \omega_1 - \tau_2 \omega_2)} \langle T_\tau f_{i_1 s_1}(\tau_1) f_{i_2 s_2}(\tau_2) f_{i'_2 s'_2}^\dagger(\tau'_2) f_{i'_1 s'_1}^\dagger(\tau'_1) \rangle, \quad (4.36)$$

where the multi-index $x = (i, s, \omega)$ collects site, spin, and Matsubara frequency.

As an example, consider the local U(1) gauge redundancy acting as

$$f_{is}^\dagger \xrightarrow{\text{U(1)}} e^{i\phi_i} f_{is}^\dagger \quad f_{is} \xrightarrow{\text{U(1)}} e^{-i\phi_i} f_{is}, \quad (4.37)$$

for an arbitrary local phase ϕ_i . Requiring invariance of $G(x'_1; x_1)$ under this transformation yields

$$G(x'_1; x_1) = G(x'_1; x_1) \exp(i\phi_{i'_1} - i\phi_{i_1}) \implies G(x'_1; x_1) = G(x'_1; x_1) \delta_{i'_1, i_1}, \quad (4.38)$$

which shows that the one-particle correlator depends only on a single site index.

Table 4.1 – Symmetries, gauge transformations, and their action on the fermionic Hilbert space. The transformations U(1) and PH are gauge redundancies introduced by the pseudo-fermion mapping. The other transformations are intrinsic symmetries of the spin Hamiltonian in Eq. (4.26). s_i appearing as a factor should be interpreted as $s_i = \uparrow = 1$ and $s_i = \downarrow = -1$ and $\bar{s}_i = -s_i$ denotes a flipped spin. Some symmetries are accompanied by complex conjugation of all complex numbers.

| Symmetry/Transformation | Label | Action on fermions | | Complex conjugation |
|-------------------------|-------|--|---|---------------------|
| Local U(1) | U(1) | $f_{is}^\dagger \rightarrow e^{-i\phi_i} f_{is}^\dagger$ | $f_{is} \rightarrow e^{i\phi_i} f_{is}$ | ✗ |
| Local particle-hole | PH | $f_{is}^\dagger \rightarrow s f_{i\bar{s}}$ | $f_{is} \rightarrow s f_{i\bar{s}}^\dagger$ | ✗ |
| Time-reversal | TR | $f_{is}^\dagger \rightarrow e^{i\pi s/2} f_{i\bar{s}}^\dagger$ | $f_{is} \rightarrow e^{-i\pi s/2} f_{i\bar{s}}$ | ✓ |
| Lattice | L | $f_{is}^\dagger \rightarrow f_{L(i)s}^\dagger$ | $f_{is} \rightarrow f_{L(i)s}$ | ✗ |
| Hermiticity | H | | ✗ | ✓ |
| Time-translation | TT | | ✗ | ✗ |
| Crossing | X | | ✗ | ✗ |

Analogous constraints can be derived for the two-particle correlator and other symmetries. The action of all symmetries and gauge redundancies is summarized in Table 4.1, while the resulting constraints are listed in Table 4.2. We also include the *crossing symmetry* (X),

$$G(x'_1, x'_2; x_1, x_2) = -G(x'_2, x'_1; x_1, x_2) = -G(x'_1, x'_2; x_2, x_1) = G(x'_2, x'_1; x_2, x_1), \quad (4.39)$$

which follows directly from fermionic anticommutation, or equivalently from the Grassmann functional integral representation (4.9). Detailed derivations can be found, e.g., in Refs. [14, 118].

4.2.3 Vertex parametrization

To efficiently exploit symmetry constraints in pf-FRG, it is convenient to adopt parametrizations of the correlators that already encodes most of them. For the one-particle correlation function, the local U(1) gauge redundancy and imaginary-time-translation invariance implies it is diagonal in the site and frequency indices. Expanding the spin-dependence in Pauli matrices θ^μ , (which together with the identity form a basis of the complex vector space of complex 2×2 matrices) we obtain the symmetry constraints parametrization

$$G(x'; x) = \sum_{\mu=0}^3 G_i^\mu(\omega) \theta_{s's}^\mu \delta_{i'i} \delta_{\omega'\omega}, \quad (4.40)$$

where $\mu = 0, 1, 2, 3 \equiv d, x, y, z$ and $\theta^0 = \mathbb{1}_2$. Plugging this parametrization into the symmetry relations for the one-particle correlation function in Table 4.1, we obtain the constraints on $G^{\mu\nu}$ which are summarized in Table 4.3. In the derivation, we defined

$$\xi(\mu) = \begin{cases} +1 & \text{if } \mu = 0 \\ -1 & \text{otherwise} \end{cases} \quad (4.41)$$

and made use of the relation

$$s's \theta_{\bar{s}\bar{s}'} = \xi(\mu) \theta_{s's}^\mu, \quad (4.42)$$

Table 4.2 – Symmetry relations of the one- and two-particle correlation functions. The label denotes the symmetry or gauge transformation as defined in Table 4.1, where PH1 (PH2) refers to PH applied to the site i_1 (i_2). The multi indices are $x_j = (i_j, s_j, \omega_j)$ and we define $x_j^* = (i_j, s_j, -\omega_j)$, $\mathcal{T}x_j = (i_j, \bar{s}_j, -\omega_j)$ and $Lx_j = (L(i_j), s_j, \omega_j)$. s_i appearing as a factor should be interpreted as $s_i = \uparrow = 1$ and $s_i = \downarrow = -1$ and $\bar{s}_i = -s_i$ denotes a flipped spin.

| | |
|---|--------|
| $G(x'_1; x_1) = G(x'_1; x_1) \delta_{i'_1 i_1}$ | (U(1)) |
| $G(x'_1; x_1) \delta_{i'_1 i_1} = G(Lx'_1; Lx_1) \delta_{i'_1 i_1}$ | (L) |
| $G(x'_1; x_1) \delta_{i'_1 i_1} = G(x'_1; x_1) \delta_{i'_1 i_1} \delta_{\omega'_1, \omega_1}$ | (TT) |
| $G(x'_1; x_1) \delta_{i'_1 i_1} = -s'_1 s_1 G(\mathcal{T}x_1; \mathcal{T}x'_1) \delta_{i'_1 i_1}$ | (PH) |
| $G(x'_1; x_1) \delta_{i'_1 i_1} = s'_1 s_1 G(\mathcal{T}x'_1; \mathcal{T}x_1)^* \delta_{i'_1 i_1}$ | (TR) |
| $G(x'_1; x_1) \delta_{i'_1 i_1} = G(x_1^*; x_1'^*)^* \delta_{i'_1 i_1}$ | (H) |
| $G(x'_1, x'_2; x_1, x_2) = G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} - G(x'_2, x'_1; x_1, x_2) \delta_{i'_1 i_2} \delta_{i'_2 i_1}$ | (U(1)) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = G(Lx'_1, Lx'_2; Lx_1, Lx_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (L) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} \delta_{\omega'_1 + \omega'_2, \omega_1 + \omega_2}$ | (TT) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = -s'_1 s_1 G(\mathcal{T}x_1, x'_2; \mathcal{T}x'_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (PH1) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = -s'_2 s_2 G(x'_1, \mathcal{T}x_2; x_1, \mathcal{T}x'_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (PH2) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = s'_1 s'_2 s_1 s_2 G(x_1^*, x_2^*; x_1'^*, x_2'^*)^* \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (TR) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = G(x_1^*, x_2^*; x_1'^*, x_2'^*)^* \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (H) |
| $G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = G(x'_2, x'_1; x_2, x_1) \delta_{i'_1 i_1} \delta_{i'_2 i_2}$ | (X) |

where s appearing as a factor should be interpreted as $s = \pm 1$ and $\bar{s}_i = -s_i$ denotes a flipped spin. The first relation in Table 4.3 already implies that only the density component $G^0 \equiv G^d \equiv G$ is nonzero, rendering G diagonal also in spin space. For our pf-FRG implementation, we restrict to *Archimedean* lattices, where all sites are symmetry equivalent. In this case, the one-particle correlation function becomes completely site-independent, leading to the final parametrization

$$G(x', x) = G(\omega) \delta_{s' s} \delta_{i' i} \delta_{\omega' \omega}, \quad (4.43)$$

where $G(\omega) \in i\mathbb{R}$ is a completely imaginary function odd in frequency space $G(\omega) = G(-\omega)$.

For the two-particle correlator, applying the same symmetries and expanding in Pauli matrices yields

$$G(x'_1, x'_2; x_1, x_2) = \sum_{\mu, \nu=0}^3 G_{i_1 i_2}^{\mu\nu}(s, t, u) \theta_{s'_1 s_1}^{\mu} \theta_{s'_2 s_2}^{\nu} \delta_{i'_1 i_1} \delta_{i'_2 i_2} \delta_{\omega'_1 + \omega'_2, -\omega_1 - \omega_2}, \quad (4.44)$$

where we defined the bosonic transfer frequencies

$$s = \omega'_1 + \omega'_2, \quad t = \omega'_1 - \omega_1, \quad u = \omega'_1 - \omega_2, \quad (4.45)$$

a choice that we motivate when discussing the pf-FRG flow equations in the next section.

Inserting this parametrization into the symmetry relations of Table 4.1 again leads to constraints on $G_{i_1 i_2}^{\mu\nu}(s, t, u)$ which are also summarized in Table 4.3. Most notably, we find

$$G_{i_1 i_2}^{\mu\nu} \in \begin{cases} \mathbb{R} & \text{if } \xi(\mu)\xi(\nu) = 1 \\ i\mathbb{R} & \text{if } \xi(\mu)\xi(\nu) = -1 \end{cases}, \quad (4.46)$$

Table 4.3 – Symmetry relations on the parametrized one- and two-particle correlation functions. These relations are obtained by plugging the parametrization in Eqs. (4.40,4.44) in the symmetry relations for the full correlation functions defined in Table 4.2. The symmetries/gauge transformations used in the derivation are indicated by the labels on the right, which are defined in Table 4.1, where PH1 (PH2) refers to PH applied to the site i_1 (i_2).

| | |
|---|---|
| $G_i^\mu(\omega) = \xi(\mu)G_i^\mu(\omega)$ | (H \circ TR) |
| $G_i^\mu(\omega) = -\xi(\mu)G_i^\mu(-\omega)$ | (PH) |
| $G_i^\mu(\omega) = -G_i^\mu(\omega)^*$ | (TR \circ PH) |
| $G_{i_1 i_2}^{\mu\nu}(s, t, u) = G_{i_2 i_1}^{\nu\mu}(-s, t, u)$ | (X \circ H \circ TR \circ PH1 \circ PH2) |
| $G_{i_1 i_2}^{\mu\nu}(s, t, u) = \xi(\mu)\xi(\nu)G_{i_1 i_2}^{\mu\nu}(s, -t, u)$ | (H \circ TR) |
| $G_{i_1 i_2}^{\mu\nu}(s, t, u) = \xi(\mu)\xi(\nu)G_{i_2 i_1}^{\nu\mu}(s, t, -u)$ | (X \circ H \circ TR) |
| $G_{i_1 i_2}^{\mu\nu}(s, t, u) = -\xi(\nu)G_{i_1 i_2}^{\mu\nu}(u, t, s)$ | (PH2) |
| $G_{i_1 i_2}^{\mu\nu}(s, t, u) = \xi(\mu)\xi(\nu)G_{i_1 i_2}^{\mu\nu}(s, t, u)^*$ | (TR \circ H \circ TR \circ PH1 \circ PH2) |

so that diagonal components $G_{i_1 i_2}^{\mu\mu}$ are real and off-diagonal ones purely imaginary. Moreover, negative frequencies can be mapped to positive ones, and s and u may be interchanged.

Finally, we note that the same symmetry constraints carry over from the disconnected correlation functions to the self-energy Σ and the vertex Γ . This follows directly from their relations in Eqs. (4.16, 4.17), together with the fact that the one-particle correlation function is fully diagonal. We therefore parametrize the self-energy as

$$\Sigma(x', x) = \Sigma(\omega)\delta_{s's}\delta_{i'i}\delta_{\omega'\omega}, \quad (4.47)$$

and the vertex as

$$\Gamma(x'_1, x'_2; x_1, x_2) = \sum_{\mu, \nu=0}^3 \Gamma_{i_1 i_2}^{\mu\nu}(s, t, u)\theta_{s'_1 s_1}^\mu\theta_{s'_2 s_2}^\nu\delta_{i'_1 i_1}\delta_{i'_2 i_2}\delta_{\omega'_1 + \omega'_2, -\omega_1 - \omega_2}, \quad (4.48)$$

which both satisfy the same symmetry relation as the disconnected correlation functions stated in Table 4.3.

4.2.4 pf-FRG flow equations

We now derive the flow equations for the parametrized self-energy and vertex within the Katanin truncation. This requires inserting the parametrizations into the general flow equations defined in Eqs. (4.22, 4.23), and collecting terms with the same structure in the site and frequency indices (as dictated by the Kronecker deltas).

For the self-energy, we can additionally evaluate the sum over the spin-indices analytically. Using the vertex symmetries and properties of the Pauli matrices (a detailed derivation can be found in the appendix of Ref. [118]), we obtain

$$\frac{d}{d\Lambda}\Sigma^\Lambda(\omega) = -\frac{1}{2\pi} \int d\omega' \left[2 \sum_j \Gamma_{ij}^{\Lambda dd}(\omega' + \omega, \omega' - \omega, 0) - \sum_\mu \Gamma_{ii}^{\Lambda\mu\mu}(\omega' + \omega, 0, \omega - \omega') \right] S^\Lambda(\omega'), \quad (4.49)$$

where i is an arbitrary site, which shows that only diagonal vertex components $\Gamma^{\mu\mu}$ contribute to the self-energy flow.

For the vertex, resolving the δ -distributions in frequency space results in a flow equation that can be split into three channels

$$\frac{d}{d\Lambda} \Gamma^\Lambda = \dot{g}_s^\Lambda + \dot{g}_t^\Lambda + \dot{g}_u^\Lambda, \quad (4.50)$$

where the channel \dot{g}_c^Λ depends explicitly on the transfer frequency $c = s, t, u$ defined in Eq. (4.45). Explicitly, the s -channel is defined as

$$\begin{aligned} \dot{g}_{s i_1 i_2}^\Lambda(x'_1, x'_2; x_1, x_2) = & -\frac{1}{2\pi} \sum_{\mu\nu\kappa\eta} \sum_{s_3, s_4} \int d\omega \left[P(s, s - \omega) \right. \\ & \left. \Gamma_{i_1 i_2}^{\Lambda\mu\nu}(s, \omega - \omega_1, \omega - \omega_2) \Gamma_{i_1 i_2}^{\Lambda\kappa\eta}(s, \omega'_1 - \omega, \omega + \omega'_1 - s) \theta_{s_3 s_1}^\mu \theta_{s_4 s_2}^\nu \theta_{s'_1 s_3}^\kappa \theta_{s'_2 s_4}^\eta \right], \end{aligned} \quad (4.51)$$

the t -channel as

$$\begin{aligned} \dot{g}_{t i_1 i_2}^\Lambda(x'_1, x'_2; x_1, x_2) = & -\frac{1}{2\pi} \sum_{\mu\nu\kappa\eta} \sum_{s_3, s_4} \int d\omega \left[P(\omega, \omega - t) \right. \\ & \sum_j \Gamma_{i_1 j}^{\Lambda\mu\nu}(\omega + \omega'_1 - t, t, \omega'_1 - \omega) \Gamma_{j i_2}^{\Lambda\kappa\eta}(\omega'_2 + \omega, t, \omega - \omega_2) \theta_{s'_1 s_1}^\mu \theta_{s_4 s_3}^\nu \theta_{s_3 s_4}^\kappa \theta_{s'_2 s_2}^\eta \\ & \Gamma_{i_1 i_2}^{\Lambda\mu\nu}(\omega + \omega'_1 - t, t, \omega'_1 - \omega) \Gamma_{i_2 i_2}^{\Lambda\kappa\eta}(\omega'_2 + \omega, \omega - \omega_2, t) \theta_{s'_1 s_1}^\mu \theta_{s_4 s_3}^\nu \theta_{s_3 s_2}^\kappa \theta_{s'_2 s_4}^\eta \\ & \left. \Gamma_{i_1 i_1}^{\Lambda\mu\nu}(\omega + \omega'_1 - t, t - \omega + \omega_1, t) \Gamma_{i_1 i_2}^{\Lambda\kappa\eta}(\omega + \omega'_2, t, \omega - \omega_2) \theta_{s'_1 s_3}^\mu \theta_{s_4 s_1}^\nu \theta_{s_3 s_4}^\kappa \theta_{s'_2 s_2}^\eta \right], \end{aligned} \quad (4.52)$$

and the u -channel as

$$\begin{aligned} \dot{g}_{u i_1 i_2}^\Lambda(x'_1, x'_2; x_1, x_2) = & -\frac{1}{2\pi} \sum_{\mu\nu\kappa\eta} \sum_{s_3, s_4} \int d\omega \left[P(u, u + \omega) \right. \\ & \Gamma_{i_1 i_2}^{\Lambda\mu\nu}(\omega + \omega'_2 + u, u + \omega - \omega_1, u) \Gamma_{i_1 i_2}^{\Lambda\kappa\eta}(\omega'_1 + \omega, \omega'_1 - \omega - u, u) \theta_{s_3 s_1}^\mu \theta_{s'_2 s_4}^\nu \theta_{s'_1 s_3}^\kappa \theta_{s_4 s_2}^\eta \left. \right], \end{aligned} \quad (4.53)$$

where the $x_j = (i_j, \omega_j)$ now only contain the site and frequency index. Here, terms with exchanged indices ($x_3 \leftrightarrow x_4$) were absorbed into the product $G(x_3, x'_3)S(x_4, x'_4)$, giving rise to the *propagator bubble*

$$P^\Lambda(\omega, \omega') = G^\Lambda(\omega)S_{\text{kat}}^\Lambda(\omega') + S_{\text{kat}}^\Lambda(\omega)G^\Lambda(\omega') = \partial_\Lambda [G^\Lambda(\omega)G^\Lambda(\omega')], \quad (4.54)$$

and the single-scale propagator has been replaced by its Katanin form defined in Eq. (4.25). Flow equations for a specific component $\Gamma_{i_1 i_2}^{\Lambda\mu\nu}(s, t, u)$ are obtained by summing over the spin indices on the right-hand side and keeping only the terms proportional to $\theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\nu$, which we do numerically (a full diagrammatic representation of the resulting flow equations for all $\Gamma^{\mu\nu}$ components is provided in the appendix of Ref. [118]). The flow equations include integrals over continuous Matsubara frequencies and, in the t -channel, a sum over the infinite lattice. We discuss the numerical solution of these equations in Sec. 4.4.

Here, we can already observe a key advantage of the pf-FRG: the vertex depends only on two sites, and—thanks to the translational invariance of Bravais lattices—we can further restrict one of these indices to a reference unit cell (as we will show in Sec. 4.4 below). The flow equations additionally contain a single site summation. Consequently, the computational cost scales

roughly as N^2 with the number of sites N . This polynomial scaling is vastly more favorable than the exponential growth of the Hilbert space in exact treatments, and makes pf-FRG numerically feasible even for three-dimensional spin models.

Finally, the initial conditions follow from inserting the vertex parametrization into general FRG initial conditions (4.19) and antisymmetrizing the pseudo-fermion Hamiltonian, yielding

$$\Sigma^{\Lambda \rightarrow \infty}(\omega) = 0, \quad (4.55)$$

$$\Gamma_{i_1 i_2}^{\Lambda \rightarrow \infty \mu \nu}(s, t, u) = \begin{cases} \frac{1}{4} J_{i_1 i_2}^{\mu \nu} & \text{for } \mu, \nu > 0 \\ 0 & \text{otherwise} \end{cases} \quad (4.56)$$

which are simply the bare couplings of the spin Hamiltonian (4.26).

4.2.5 Regulator

To fully specify the flow equations, we still need to define how the infrared cutoff is explicitly implemented in the bare propagator G_0 . Since the pseudo-fermion Hamiltonian has no kinetic term, the definition of G_0 in Eq. (4.7) implies that it is diagonal in all indices

$$G_0(x'; x) = G_0(\omega) \delta_{i, i'} \delta_{\omega, \omega'} \delta_{s, s'}, \quad (4.57)$$

with frequency dependence

$$G_0(\omega) = \frac{1}{i\omega}. \quad (4.58)$$

It is therefore natural to implement the infrared cutoff in frequency space by a multiplication with a regulator $R(\Lambda, \omega)$ as

$$G_0^\Lambda(\omega) = R(\Lambda, \omega) G_0(\omega) \quad (4.59)$$

so that G_0^Λ satisfies the boundary conditions of Eq. (4.18) required for the derivation of the FRG flow equations. Using the Dyson equation (4.16) and the fact that the self-energy is initially zero in the FRG flow, the full propagator takes the form

$$G^\Lambda(\omega) = \frac{R(\omega, \Lambda)}{i\omega - \Sigma^\Lambda(\omega)}. \quad (4.60)$$

In earlier implementations, the regulator was chosen as a sharp Heaviside step function $R(\Lambda, \omega) = \Theta(|\omega| - \Lambda)$. However, this produces a Dirac-delta distribution $\delta(|\omega| - \Lambda)$ in the single-scale propagator $S^\Lambda = \partial_\Lambda^\Sigma G^\Lambda$. While in the $L2$ truncation this δ -distribution simplifies the frequency integrals on the right-hand side of the flow equations to discrete sums, in the Katanin truncation the additional term arising from the substitution $S^\Lambda \rightarrow S_{\text{kat}}^\Lambda$ [Eq. (4.25)] cancels this simplification. Moreover, the resulting singularities in S^Λ introduce numerical artifacts, leading to unphysical oscillatory behavior in the pf-FRG flow, as observed in several early works [13, 134].

To overcome these issues, we employ the smooth regulator [14, 129]

$$R(\omega, \Lambda) = 1 - e^{-\frac{\omega^2}{\Lambda^2}} \quad (4.61)$$

for which the single-scale propagator evaluates to

$$iS^\Lambda(\omega) = \frac{2e^{-\frac{\omega^2}{\Lambda^2}}}{[\omega - i\Sigma^\Lambda(\omega)]^2} \frac{\omega^3}{\Lambda^3}. \quad (4.62)$$

This expression is smooth and typically features two symmetric peaks around $\omega = 0$ at frequencies $|\omega_p| < \Lambda$ (visible already for $\Sigma = 0$). Compared to the sharp cutoff, the smooth regulator significantly improves numerical stability of the pf-FRG flow.

Table 4.4 – Symmetry relations of the channel-resolved two-particle vertex in the asymptotic frequency parametrization. Shown are transformation properties of $g_{c,i_1i_2}^{\mu\nu}$ in the s -, t -, and u -channels that can be derived by inserting the channel decomposition Eq. (4.63) into the symmetry relations for the full vertex in Table 4.3.

| |
|---|
| $g_{s,i_1i_2}^{\mu\nu}(\omega_s, \nu_s, \nu'_s) = g_{s,i_2i_1}^{\nu\mu}(-\omega_s, \nu_s, \nu'_s)$ |
| $g_{s,i_1i_2}^{\mu\nu}(\omega_s, \nu_s, \nu'_s) = -\xi(\mu) g_{u,i_2i_1}^{\nu\mu}(\omega_s, -\nu_s, \nu'_s)$ |
| $g_{s,i_1i_2}^{\mu\nu}(\omega_s, \nu_s, \nu'_s) = -\xi(\nu) g_{u,i_1i_2}^{\mu\nu}(\omega_s, \nu_s, -\nu'_s)$ |
| $g_{s,i_1i_2}^{\mu\nu}(\omega_s, \nu_s, \nu'_s) = g_{s,i_2i_1}^{\nu\mu}(\omega_s, \nu'_s, \nu_s)$ |
| |
| $g_{t,i_1i_2}^{\mu\nu}(\omega_t, \nu_t, \nu'_t) = \xi(\mu)\xi(\nu) g_{t,i_1i_2}^{\mu\nu}(-\omega_t, \nu_t, \nu'_t)$ |
| $g_{t,i_1i_2}^{\mu\nu}(\omega_t, \nu_t, \nu'_t) = \xi(\mu) g_{t,i_1i_2}^{\mu\nu}(\omega_t, -\nu_t, \nu'_t)$ |
| $g_{t,i_1i_2}^{\mu\nu}(\omega_t, \nu_t, \nu'_t) = \xi(\nu) g_{t,i_1i_2}^{\mu\nu}(\omega_t, \nu_t, -\nu'_t)$ |
| $g_{t,i_1i_2}^{\mu\nu}(\omega_t, \nu_t, \nu'_t) = g_{t,i_2i_1}^{\nu\mu}(\omega_t, \nu'_t, \nu_t)$ |
| |
| $g_{u,i_1i_2}^{\mu\nu}(\omega_u, \nu_u, \nu'_u) = \xi(\mu)\xi(\nu) g_{u,i_2i_1}^{\nu\mu}(-\omega_u, \nu_u, \nu'_u)$ |
| $g_{u,i_1i_2}^{\mu\nu}(\omega_u, \nu_u, \nu'_u) = \xi(\nu) g_{u,i_2i_1}^{\nu\mu}(\omega_u, -\nu_u, \nu'_u)$ |
| $g_{u,i_1i_2}^{\mu\nu}(\omega_u, \nu_u, \nu'_u) = -\xi(\nu) g_{u,i_1i_2}^{\mu\nu}(\omega_u, \nu_u, -\nu'_u)$ |
| $g_{u,i_1i_2}^{\mu\nu}(\omega_u, \nu_u, \nu'_u) = g_{u,i_2i_1}^{\nu\mu}(\omega_u, \nu'_u, \nu_u)$ |

4.2.6 Asymptotic frequency parametrization

The bosonic transfer-frequency parametrization (s, t, u) introduced above is convenient for implementing symmetries and was used in most early pf-FRG implementations. However, as pointed out in Ref. [135], it leads to a nontrivial asymptotic structure of the vertex: when one transfer frequency becomes large while the others remain finite, the vertex does not decay to zero but approaches a constant depending on the remaining frequencies. A numerical implementation needs to accurately capture this asymptotic behavior to avoid accumulation of large errors. To overcome this issue, Ref. [135] proposed an alternative *asymptotic frequency parametrization*, which we also adopt and briefly review here.

The general idea is to decompose the full vertex into the three channels (already introduced for the flow equations in Eq. (4.50)) as

$$\Gamma^\Lambda(s, t, u) = \Gamma^{\Lambda \rightarrow \infty} + \sum_{c=(s,t,u)} g_c(\omega_c, \nu_c, \nu'_c), \quad (4.63)$$

where we only specify the frequency dependence for brevity, and $\Gamma^{\Lambda \rightarrow \infty}$ are the frequency independent initial conditions. Each channel is described by one bosonic frequency ω_c (identified with the corresponding transfer frequency $c = s, t, u$) and two associated fermionic frequencies (ν_c, ν'_c) , whose definitions differ between channels. A convenient choice leading to simple symmetry constraints [129] is

$$\begin{aligned} \omega_s &= s & \nu_s &= \omega_1 - \frac{s}{2} & \nu'_s &= \frac{s}{2} - \omega'_1 \\ \omega_t &= t & \nu_t &= \omega_1 + \frac{t}{2} & \nu'_t &= \omega'_1 - \frac{t}{2} \\ \omega_u &= u & \nu_u &= \omega_1 - \frac{u}{2} & \nu'_u &= \omega'_1 - \frac{u}{2}. \end{aligned} \quad (4.64)$$

It can then be shown that the channel decomposes further into different *kernels* [135] as

$$g_c(\omega_c, \nu_c, \nu'_c) = K_1^c(\omega_c) + K_2^c(\omega_c, \nu_c) + \bar{K}_2^c(\omega_c, \nu'_c) + K_3^c(\omega_c, \nu_c, \nu'_c). \quad (4.65)$$

and each kernel vanishes once *any* of its frequency arguments tends to infinity. Specifically, K_1^c captures the asymptotics for $\nu_c, \nu'_c \rightarrow \infty$, K_2^c and \bar{K}_2^c describe the limits $\nu'_c \rightarrow \infty$ and $\nu_c \rightarrow \infty$,

respectively, while K_3^c retains the full frequency dependence but decays rapidly as soon as one frequency becomes large.

There are two strategies of implementing this numerically: First, we could define separate frequency grids for each kernel, enabling us to use a large grid with many frequencies for K_1^c , an intermediate grid for K_2^c and \bar{K}_2^c , and only a relatively small but dense grid for K_3^c . However, when evaluating the channel inside the region of the K_3^c kernel, all other kernels need to be evaluated as well. As we discuss in Sec. 4.4 below, each evaluation requires a relatively expensive interpolation when frequencies fall outside the discrete numerical frequency grid. As an alternative, we adopt a strategy that is build on defining new kernels [129]

$$\begin{aligned} Q_{1c}^\Lambda(\omega_c) &\equiv \lim_{|\nu_c|, |\nu'_c| \rightarrow \infty} \dot{g}_c^\Lambda(\omega_c, \nu_c, \nu'_c), \\ Q_{2c}^\Lambda(\omega_c, \nu_c) &\equiv \lim_{|\nu'_c| \rightarrow \infty} \dot{g}_c^\Lambda(\omega_c, \nu_c, \nu'_c), \\ \bar{Q}_{2c}^\Lambda(\omega_c, \nu'_c) &\equiv \lim_{|\nu_c| \rightarrow \infty} \dot{g}_c^\Lambda(\omega_c, \nu_c, \nu'_c), \\ Q_{3c}^\Lambda(\omega_c, \nu_c, \nu'_c) &\equiv \dot{g}_c^\Lambda(\omega_c, \nu_c, \nu'_c), \end{aligned} \quad (4.66)$$

which all use the same frequency grid and are simply the sums over the original kernels

$$\begin{aligned} Q_{1c}^\Lambda(\omega_c) &= K_{1c}^\Lambda(\omega_c), \\ Q_{2c}^\Lambda(\omega_c, \nu_c) &= K_{1c}^\Lambda(\omega_c) + K_{2c}^\Lambda(\omega_c, \nu_c), \\ \bar{Q}_{2c}^\Lambda(\omega_c, \nu'_c) &= K_{1c}^\Lambda(\omega_c) + \bar{K}_{2c}^\Lambda(\omega_c, \nu'_c), \\ Q_{3c}^\Lambda(\omega_c, \nu_c, \nu'_c) &= K_{1c}^\Lambda(\omega_c) + K_{2c}^\Lambda(\omega_c, \nu_c) + \bar{K}_{2c}^\Lambda(\omega_c, \nu'_c) + K_{3c}^\Lambda(\omega_c, \nu_c, \nu'_c). \end{aligned} \quad (4.67)$$

Here, an evaluation of the channel \dot{g}_c only requires just one interpolation, which in our cases was more numerically efficient compared to a parametrization in the original kernels, even when a slightly larger frequency grid is required to correctly resolve Q_{3c}^Λ . In this approach, the lower-order kernels are then invoked when one (or more) of the three frequencies lies out of the numerical frequency grid, and thus act as efficient buffers of the asymptotic behavior of the channel.

The spin and site structure of the channels remains identical to that of the full vertex, and we can therefore employ the same parametrization (4.48) we used for the vertex for each channel as

$$g_c(x'_1, x'_2; x_1, x_2) = \sum_{\mu\nu=0}^3 g_{c,i_1 i_2}^{\mu\nu}(s, t, u) \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\nu \delta_{i'_1 i_1} \delta_{i'_2 i_2} \delta_{\omega'_1 + \omega'_2, -\omega_1 - \omega_2}. \quad (4.68)$$

The frequency structure of the symmetries, however, translate into new symmetries in terms of ω_c, ν_c, ν'_c , which are summarized in Table 4.4. Most notably, the exchange in $s \leftrightarrow u$ maps the respective channels onto each. Additionally, all positive frequencies can be mapped to their negative counter part and thus only components with positive frequencies need to be explicitly computed. Furthermore, we can exchange the fermionic frequencies $\nu_c \leftrightarrow \nu'_c$, allowing us to only compute the vertex for $\nu_c \leq \nu'_c$.

Finally, the frequency arguments in the flow equations must be reformulated in this parametrization. Their explicit form can be found in the appendix of Ref. [14].

4.2.7 Symmetries in spin space

When defining the spin Hilbert space, we have to make a basis choice which spin direction corresponds to the quantization axis (most often we call this axis the z -axis). This choice is, of

course, arbitrary as the commutation relations are invariant under global $SU(2)$ transformations that rotate this axis in spin space. Many spin models we consider, will also be invariant under such transformations. For instance, all pure Heisenberg models are invariant under *any* $SU(2)$ rotation in spin space. We can utilize these additional symmetries, which again strongly constrain the correlation functions. For the Heisenberg case, for instance, only the density $\Gamma^d = \Gamma^{dd}$ and the spin vertex $\Gamma^s = \Gamma^{xx} = \Gamma^{yy} = \Gamma^{zz}$ are nonzero. Other models may be invariant only under a subgroup of $SU(2)$ (e.g. $U(1)$), which again reduces the number of spin components of the vertex that need to be considered. Since we have not found a detailed derivation of these constraints in the literature (apart from my master thesis in the context of spin-valley models), we include one here. For brevity, we suppress all indices except the spin labels.

First, we need to understand how an “ $SU(2)$ rotation” acts on the spin and fermionic operators. For simplicity of the derivation, we first only consider a rotation around a single axis by an angle φ , for which, without loss of generality, we choose the z -axis. For the Hilbert space of spin-1/2 particles in the standard basis, the corresponding $SU(2)$ matrix becomes

$$\mathbf{U} = e^{i\varphi\theta^z/2} = \begin{pmatrix} e^{i\varphi/2} & 0 \\ 0 & e^{-i\varphi/2} \end{pmatrix} \in SU(2), \quad (4.69)$$

which acts on states as $|\psi\rangle \rightarrow U|\psi\rangle$. This translates into a transformation of the fermionic operators as

$$f_s^\dagger \rightarrow f_{s'}^\dagger U_{s's}^\dagger \quad f_s \rightarrow U_{ss'} f_{s'}. \quad (4.70)$$

Note that these *global* transformations differ from the *local* $SU(2)$ gauge redundancy discussed in Sec. 4.2.1, which acted by a right multiplications of \mathbf{U} on f , instead of left multiplications as is the case here. Using the pseudo-fermion mapping (4.27), the spin operators transform as

$$S^a = f_{is_1}^\dagger \theta_{s_1 s_2}^a f_{is_2} \rightarrow f_{s'_1}^\dagger U_{s'_1 s_1}^\dagger \theta_{s_1 s_2}^a U_{s_2 s'_2} f_{s'_2} = \sum_b [R]^{ab} S^b \implies \mathbf{S} \rightarrow \mathbf{R} \mathbf{S}, \quad (4.71)$$

In the second step, we used the identity $\mathbf{U}^\dagger \boldsymbol{\theta} \mathbf{U} = \mathbf{R} \boldsymbol{\theta}^2$, where $\boldsymbol{\theta} = (\theta^x, \theta^y, \theta^z)^T$ is the Pauli matrix vector and \mathbf{R} is the $SO(3)$ rotation matrix corresponding to the $SU(2)$ rotation. For a rotation around the z -axis by the angle φ , this reads

$$\mathbf{R}^\varphi = \begin{pmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (4.72)$$

Spin Hamiltonians consequently transform as

$$\mathbf{S}^T \mathbf{J} \mathbf{S} \rightarrow \mathbf{S}^T \mathbf{R}^T \mathbf{J} \mathbf{R} \mathbf{S}. \quad (4.73)$$

A spin models is thus invariant under the specific $SU(2)$ rotation if

$$\mathbf{J} = \mathbf{R}^T \mathbf{J} \mathbf{R}. \quad (4.74)$$

For invariance under arbitrary $SU(2)$ rotations, this condition is satisfied only for the Heisenberg case $J^{ab} \sim \delta^{ab}$. Less restrictive symmetries, on the other hand, allow for other couplings.

²This can be derived by evaluating the matrix exponential in the definition of \mathbf{U} by the Baker–Campbell–Hausdorff formula, and using the commutation relations of the $SU(2)$ Lie algebra

Next, we need to determine how the correlation functions transform. For the one-particle function we find

$$G(s'; s) = -\langle f_s f_{s'}^\dagger \rangle \rightarrow -U_{sr} \langle f_r f_{r'}^\dagger \rangle U_{r's'}^\dagger = U_{r's'}^\dagger G(r', r) U_{sr}. \quad (4.75)$$

An expansion in terms of Pauli matrices yields

$$\sum_{\mu=0}^3 G^\mu \theta_{s's}^\mu \rightarrow \sum_{\mu=0}^3 G^\mu U_{s'r'}^\dagger \theta_{r'r}^\mu U_{rs} = G^0 + \sum_{a,b=1}^3 G^a R^{ab} \theta_{s's}^b, \quad (4.76)$$

which is invariant if

$$G^a = \sum_{b=1}^3 G^b R^{ba}. \quad (4.77)$$

For SU(2) spins in pf-FRG these components already vanish by other symmetries, but this relation becomes important for generalizations to SU(4) models discussed in the next section.

For the two-particle correlation functions, a similar derivation gives

$$G(s'_1, s'_2; s_1 s_2) \rightarrow U_{r'_1 s'_1}^\dagger U_{r'_2 s'_2}^\dagger G(r'_1, r'_2; r_1 r_2) U_{s_1 r_1} U_{s_2 r_2} \quad (4.78)$$

and in the Pauli basis

$$\sum_{\mu\nu} G^{\mu\nu} \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\nu \rightarrow \sum_{\mu\nu\kappa\eta} R^{\mu\kappa} R^{\nu\eta} G^{\kappa\eta} \theta_{s'_1 s_1}^\kappa \theta_{s'_2 s_2}^\eta, \quad (4.79)$$

where here \mathbf{R} needs to be understood as a 4×4 matrix that acts trivially on the $\mu = 0$ components (i.e. $R^{0\mu} = R^{\mu 0} = \delta_{\mu 0}$). Invariance therefore requires

$$G^{\mu\nu} = \sum_{\kappa\eta} [R^{\mu\kappa}]^T G^{\kappa\eta} R^{\eta\nu}. \quad (4.80)$$

These constraints again carry over directly to the self-energy and the vertex.

We now explicitly examine what this implies for spin models that are symmetric under varying subgroups of SU(2), namely U(1), \mathbb{Z}_2 , $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$ (completely diagonal models), and SU(2) (the Heisenberg model).

U(1) symmetry We start with U(1)-symmetric models, which are invariant under arbitrary rotations about a *single* axis. We can therefore continue with the example of a rotation about the z -axis. According to Eq. (4.74) with the rotation matrix \mathbf{R}^φ defined in Eq. (4.72), a spin model is U(1)-invariant if its coupling matrix takes the form

$$\mathbf{J} = \begin{pmatrix} J^{xx} & J^{xy} & 0 \\ -J^{xy} & J^{xx} & 0 \\ 0 & 0 & J^{zz} \end{pmatrix}, \quad (4.81)$$

which can be derived by considering an infinitesimal $\varphi = \epsilon$ and expanding \mathbf{R}^ϵ to first order. This includes, XXZ and Dzyaloshinskii–Moriya (DM) interactions.

Enforcing invariance of the one-particle correlation functions via Eq. (4.77) gives

$$G^x = G^y = 0. \quad (4.82)$$

For the two-particle correlators, Eq. (4.80) yields

$$G^{\mu\nu} = \begin{pmatrix} G^{dd} & 0 & 0 & G^{zd} \\ 0 & G^{xx} & G^{xy} & 0 \\ 0 & -G^{xy} & G^{xx} & 0 \\ G^{dz} & 0 & 0 & G^{zz} \end{pmatrix}^{\mu\nu}. \quad (4.83)$$

The terms involving density indices $d = 0$ are symmetry-allowed, even though they are absent in the initial conditions of the flow equations. During the FRG flow, however, they indeed acquire finite values. Exploiting the U(1) symmetry thus reduces the sixteen possible vertex components $\Gamma^{\mu\nu}$ to only six independent ones.

\mathbb{Z}_2 symmetry Another common discrete symmetry is invariance under a rotation by $\varphi = \pi$ about a single axis. Again choosing the z -axis, the corresponding rotation matrix is

$$\mathbf{R}^{\pi/2} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (4.84)$$

which flips $S^x \rightarrow -S^x$ and $S^y \rightarrow -S^y$. Hamiltonians invariant under this rotation have coupling matrices of the form

$$\mathbf{J} = \begin{pmatrix} J^{xx} & J^{xy} & 0 \\ J^{yx} & J^{yy} & 0 \\ 0 & 0 & J^{zz} \end{pmatrix}, \quad (4.85)$$

which is slightly less restrictive than U(1) invariance and, for instance, includes models with symmetric off-diagonal exchange.

We emphasize that a transformation flipping only a *single* spin component (e.g. $S^z \rightarrow -S^z$) is *not* a valid spin symmetry, since it cannot be represented by an SU(2) matrix and therefore does not preserve the SU(2) commutation relations. Although the above Hamiltonian is invariant under $S^z \rightarrow -S^z$, this does not enforce additional constraints on correlation functions. The correct \mathbb{Z}_2 symmetry instead implies

$$G^x = G^y = 0, \quad (4.86)$$

and

$$G^{\mu\nu} = \begin{pmatrix} G^{dd} & 0 & 0 & G^{zd} \\ 0 & G^{xx} & G^{xy} & 0 \\ 0 & G^{yx} & G^{yy} & 0 \\ G^{dz} & 0 & 0 & G^{zz} \end{pmatrix}^{\mu\nu}, \quad (4.87)$$

leaving eight independent components. Notably, the mixed components G^{dz} and G^{zd} are symmetry-allowed and indeed emerge during the FRG flow, even though a naive application of the $S^z \rightarrow -S^z$ transformation might suggest they should vanish.

$\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry (diagonal spin models) In practice, \mathbb{Z}_2 symmetry most often appears simultaneously in all spin directions. All relations derived for the z direction above of course equally hold also for the x and y directions. Combining the constraints for x , y , and z , the coupling matrix is forced to be take the diagonal form

$$\mathbf{J} = \begin{pmatrix} J^{xx} & 0 & 0 \\ 0 & J^{yy} & 0 \\ 0 & 0 & J^{zz} \end{pmatrix}, \quad (4.88)$$

which includes the famous Kitaev model. The corresponding constraints on the correlation functions are

$$G^x = G^y = G^z = 0, \quad (4.89)$$

leaving only finite G^d , and

$$G^{\mu\nu} = \begin{pmatrix} G^{dd} & 0 & 0 & 0 \\ 0 & G^{xx} & 0 & 0 \\ 0 & 0 & G^{yy} & 0 \\ 0 & 0 & 0 & G^{zz} \end{pmatrix}^{\mu\nu}, \quad (4.90)$$

so that only diagonal vertex components remain finite.

Heisenberg models Finally, the full SU(2) symmetry requires invariance under arbitrary rotations, which enforces

$$\mathbf{J} = \begin{pmatrix} J & 0 & 0 \\ 0 & J & 0 \\ 0 & 0 & J \end{pmatrix}, \quad (4.91)$$

and the constraints

$$G^x = G^y = G^z = 0, \quad (4.92)$$

and

$$G^{\mu\nu} = \begin{pmatrix} G^d & 0 & 0 & 0 \\ 0 & G^s & 0 & 0 \\ 0 & 0 & G^s & 0 \\ 0 & 0 & 0 & G^s \end{pmatrix}^{\mu\nu}, \quad (4.93)$$

so that only two independent vertex components already mentioned in the introduction remain.

4.2.8 Single occupation constraint

Expressing spin operators in terms of fermions as in Eq. (4.27) is only a faithful mapping if the single-occupation constraint (4.29) is exactly enforced. Let us briefly comment on how this can be done in practice.

In principle, one can enforce the constraint exactly at the path-integral level by introducing a Lagrange multiplier in the form of an SU(2) gauge field in the action [14]. This approach, however, requires keeping track of an additional non-Abelian field, which significantly complicates the flow equations and has not been systematically pursued so far [136].

At finite temperatures, Popov and Fedotov showed that the unphysical states can be projected out by adding a finite imaginary chemical potential $\mu = i\pi T/2$ [137]. This additional term, however, breaks the invariance of the pseudo-fermion Hamiltonian under Hermitian conjugation, time-reversal and particle-hole conjugation (where the term changes sign). For this reason, the method was originally deemed numerically impractical in the context of the pf-FRG [118]. More recently, it was demonstrated [138] that combinations of these transformations remain exact symmetries in the presence of the Popov–Fedotov potential. Compared to conventional finite temperature pf-FRG, this increases the numerical cost by only roughly a factor of four. This is balanced by the fact that at finite temperature Matsubara frequencies are discrete, which at high T can even render the calculation numerically cheaper than convectional $T = 0$ pf-FRG. At sufficiently high temperatures $T \gtrsim 0.3J$, the authors of Ref. [138] additionally demonstrated that quantitatively exact results (compared to ED calculations) could be obtained for small spin

clusters. At lower temperatures, however, the constraint is no longer enforced exactly—likely a consequence of the truncation of the flow equations becoming less accurate once J/T is no longer small—and the calculations indeed deviate from ED. At $T = 0$, the Popov–Fedotov potential vanishes identically, rendering this method unsuitable for zero-temperature calculations.

Since we are primarily interested in quantum effects that occur at low temperatures, however, we work at $T = 0$. Here, the single-occupation constraint cannot easily be enforced exactly, but is enforced on average simply by setting the chemical potential to zero. At zero chemical potential, the Hamiltonian is invariant under particle–hole conjugation (see Sec. 4.2.1), and this invariance is preserved throughout the FRG flow, guaranteeing $\langle n_i \rangle = 1$ at all scales [14]. Nevertheless, local particle-number fluctuations around this mean value remain, which in principle results in contributions from the unphysical states. It has been argued, however, that for most spin models the unphysical states are higher in energy than the physical ones, and are therefore not populated at $T = 0$ [13]. This argument can be reinforced by introducing a level-repulsion term of the form

$$H \rightarrow H - A \sum_i \mathbf{S}_i^2. \quad (4.94)$$

For the physical $S = 1/2$ sector this is simply $S_i^2 = S(S + 1)$. The unphysical states, however, both have $S = 0$ and are therefore energetically penalized. If the presence of unphysical states played a dominant role in the FRG flow, adding such a term should substantially alter the results. In practice, however, studies [67, 104] have shown that the flow is only rescaled and otherwise behaves identically for large A and $A = 0$, providing evidence that particle-number fluctuations do not qualitatively affect pf-FRG results.

Away from the thermodynamic limit, however, a more recent study demonstrated that for small spin clusters this no longer holds [139]. In this case, the exact ground state of the pseudo-fermion Hamiltonian without the constraint has weight in the unphysical sector, and particle-number fluctuations spoil the validity of the pf-FRG. A systematic analysis of the precise impact of constraint violations is still lacking. In this thesis, we employ the conventional $T = 0$ pf-FRG with $A = 0$ and the chemical potential set to zero, while keeping in mind that our results may be affected by such violations.

Finally, let us note an elegant way to avoid the constraint altogether: a representation of spins in terms of (real) Majorana fermions instead of complex fermions provides a mapping between the spin and fermionic Hilbert spaces without the introduction of unphysical states. The resulting framework, called pseudo-Majorana FRG (pm-FRG) [140, 141], will be discussed further in Sec. 4.6.

4.3 Generalization to SU(4) spin models

In this section, we discuss how the pseudo-fermion FRG, developed above for SU(2) systems, can be generalized to a broad class of SU(4) spin models. Such models arise naturally in the context of Moiré materials (where they are often referred to as spin-valley models) and in spin–orbit entangled Mott insulators (commonly termed spin–orbital models) as motivated in Sec. 2.2.5. Since our later applications focus on Moiré systems (Chapter 7), we will adopt the spin-valley terminology throughout this section. The presentation follows Ref. [P3], building on conceptual work already developed during my master thesis.

We have already introduced SU(4) spin models and their representations in terms of spin-valley operators in Sec. 3.2. Because this construction is essential for what follows, we begin by briefly

recalling the key points here. A general SU(4) Hamiltonian takes the form

$$H = \frac{1}{2} \sum_{ij} \sum_{a,b=1}^{15} T_i^a J_{ij}^{ab} T_j^b, \quad (4.95)$$

where T_i^a are a basis of the $4^2 - 1 = 15$ generators of SU(4) (i.e., a basis of the Lie Algebra $\mathfrak{su}(4)$). As before, for an efficient treatment via the FRG, the first step is to represent them in terms of fermions. For the models of interest in this thesis, this is most conveniently done by introducing spin-valley operators defined by the parton decomposition

$$T^a \rightarrow \sigma_i^\mu \tau_i^\kappa = \sum_{ss',ll'} f_{isl}^\dagger \theta_{ss'}^\mu \theta_{ll'}^\kappa f_{is'l'}, \quad (4.96)$$

which provides a faithful basis of $\mathfrak{su}(4)$ if the fermion number is fixed to

$$\sum_{s,l} f_{isl}^\dagger f_{isl} = n \quad (4.97)$$

on each lattice site. Following our discussion of the single occupation constraint for SU(2) models in Sec. 4.2.8, we enforce this constraint on average by setting the chemical potential to zero, which restricts us to half-filling, i.e. $n = 2$. According to our discussion in Sec. 3.2, the dimension of the local Hilbert space is then $\binom{4}{2=6}$, which means we consider the six-dimensional representation of $\mathfrak{su}(4)$. Let us emphasize again that, in this representation, the spin operators $\sigma_i^\mu \tau_i^\kappa$ should be interpreted as a *single* operator acting on the local Hilbert space at site i , and can *not* be factorized to a tensor product $\sigma_i^\mu \tau_i^\kappa \neq \sigma_i^\mu \otimes \tau_i^\kappa$.

The idea of applying the pf-FRG to spin-valley models is now the same as before. The first step is to rewrite the general SU(4) Hamiltonian in terms of spin-valley operators which yields

$$H = \frac{1}{2} \frac{1}{8} \sum_{ij} \sum_{\mu\nu\kappa\eta} (\sigma_i^\mu \tau_i^\kappa) J_{ij}^{\mu\nu\kappa\eta} (\sigma_j^\nu \tau_j^\eta), \quad (4.98)$$

where we have to set $J^{00\kappa\eta} = J^{\mu\nu 00} = 0$ to have an exact mapping to the SU(4) Hamiltonian (4.95) and the additional prefactor 1/8 ensures consistency with the conventions often used in spin-valley models. The next steps follow the SU(2) logic: analyze the symmetry structure, derive symmetry-constrained parametrizations of the self-energy and vertex, and finally obtain the corresponding flow equations.

Let us summarize the central result before going into details: For the very general spin-valley Hamiltonian introduced above, straightforward generalizations of the particle-hole conjugation and time-reversal are no longer symmetries of the pseudo-fermion Hamiltonian, which would dramatically increase the numerical complexity. This can be remedied by restricting to Hamiltonians where all couplings with an odd number of indices (μ, ν, κ, η) equal to zero vanish, and which are either diagonal in spin or valley space, i.e.

$$J^{\mu\nu\kappa\eta} = J^{\mu\nu\kappa\eta} \delta_{\mu\nu} \quad \text{or} \quad J^{\mu\nu\kappa\eta} = J^{\mu\nu\kappa\eta} \delta_{\kappa\eta}. \quad (4.99)$$

For example, if the Hamiltonian is diagonal in the spin sector, the most general Hamiltonian we can treat with the pf-FRG takes the form

$$H = \frac{1}{16} \sum_{ij} \left[\sum_{a,b,c=1}^3 (\sigma_i^a \tau_j^b) J_{ij}^{abc} (\sigma_j^a \tau_j^c) + \sum_a \sigma_i^a J_{\sigma,ij}^a \sigma_j^a + \sum_{b,c} \tau_i^b J_{\tau,ij}^{bc} \tau_j^c \right]. \quad (4.100)$$

Here, we introduced $J_{\sigma,ij}^a$ coupling the pure spin terms $\sigma_i^a \equiv \sigma_i^a \tau_i^0$, $J_{\tau,ij}^{bc}$ coupling the pure valley terms $\tau_i^b \equiv \sigma_i^0 \tau_i^b$, and J_{ij}^{abc} coupling the mixed spin-valley terms with $a, b, c > 0$. While this restriction reduces the generality of the approach, it is well aligned with many effective models relevant for moiré materials, which indeed take precisely this form. Prominent examples include twisted bilayer graphene (TBG) [48] and trilayer graphene aligned with hexagonal boron nitride (TG/h-BN) [49]. For Hamiltonians of this structure, almost all symmetries present in SU(2) spin models remain intact, allowing us to employ essentially the same self-energy and vertex parametrization, symmetry relations, and flow equations as in the SU(2) case. The only differences are that vertices now carry additional valley indices, $\Gamma^{\mu\nu} \rightarrow \Gamma^{\mu\nu\kappa\eta}$, and certain frequency symmetries are no longer preserved. Fortunately, the missing symmetry constraint increases the numerical cost by only about a factor of two.

In the following, we derive this result step by step, closely following the SU(2) logic: we first generalize the relevant symmetry transformations to spin-valley Hamiltonians, then work out the constraints they impose on correlation functions, and finally formulate efficient parametrizations of the self-energy and vertex and derive the corresponding flow equations.

4.3.1 Symmetry constraints on the correlation functions

Let us start the symmetry analysis of the generic spin-valley Hamiltonian (4.98) by noting which symmetries remain unchanged compared to the SU(2) pseudo-fermion Hamiltonian. These are all transformations that do not act on the spin or valley quantum numbers of the fermionic operators: the local U(1) gauge redundancy, imaginary time-translation invariance, Hermitian conjugation, crossing symmetry (which now also exchanges valley indices), and lattice symmetries. Each of these continues to leave the spin-valley Hamiltonian (or the action) invariant, and therefore imposes exactly the same constraints on the correlation functions as listed in Table 4.2.

In contrast, particle-hole conjugation and time-reversal symmetry act directly on the spin indices. To extend them to the spin-valley case, we define straightforward generalizations that act in the same way on both spin and valley indices. For particle-hole conjugation this yields

$$f_{isl}^\dagger \xrightarrow{\text{PH}} sl f_{i\bar{s}\bar{l}} \quad f_{isl} \xrightarrow{\text{PH}} sl f_{i\bar{s}\bar{l}}^\dagger, \quad (4.101)$$

where spin and valley indices take values $s = \pm 1$ (corresponding to $s = \uparrow, \downarrow$) and $l = \pm 1$ (corresponding to $l = +, -$) and we define $\bar{l} = -l$ and $\bar{s} = -s$ as before. Inserting this into the definition of the spin-valley operators (4.96) and using Eq. (4.42), we obtain

$$\sigma_i^\mu \tau_i^\kappa \xrightarrow{\text{PH}} -\xi(\mu)\xi(\kappa) \sigma_i^\mu \tau_i^\kappa, \quad (4.102)$$

with the sign function $\xi(\mu)$ defined in Eq. (4.41). Spin-valley operators with either $\mu = 0$ or $\kappa = 0$ reduce to the familiar SU(2) spin operators (acting in the enlarged SU(4) Hilbert space) and are invariant, as expected. General spin-valley operators, however, may change sign. The Hamiltonian is therefore not invariant under the *local* PH conjugation. This can be remedied by promoting the transformation to a *global* one, applied simultaneously to all lattice sites. A pairwise interaction of spin-valley operators then transforms as

$$(\sigma_i^\mu \tau_i^\kappa)(\sigma_j^\nu \tau_j^\eta) \xrightarrow{\text{PH}} \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) (\sigma_i^\mu \tau_i^\kappa)(\sigma_j^\nu \tau_j^\eta). \quad (4.103)$$

Whenever an odd number of indices μ, ν, κ, η are nonzero, the interaction acquires a minus sign, otherwise it remains unchanged. This motivates the restricted Hamiltonian introduced in Eq. (4.100), where precisely these sign-changing terms are excluded, ensuring invariance under

global particle–hole conjugation. In the same way we derived the resulting constraints on the correlation functions in Sec. 4.2.2, we find that this implies

$$G(x'_1; x_1) \delta_{i'_1 i_1} = -ss' ll' G(\mathcal{T}x_1; \mathcal{T}x'_1) \delta_{i'_1 i_1} \quad (4.104)$$

$$G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = s'_1 s_1 l'_1 l_1 s'_2 s_2 l'_2 l_2 G(\mathcal{T}x_1, \mathcal{T}x_1; \mathcal{T}x'_1, \mathcal{T}x'_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2}, \quad (4.105)$$

where we again use the definition $\mathcal{T}x_j = (i_j, -\omega_j, \bar{s}, \bar{l})$.

We can also define a generalized anti-unitary time-reversal symmetry, which acts by complex conjugation on all numbers and on the fermionic operators as

$$f_{isl}^\dagger \xrightarrow{\text{TR}} e^{i\pi s/2} e^{i\pi l/2} f_{i\bar{s}\bar{l}}^\dagger \quad f_{isl} \xrightarrow{\text{TR}} e^{-i\pi s/2} e^{-i\pi l/2} f_{i\bar{s}\bar{l}}. \quad (4.106)$$

Using the identity $e^{i\pi/2(s-s')} = ss'$, it is straightforward to show that the spin-valley operator transforms as

$$\sigma_i^\mu \tau_i^\kappa \xrightarrow{\text{TR}} \xi(\mu) \xi(\kappa) \sigma_i^\mu \tau_i^\kappa. \quad (4.107)$$

For $\mu = 0$ or $\kappa = 0$, this reduces to the conventional SU(2) time-reversal symmetry, while in general some components acquire an additional minus sign. Pairwise interactions, however, transform exactly as under global particle–hole conjugation, and thus the restricted Hamiltonian (4.100) remains invariant by the same arguments. The induced symmetry relations for the correlation functions are

$$G(x'_1; x_1) \delta_{i'_1 i_1} = -ss' ll' G(\mathcal{T}x'_1; \mathcal{T}x_1) \delta_{i'_1 i_1} \quad (4.108)$$

$$G(x'_1, x'_2; x_1, x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2} = s'_1 s_1 l'_1 l_1 s'_2 s_2 l'_2 l_2 G(\mathcal{T}x'_1, \mathcal{T}x'_2; \mathcal{T}x_1, \mathcal{T}x_2) \delta_{i'_1 i_1} \delta_{i'_2 i_2}, \quad (4.109)$$

Comparing these relations with their SU(2) counterparts in Table 4.2, we find that the constraints on the correlation functions remain essentially the same. The only difference is that the local particle–hole conjugations (PH1, PH2) no longer appear separately. Instead, only the global PH acting simultaneously on both sites survives. We discuss the consequences of this restriction for the vertex parametrization in the next section.

4.3.2 Vertex parametrization

We now define the parametrization of the correlation functions in direct analogy to the SU(2) case [see Eqs. (4.40, 4.44)]. For the one-particle correlator this reads

$$G(x', x) = \sum_{\mu, \kappa=0}^3 G_i^{\mu\kappa}(\omega) \theta_{s's}^\mu \theta_{l'l}^\kappa \delta_{i'i} \delta_{\omega'\omega}, \quad (4.110)$$

and for the two-particle correlator

$$G(x'_1, x'_2; x_1, x_2) = \sum_{\mu, \nu, \kappa, \eta=0}^3 G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\nu \theta_{l'_1 l_1}^\kappa \theta_{l'_2 l_2}^\eta \delta_{i'_1 i_1} \delta_{i'_2 i_2} \delta_{\omega'_1 + \omega'_2, -\omega_1 - \omega_2}. \quad (4.111)$$

We derive symmetry constraints by inserting this parametrization into the relations given in Table 4.2, but with PH1 and PH2 replaced by the global PH, and TR replaced by the generalized TR defined above. The resulting symmetry constraints on $G_i^{\mu\kappa}$ and $G_{i_1 i_2}^{\mu\nu\kappa\eta}$ are summarized in

Table 4.5 – Symmetry relations on the parametrized one- and two-particle correlation functions for spin-valley models. These relations hold for the slightly restricted spin-valley Hamiltonian (4.100). The corresponding symmetry relations for SU(2) spin models are shown in Table 4.3. The symmetries/gauge transformations used in the derivation are indicated by the labels on the right, which are defined in Table 4.1 except that PH now denotes the global particle-hole conjugation defined in (4.101) and TR the generalized time-reversal defined in (4.106). The $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$ transformation is only a symmetry of Hamiltonians with couplings diagonal in the spin indices.

| | |
|--|---|
| $G_i^{\mu\kappa}(\omega) = \xi(\mu)\xi(\kappa)G_i^{\mu\kappa}(\omega)$ | $(H \circ TR)$ |
| $G_i^{\mu\kappa}(\omega) = -G_i^{\mu\kappa}(-\omega)$ | $(H \circ TR \circ PH)$ |
| $G_i^{\mu\kappa}(\omega) = -G_i^{\mu\kappa}(\omega)^*$ | $(TR \circ PH)$ |
| $G_i^{\mu\kappa}(\omega) = G_i^{00}(\omega)\delta_{\mu,0}\delta_{\kappa,0}$ | $(\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma)$ |
| $G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u)^*$ | $(TR \circ PH \circ H \circ TR)$ |
| $G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) = G_{i_2 i_1}^{\nu\mu\eta\kappa}(-s, t, u)$ | $(H \circ TR \circ PH \circ X)$ |
| $G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, -t, u)$ | $(H \circ TR)$ |
| $G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) G_{i_2 i_1}^{\nu\mu\eta\kappa}(s, t, -u)$ | $(H \circ TR \circ X)$ |
| $G_{i_1 i_2}^{\mu\nu\kappa\eta}(s, t, u) = G_{i_1 i_2}^{\mu\mu\kappa\eta}(s, t, u)\delta_{\mu\nu}$ | $(\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma)$ |

Table 4.5. Compared to the relations for SU(2) spin models in Table 4.3, there are two main differences.

First, the two-particle correlation function does not possess the symmetry that exchanges $s \leftrightarrow u$. In practice, this simply means we have to compute the vertex for twice as many frequencies.

Second, and more severe, the one-particle correlator $G_i^{\mu\kappa}$ is not automatically diagonal in spin and valley indices due to the appearance of *two* sign functions $\xi(\mu)\xi(\kappa)$ in the first symmetry relation, which allow components with both $\mu > 0$ and $\kappa > 0$ to be nonzero. This would necessitate tracking multiple independent self-energy components in the flow equations, significantly increasing the computational cost. In addition, the argument in Sec. 4.2.3 that the symmetries of the disconnected correlators carry over to the self-energy and vertex relied on G being fully diagonal. Without this property, an additional dedicated symmetry analysis of the self-energy would be required.

Fortunately, most material-relevant models possess additional spin- and valley-space symmetries that themselves impose diagonality on one-particle terms. As shown in Sec. 4.2.7, for an SU(2) model with only diagonal couplings the $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry enforces $G^\mu = G^0\delta_{\mu,0}$. It is straightforward to show that the same derivation goes through exactly the same way for a spin-valley model that is diagonal either in spin space ($J^{\mu\nu\kappa\eta} = J^{\mu\nu\kappa\eta}\delta_{\mu\nu}$) or in valley space ($J^{\mu\nu\kappa\eta} = J^{\mu\nu\kappa\eta}\delta_{\kappa\eta}$). In the following, we restrict to models diagonal in spin space, denoting the symmetry by $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$. This symmetry enforces

$$G_i^{\mu\kappa}(\omega) = G_i^{0\kappa}(\omega)\delta_{0,\mu}, \quad (4.112)$$

which, combined with the first symmetry relation in Table 4.5, further reduces to

$$G_i^{\mu\kappa} = G_i^{00}(\omega)\delta_{0,\mu}\delta_{0,\kappa}, \quad (4.113)$$

and thus a completely diagonal one-particle correlator. We can then employ the full SU(2)-based pf-FRG formalism from the previous section. Restricting further to Archimedean lattices renders

G also site-independent, giving the final parametrization

$$G(x', x) = G(\omega) \delta_{s' s} \delta_{l' l} \delta_{i' i} \delta_{\omega' \omega}, \quad (4.114)$$

where, as in the SU(2) case, $G(\omega) \in i\mathbb{R}$ and $G(\omega) = -G(-\omega)$.

Again following the same reasoning as in Sec. 4.2.7, the $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$ symmetry constrains the two-particle correlation function to be diagonal in spin space as

$$G_{i_1 i_2}^{\mu \nu \kappa \eta}(s, t, u) = G_{i_1 i_2}^{\mu \mu \kappa \eta}(s, t, u) \delta_{\mu \nu} \equiv G_{i_1 i_2}^{\mu \kappa \eta}(s, t, u) \delta_{\mu \nu} \quad (4.115)$$

Accordingly, we parametrize it as

$$G(x'_1, x'_2; x_1, x_2) = \sum_{\mu, \kappa, \eta=0}^3 G_{i_1 i_2}^{\mu \kappa \eta}(s, t, u) \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\mu \theta_{l'_1 l_1}^\kappa \theta_{l'_2 l_2}^\eta \delta_{i'_1 i_1} \delta_{i'_2 i_2} \delta_{\omega'_1 + \omega'_2, -\omega_1 - \omega_2}. \quad (4.116)$$

In this case the symmetry relations of Table 4.5 simplify via $\xi(\mu)\xi(\nu) \rightarrow \xi(a)\xi(a) = 1$, and coincide with those for SU(2) models (Table 4.3), except for the absence of the $s \leftrightarrow u$ symmetry. Most notably, we again find

$$G_{i_1 i_2}^{\mu \kappa \eta} \in \begin{cases} \mathbb{R} & \text{if } \xi(\kappa)\xi(\eta) = 1 \\ i\mathbb{R} & \text{if } \xi(\kappa)\xi(\eta) = -1 \end{cases}, \quad (4.117)$$

and all negative frequencies can be mapped to positive ones. By the same arguments as in Sec. 4.2.3, the self-energy and vertex inherit these parametrizations and symmetry relations.

4.3.3 Flow equations

We have shown that for $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$ symmetric spin-valley models, the parametrizations of the self-energy and vertex are essentially identical to those in the SU(2) case. Consequently, the flow equations derived in Sec. 4.2.4 carry over with only slight modifications due to additional valley indices.

For the channel-decomposed vertex flow equations [Eqs. (4.51, 3.11, 4.53)], we did not explicitly perform the sums over spin indices. The generalization to spin-valley models is therefore simply obtained by the substitution

$$\Gamma^{\mu \nu} \rightarrow \Gamma^{\mu \kappa \eta}, \quad \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\nu \rightarrow \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\mu \theta_{l'_1 l_1}^\kappa \theta_{l'_2 l_2}^\eta, \quad (4.118)$$

which directly yields the corresponding equations for the spin-valley case.

For the self-energy flow, by contrast, the sums over spin and valley indices can be carried out analytically. Many terms vanish due to properties of the Pauli matrices and the symmetry constraints on the vertex. The calculation is relatively involved, and here we only state the final result, which reads

$$\frac{d}{d\Lambda} \Sigma^\Lambda(\omega) = -\frac{1}{2\pi} \int d\omega' \left[4 \sum_j \Gamma_{ij}^{\Lambda 000}(\omega' + \omega, \omega' - \omega, 0) - \sum_{\mu, \kappa} \Gamma_{ii}^{\Lambda \mu \kappa \kappa}(\omega' + \omega, 0, \omega - \omega') \right] S^\Lambda(\omega'), \quad (4.119)$$

where i is an arbitrary reference site. Details of the calculation can be found in the appendix of Ref. [P3]. Consistent with the self-energy flow equation (4.49) for the SU(2) case, only vertex components diagonal in spin and valley contribute to the flow.

Table 4.6 – Symmetry relations of the channel-resolved two-particle vertex in the asymptotic frequency parametrization for spin-valley models. Shown are transformation properties of $g_{c,i_1 i_2}^{\mu \nu \kappa \eta}$ in the s -, t -, and u -channels, derived from the frequency symmetries of the full vertex in Table 4.5. The corresponding relations for SU(2) spin models are shown in Table 4.4.

| |
|--|
| $g_{s,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(s, \nu_s, \nu'_s) = g_{s,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(-s, \nu_s, \nu'_s)$ |
| $g_{s,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(s, \nu_s, \nu'_s) = g_{s,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(s, -\nu_s, -\nu'_s)$ |
| $g_{s,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(s, \nu_s, \nu'_s) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{s,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(s, \nu'_s, \nu_s)$ |
| $g_{t,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(t, \nu_t, \nu'_t) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{t,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(-t, \nu_t, \nu'_t)$ |
| $g_{t,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(t, \nu_t, \nu'_t) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{t,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(t, -\nu_t, -\nu'_t)$ |
| $g_{t,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(t, \nu_t, \nu'_t) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{t,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(t, \nu'_t, \nu_t)$ |
| $g_{u,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(u, \nu_u, \nu'_u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{u,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(-u, \nu_u, \nu'_u)$ |
| $g_{u,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(u, \nu_u, \nu'_u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{u,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(u, -\nu_u, -\nu'_u)$ |
| $g_{u,i_1 i_2}^{\Lambda, \mu \nu \kappa \eta}(u, \nu_u, \nu'_u) = \xi(\mu)\xi(\nu)\xi(\kappa)\xi(\eta) g_{u,i_2 i_1}^{\Lambda, \nu \mu \eta \kappa}(u, \nu'_u, \nu_u)$ |

The initial conditions follow from Eq. (4.19), after antisymmetrizing the interaction terms of the Hamiltonian. For the generic spin-valley Hamiltonian (4.98), this yields

$$\begin{aligned} \Sigma^{\Lambda \rightarrow \infty} &= 0 \\ \Gamma_{ij}^{\Lambda \rightarrow \infty \mu \kappa \eta} &= \frac{1}{8} J_{ij}^{\mu \mu \kappa \eta}, \end{aligned} \quad (4.120)$$

while for the restricted $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$ symmetric Hamiltonian (4.100), this translates to

$$\Gamma_{ij}^{\Lambda \rightarrow \infty \mu \kappa \eta} = \frac{1}{8} \begin{cases} J_{ij}^{\mu \kappa \eta} & \text{if } \mu, \kappa, \eta > 0 \\ J_{\sigma,ij}^\mu & \text{if } \mu > 0 \text{ and } \kappa = \eta = 0 \\ J_{\tau,ij}^{\kappa \eta} & \text{if } \kappa, \eta > 0 \text{ and } \mu = 0 \\ 0 & \text{if } \mu = \kappa = \eta = 0 \end{cases}, \quad (4.121)$$

Since the initial conditions are frequency-independent, we have omitted the frequency arguments of Σ and Γ above.

4.3.4 Asymptotic frequency parametrization

Also for spin-valley models it is beneficial to employ the asymptotic frequency parametrization discussed in Sec. 4.2.6. We, therefore, also need to state the frequency symmetries in this parametrization, which slightly change compared to SU(2) models due to the absence of the $s \leftrightarrow u$ symmetry. The results are summarized in Table 4.6. Note that if we include the $\mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma \times \mathbb{Z}_2^\sigma$ symmetry (valid for models diagonal in spin space) they again significantly simplify by setting $\mu = \nu$.

The most notable difference to the SU(2) case is that instead of *separate* relations reversing the sign of the fermionic frequencies, $\nu_c \rightarrow -\nu_c$ and $\nu'_c \rightarrow -\nu'_c$, only the *combined* transformation $\nu_c, \nu'_c \rightarrow -\nu_c, -\nu'_c$ remains. Consequently, while for SU(2) it was sufficient to restrict both fermionic frequencies to be positive, here we must also allow one of them to be negative. Here, we choose ν'_c . Furthermore, the s - and u -channels are now independent and no longer related by symmetry. Taken together, these modifications increase the numerical cost by roughly a factor of two. What remains unchanged is that only positive bosonic frequencies ω_c need to be considered, and the exchange symmetry $\nu_c \leftrightarrow \nu'_c$ still allows us to restrict the calculation to $\nu_c \leq |\nu'_c|$.

4.4 Numerical implementation

The symmetry analysis presented in Sec. 4.2 and Sec. 4.3 already allowed us to substantially reduce the number of self-energy and vertex components required in solving the pf-FRG flow equations. Nevertheless, even after fully exploiting the symmetries of the pseudo-fermion Hamiltonian, significant numerical challenges remain. First, the vertex flow equations (4.2.4) contain sums over an infinite lattice that must be approximated. Second, the vertex depends on three continuous Matsubara frequencies, requiring a suitable discretization. Third, the flow equations involve frequency integrals over products of vertices and propagator bubbles, which become increasingly singular as the cutoff Λ is lowered. To accurately perform these integrals, the vertex must be evaluated at frequencies not on the discrete frequency grid, necessitating (multi-)linear interpolation.

All of these ingredients must be implemented with both high accuracy—to capture the often subtle structures of pf-FRG vertices—and high efficiency, as the computational cost otherwise quickly grows prohibitively large. For example, a typical frequency grid of size $40 \times 35 \times 35$ and a lattice of a few hundred sites already yields $\mathcal{O}(10^7)$ independent vertex components, and therefore an equally large system of coupled differential equations. Solving such systems requires HPC resources and a highly parallelized code.

Fortunately for me, I did not have to write such an implementation from scratch, but could instead extend the Julia package `PFFRGSolver.jl`[C3], originally written by my former colleague Dominik Kiese, and Tobias Müller. This package implements state-of-the-art integration routines for the pf-FRG flow equations, employing an error-controlled adaptive ODE solver together with adaptive frequency grids and adaptive frequency integration. The code is efficiently parallelized at the shared-memory level, using multiple threads within a single compute node, and shows near-perfect scaling up to 128 cores. Despite these optimizations, typical computation times remain substantial: for example, a high-resolution calculation with a frequency grid of $40 \times 35 \times 35$ and about 700 correlated sites requires roughly 5 hours (≈ 640 core-hours) for highly symmetric Heisenberg models, but up to 7 days ($\approx 22\,000$ core-hours) for less symmetric spin models (such as the anisotropic nearest-neighbor pyrochlore studied in Chapter 5) on an AMD EPYC Milan node with 128 cores. For full phase diagrams, the total computational cost can easily reach several million core-hours.

A detailed discussion of the numerical algorithms used in the `PFFRGSolver.jl` package is presented in Dominik Kiese’s PhD thesis [116], with a more concise presentation available in Ref. [129]. In the following, we therefore only shortly summarize the employed algorithms and discretization schemes, before turning to a methodological development that I contributed in the course of this thesis: an implementation of combined spin–lattice symmetries tailored for spin-orbit entangled Mott insulators as will be studied in Chapter 5. In such systems, strong spin-orbit coupling leads to bond-dependent, nondiagonal interactions that individually break the lattice space-group symmetries. Crucially, these symmetries are restored only when combined with appropriate spin rotations. A prototypical model where this is the case is the Kitaev honeycomb model. In pf-FRG, exploiting these combined symmetries is absolutely essential for obtaining reasonable runtimes, particularly for complex three-dimensional models.

We emphasize that we are not the first to address this problem, and similar models where already studied with the pf-FRG before [118–122]. To our knowledge, however, there is no *generic*, publicly available implementation capable of handling essentially arbitrary spin models of the form (4.26). While Buessen’s `SpinParser` software [117] features the implementation of combined spin and lattice symmetries, it is restricted to transformations that simply permute the spin indices. This suffices for models such as the Kitaev model, but fails for more gen-

eral Hamiltonians—such as the nearest-neighbor models for rare-earth pyrochlores studied in Chapter 5—where more general spin rotations need to be considered.

Closely building on existing algorithms [14, 116, 117], we developed a routine that, for an arbitrary spin model on any lattice graph, automatically detects all combined spin–lattice transformations that leave the Hamiltonian invariant, and generates the necessary information to exploit these symmetries in pf-FRG (and related) calculations.

4.4.1 Algorithms in the PFFRGSolver.jl package

Before discussing new implementation of lattice symmetries, let us briefly summarize the core algorithms implemented in `PFFRGSolver.jl`. These algorithms provide the basis for solving the pf-FRG flow equations with both accuracy and efficiency.

ODE solver While the pf-FRG flow equations constitute a massive system of coupled differential equations, the overall structure is, fortunately, comparably simple: the equations are simple, first-order, ordinary differential equations of the form

$$\partial_\Lambda \gamma(\Lambda) = f(\Lambda, \gamma(\Lambda)), \quad (4.122)$$

where $\gamma(\Lambda)$ collects all components of the self-energy and vertex. Since the flow spans multiple energy scales, it is advantageous to use an ODE solver with adaptive step size. Early implementations relied on simple adaptive Euler algorithms [13, 117], which already proved successful. However, more recent work [142]—and our own experience—shows that multistep Runge-Kutta methods achieve similar accuracy but with significantly improved efficiency. In practice, we employ the Bogacki-Shampine RK3(2) method [143] with error-controlled step-size adjustment as described in Ref. [144], using a tolerance of 10^{-8} .

Frequency discretization The continuous frequency variables of the vertex ω_c, ν_c, ν'_c and the frequency of the self-energy ω are approximated on discrete grids with a fixed number of points n_ω, n_ν, n'_ν and n_Σ . For the smooth cutoff $e^{-\omega/\Lambda^2}$ introduced in Sec. 4.2.5, sharp features appear in the propagator bubble $P(\omega, \nu)$ [defined in Eq. (4.54)] near $\omega \approx \Lambda$ and shift towards lower frequencies as Λ decreases. It is therefore essential to adapt the grid dynamically during the flow. In our implementation, we use a mixed linear–logarithmic scheme: for a grid of n points, the first 40% are distributed linearly from the origin up to ω_{lin} , while the remaining points are spaced logarithmically up to a maximum frequency ω_{max} . Both ω_{lin} and ω_{max} are updated throughout the flow by a heuristic procedure that scans for sharp features in the vertex and adjusts the grid accordingly (for details, see Ref. [116]). In practice, convergence of spin–spin correlations is typically reached with grids of size $40 \times 35 \times 35$, although qualitatively stable results are already obtained for somewhat smaller grids.

Frequency integration The frequency integration in the self-energy flow is performed using the adaptive Gauss-Kronrod quadrature provided by the Julia package `QuadGK.jl`. For the channel integrals of the vertex, it is more efficient to employ an adaptive Simpson rule [116], where the number of integration points is doubled until either an absolute tolerance of 10^{-8} or a relative tolerance of 10^{-3} is achieved. During the frequency integration, vertices must often be evaluated at points lying outside the numerical frequency grid. In such cases, we approximate their values using multilinear interpolation.

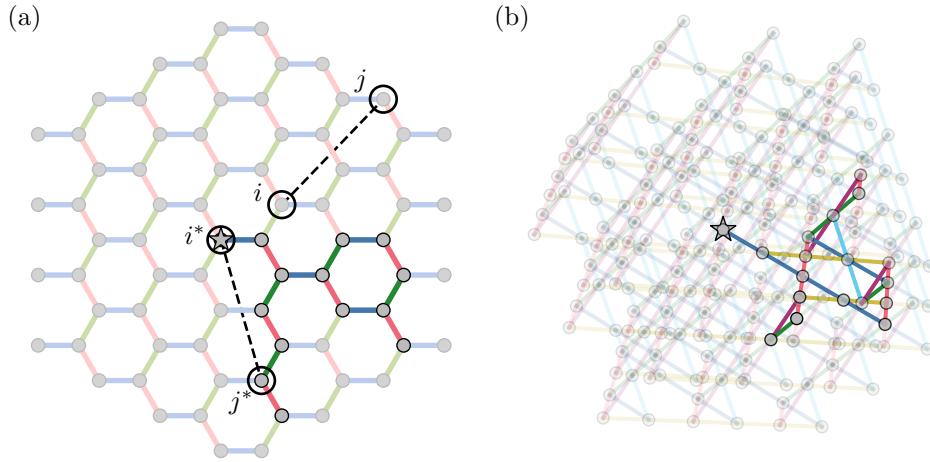


Figure 4.1 – Symmetry-reduced site pairs on the honeycomb and pyrochlore lattices. Bonds are colored according to the anisotropic couplings present in models such as the Kitaev model (a) or the most general symmetry-allowed nearest-neighbor model on the pyrochlore lattice (b). Both lattices feature a single symmetry-inequivalent reference site i^* in the unit cell (marked by a star). The full set of site pairs \mathcal{P} up to maximal bond distances of $L = 6$ (a) and $L = 4$ (b) is shown in pale colors. Using combined spin- and real-space symmetries together with the algorithm described in the main text, these sets reduce to minimal symmetry-inequivalent subsets \mathcal{P}^* , highlighted in saturated colors, shrinking from 74 to 14 pairs in (a) and from 220 to 18 pairs in (b). In panel (a), an example of two symmetry-equivalent pairs $(ij) \in \mathcal{P}$ and $(i^*j^*) \in \mathcal{P}^*$ is highlighted.

4.4.2 Exploiting real- and spin-space symmetries in spin-orbit coupled materials

The remaining challenge in solving the flow equations lies in treating the, in principle, infinite crystal lattice on which the self-energy and vertex are defined. We already showed that for Archimedean lattices the self-energy becomes site-independent. The vertex, however, still retains a site-dependence, and the flow equation for the t -channel (4.52) additionally contains a sum over all sites. In this section, we present a two-step strategy to make this problem numerically tractable. First, by exploiting the translational symmetry of the Bravais lattice and imposing a maximal correlation length, we reduce the infinite lattice to a finite set \mathcal{P} of site pairs (ij) for which bilocal vertex components $\Gamma_{ij}^{\mu\nu}$ must be computed. Second, by further exploiting combined real-space and spin space symmetries, we restrict this to a smaller set of symmetry inequivalent site pairs \mathcal{P}^* . The vertex on the full set \mathcal{P} can then be reconstructed from \mathcal{P}^* by symmetry transformations, often reducing the numerical cost dramatically. Examples of \mathcal{P} and \mathcal{P}^* for the Kitaev honeycomb model and for the most general symmetry-allowed nearest-neighbor model on the pyrochlore lattice are shown in Fig. 4.1, highlighting the enormous reduction in vertex components.

Although in practice we will restrict to Archimedean lattices, the arguments presented here hold more generally for arbitrary crystal lattices. Our discussion will be in terms of $SU(2)$ spin models, but can be straightforwardly generalized also to generic spin-valley models.

Combined real- and spin space symmetries To start, let us recapitulate how the vertex transforms under real-space and spin space transformations. Following the discussion of Sec. 4.3.1, the vertex transforms under lattice isomorphism \mathcal{L} (including translations, rotations and inversions) as

$$\Gamma_{ij}^{\mu\nu} \xrightarrow{\mathcal{L}} \Gamma_{L(i)L(j)}^{\mu\nu}, \quad (4.123)$$

where we suppress the frequency dependence for brevity. Since spin operators are angular momenta, they transform under spatial rotations as well. For effective spin models with anisotropic (nondiagonal) interactions, the real-space transformations of the lattice must therefore be accompanied by an $SU(2)$ rotation \mathbf{U} acting on the spin Hilbert space, which, as discussed in Sec. 4.2.7, corresponds to an $SO(3)$ rotation \mathbf{R} acting on spin operators as

$$\mathbf{S} \rightarrow \mathbf{R}\mathbf{S}. \quad (4.124)$$

Under combined lattice and spin transformations, the vertex then transforms as

$$\Gamma_{ij}^{\mu\nu} \xrightarrow{R \circ L} \sum_{\kappa\eta} R_{\mu\kappa}^T \Gamma_{L(i)L(j)}^{\kappa\eta} R^{\eta\nu}, \quad (4.125)$$

where $R^{\mu\nu}$ acts trivially on the density indices $\mu, \nu = 0$. The transformation $R \circ L$ is a symmetry of the Hamiltonian if the coupling matrices satisfy

$$\mathbf{J}_{ij} = \mathbf{R}^T \mathbf{J}_{L(i)L(j)} \mathbf{R}. \quad (4.126)$$

In this case, two site pairs (ij) and (i^*j^*) are called symmetry equivalent if they can be mapped onto each other by such a transformation. Consequently, it suffices to compute the vertex only on a minimal set of symmetry-*inequivalent* pairs, reconstructing the rest by symmetry.

In principle, arbitrary rotations \mathbf{R} could be considered, but for most spin models it turns out to be sufficient to only consider a much smaller subgroup of $SU(2)$ [or $SO(3)$], which is far more efficient numerically. Concretely, we restrict to rotations by $\pi/2$ (or multiples thereof) around the x , y , and z axes. As an example, a $\pi/2$ -rotation around the z axis maps the spin operators to

$$S^x \rightarrow S^y \quad S^y \rightarrow -S^x \quad S^z \rightarrow S^z, \quad (4.127)$$

corresponding to an odd permutation of spin indices plus a sign change. The combination of all of those rotations therefore generate:

1. all even permutations of the spin indices $xyz \rightarrow (xyz, zxy, yzx)$ accompanied by an *even* number of sign flips, and
2. all *odd* permutations of the spin indices $xyz \rightarrow (zyx, xzy, yxz)$ combined with an *odd* number of sign flips.

Geometrically, these operations correspond to all rotational symmetries of a cube. They form the octahedral rotation group $O \subset SO(3)$, which contains 24 elements. We specify elements $R = (p, \zeta) \in O$ of this group by a permutation $p(\mu)$ and a sign function $\zeta(\mu)$, that encode the permutation of spin indices and sign flips. In this notation, the vertex (and equivalently the coupling matrix) then transform as

$$\Gamma_{ij}^{\mu\nu} \xrightarrow{R \circ L} \zeta(\mu)\zeta(\nu)\Gamma_{L(i)L(j)}^{p(\mu)p(\nu)}, \quad (4.128)$$

which is numerically advantageous compared to general $SU(2)$ rotations, as each vertex component maps onto exactly one other component rather than a sum over several.

In the next steps, we explain how to construct the finite list of symmetry-*inequivalent* pairs \mathcal{P}^* by combining these restricted spin rotations with the lattice symmetries of the underlying Bravais lattice.

Lattice truncation To first reduce the infinite lattice to a finite set of site pairs \mathcal{P} , we exploit the translational invariance present on any Bravais lattice. This allows us to map pairs (ij) to equivalent pairs (i^*j^*) , where the first site i^* is fixed to an arbitrary reference unit cell (UC). This limits the number of sites i in the first argument of the vertex. To further limit the number of sites j in the second argument, we set vertices $\Gamma_{ij}^{\mu\nu} = 0$ if the bond-distance between the sites exceeds L , i.e. $\|i - j\|_b < L$, where the bond-distance $\|\cdot\|_b$ is defined by the minimal number of nearest-neighbor bonds that connect the sites i and j . This effectively enforces a maximal correlation length $\xi_L \sim L$, which is a good approximation in phases without long-range order, but will lead to artifacts when approaching a magnetically ordered phase (as we discuss in Sec. 4.5). Compared to conventional truncations where the Hamiltonian itself is restricted to a finite lattice with periodic, open, or similar boundary conditions, enforcing a maximal correlation length introduces much weaker boundary effects and still allows pf-FRG to capture ordered states that are incommensurate with any finite lattice. To summarize, for a fixed L the finite set of pairs on which we must evaluate the vertex is defined as

$$\mathcal{P} = \{(ij) \mid i \in \text{UC}, \|i - j\|_b < L\} . \quad (4.129)$$

Determining the set of symmetry inequivalent pairs In the next step, we show how to reduce \mathcal{P} to the smaller set of symmetry-inequivalent pairs \mathcal{P}^* by exploiting the combined real- and spin space symmetries of the Hamiltonian. The general algorithm consists of three steps:

1. Determine a minimal set of symmetry inequivalent *reference sites* $\mathcal{R} = \{i^*\}$ inside the reference unit cell (for Archimedean lattices this is just one site), and the corresponding symmetry transformations to map all sites i in the reference unit cell to \mathcal{R}
2. For each reference site $i^* \in \mathcal{R}$, determine all point-group symmetries (combined with spin rotations) of the Hamiltonian that leave i^* invariant.
3. Use these point-group symmetries from step 2 to obtain a minimal set of symmetry inequivalent pairs $\mathcal{P}_{i^*}^* = \{i^*j_1^*, i^*j_2^*, \dots\}$ to which any pair (i^*j) can be mapped.

We then define the final set of symmetry irreducible site pairs as

$$\mathcal{P}^* = \{(i^*j^*) \mid i^* \in \mathcal{R} \text{ and } j^* \in \mathcal{P}_{i^*}^*\} . \quad (4.130)$$

Any pair $(ij) \in \mathcal{P}$ can then be mapped onto a pair $(i^*j^*) \in \mathcal{P}^*$ by first applying the transformations determined in step 1 to map i onto a reference site i^* , yielding (i^*j') , and then applying the point-group transformation from step 2 to map the site j' to a site $j^* \in \mathcal{P}_{i^*}^*$. We now discuss the detailed implementation of each step, assuming that the lattice symmetries are not known *a priori* but are instead determined by the algorithm.

Step 1: Determining symmetry inequivalent sites in the unit cell To determine whether two sites i and j in the unit cell are symmetry equivalent, without assuming prior knowledge of the lattice symmetries, we must identify transformations that map both their positions and bond environments to each other. To this end, we perform the following steps:

1. Select two nearest neighbors (i_1, i_2) of i and two nearest neighbors (j_1, j_2) of j .
2. Check if a rotation matrix $R_{\text{space}} \in SO(3)$ (optionally combined with an inversion I) exists that maps the vectors $\Delta r_{i_1} = \mathbf{r}_{i_1} - \mathbf{r}_i$ and $\Delta r_{i_2} = \mathbf{r}_{i_2} - \mathbf{r}_i$ onto $\Delta r_{j_1} = \mathbf{r}_{j_1} - \mathbf{r}_j$ and

$\Delta r_{j_2} = \mathbf{r}_{j_2} - \mathbf{r}_j$, using Rodrigues' formula [145]. If successful, define $L = T \circ R_{\text{space}}(\circ I)$, where T is a translation by $\mathbf{r}_j - \mathbf{r}_i$, which maps i to j (including the two nearest neighbor bonds).

3. Check whether any of the 24 elements $R_{\text{spin}} \in O$ of the octahedral rotation group map the corresponding coupling matrices $\mathbf{J}_{ii_1} \xrightarrow{R_{\text{spin}}} \mathbf{J}_{jj_1}$ and $\mathbf{J}_{ii_2} \xrightarrow{R_{\text{spin}}} \mathbf{J}_{jj_2}$ onto each other.
4. If both are successful, verify that the corresponding lattice transformation $L \circ R_{\text{spin}}$ is a symmetry of the full lattice (in practice we check this on a finite test set of sites).
5. If this is also successful, store the pair $(i \rightarrow j)$ and the transformation R_{spin} .

If any of the steps fail, repeat the same procedure for all other possible nearest neighbor pairs i_1, i_2 and j_1, j_2 . Doing this for all sites i, j in the reference unit cell yields a minimal set of symmetry inequivalent sites \mathcal{R} , along with the transformations that map every site j in the unit cell to some $i^* \in \mathcal{R}$.

Step 2: Determine all point-group symmetries. For each reference site $i^* \in \mathcal{R}$, we determine point-group symmetries that leave i^* invariant using essentially the same algorithm as in step 1. The only differences that we now iterate over pairs i_1, i_2 and i'_1, i'_2 that are both nearest neighbors of i^* at every step. This procedure generates a list of point-group symmetries $\{R_{\text{spin}} \circ R_{\text{space}}\}$ that leave i^* invariant.

Step 3: Obtaining the minimal set of symmetry inequivalent pairs $P_{i^*}^*$ To identify the minimal set of symmetry-inequivalent pairs $\mathcal{P}_{i^*}^*$ for fixed $i^* \in \mathcal{R}$, we proceed iteratively:

1. Start from an arbitrary pair $(i^* j_1) \in \mathcal{P}$ and add it to $\mathcal{P}_{i^*}^*$.
2. Apply all point-group symmetries from step 2 to this pair, and remove the resulting images (i^*, j'_1) from \mathcal{P} for the subsequent steps. Record, however, the corresponding point-group symmetry and the site j_1 to which they are equivalent.

This iteration continues until \mathcal{P} is exhausted, at which point $\mathcal{P}_{i^*}^*$ contains exactly the set of symmetry-inequivalent pairs for the reference site i^* , along with the mappings to all other pairs $(i^* j) \notin \mathcal{P}_{i^*}^*$.

Implementation for the pf-FRG flow equations. Finally, let us discuss how to make practical use of the symmetry reduction described above in the pf-FRG flow equation solver.

As a first step, we enumerate all symmetry-irreducible pairs $(i^* j^*) \in \mathcal{P}^*$ by integers α , and use this index to label the vertex components we actually need to compute as $\Gamma_{i^* j^*}^{\mu\nu} \equiv \Gamma_{\alpha}^{\mu\nu}$.

When initializing a spin model, the code automatically constructs a mapping array M with entries $M[(ij)] = (\alpha, p, \zeta)$ for each pair $(ij) \in \mathcal{P}$. Each entry contains (i) the integer α of the irreducible pair $(i^* j^*) \in \mathcal{P}^*$ to which (ij) is mapped, and (ii) the corresponding spin transformation $\mathbf{R}_{\text{spin}} = (p, \zeta)$, specified by the permutation p and sign function ζ . Whenever the solver needs to evaluate a vertex $\Gamma_{ij}^{\mu\nu}$ for $(ij) \notin \mathcal{P}^*$, the mapping is retrieved via $\alpha, p, \zeta = M[(ij)]$, and the vertex is reconstructed as

$$\Gamma_{ij}^{\mu\nu} = \zeta(\mu)\zeta(\nu) \Gamma_{\alpha}^{p(\mu)p(\nu)}. \quad (4.131)$$

For the s - and u -channel flow equations (4.51, 4.53), only vertices with identical site indices appear on both sites of the flow equation, so the mapping array is not strictly needed. When

utilizing the frequency symmetries summarized in Table 4.4, however, they often contain site-exchange $(ij) \rightarrow (ji)$. Even if $(ij) \in P^*$, often $(ji) \notin P^*$ and we again need to apply a symmetry transformation. To handle this efficiently, we precompute an exchange array E whose entries $E[\alpha]$ (for α corresponding to $(ij) \in \mathcal{P}^*$) contain the index α' of the irreducible pair to which (ji) maps, together with the associated spin transformation.

The t -channel flow equation (4.52) is more complicated, as it contains a sum over the full lattice of the form

$$\sum_k \Gamma_{ik}^{\mu\nu} \Gamma_{kj}^{\kappa\eta} \quad (4.132)$$

We first restrict the sum to sites k within range of both i and j (i.e. again to pairs in \mathcal{P}). For all such k , the site pair of the left vertex is mapped as $(ik) \rightarrow (i^*k^*) \in \mathcal{P}^*$ with an accompanying spin rotation $R_{\text{spin}}^l = (p_l, \zeta_l)$, while the right vertex is mapped as $(kj) \rightarrow (k'^*j^*) \in \mathcal{P}^*$ with the corresponding rotation $R_{\text{spin}}^r = (p_r, \zeta_r)$. For many spin models, there are much more pairs in \mathcal{P} than there are in \mathcal{P}^* , and therefore the sum will contain many duplicate entries after the mapping. To utilize this, for each pair (ij) we store all *unique* combinations of $((i^*k^*), (k'^*j^*), R_{\text{spin}}^l, R_{\text{spin}}^r)$ and the number of times they appear in the sum N_{mult} in an array. The t -channel sum can then be efficiently computed as

$$\sum_k \Gamma_{ik}^{\mu\nu} \Gamma_{kj}^{\kappa\eta} = \sum_{\substack{\{(i^*k^*), (k'^*j^*) \\ R_{\text{spin}}^l, R_{\text{spin}}^r\}}} N_{\text{mult}} \zeta_l(\mu) \zeta_l(\nu) \zeta_r(\kappa) \zeta_r(\eta) \Gamma_{i^*k^*}^{p_l(\mu)p_l(\nu)} \Gamma_{k'^*j^*}^{p_r(\kappa)p_r(\eta)}, \quad (4.133)$$

where each summand implicitly carries its own multiplicity factor N_{mult} .

4.5 Calculating phase diagrams

In the previous sections we have discussed how the pf-FRG flow equations can be derived and solved numerically. We now turn to the question of how to extract physical observables and, ultimately, phase diagrams from the resulting flowing self-energy and vertex data. Importantly, the pf-FRG flow preserves all symmetries of the Hamiltonian. As a consequence, single-spin expectation values $\langle \mathbf{S}_i \rangle$ remain strictly zero throughout the flow, even if the true ground state is magnetically ordered and characterized by a finite $\langle \mathbf{S}_i \rangle$. The central observable of pf-FRG is therefore the flow of the two-spin correlation, or *susceptibility*, defined as

$$\chi_{ij}^{\Lambda ab}(\omega) = \int_0^\infty d\tau e^{i\omega\tau} \left\langle T_\tau S_i^a(\tau) S_j^b(0) \right\rangle. \quad (4.134)$$

For spin-valley models, one can similarly define spin-valley spin-valley correlations. An expression of those susceptibilities in terms of the of the self-energy and vertex is given in Appendix A. Since the frequencies ω are Matsubara frequencies along the imaginary axis, the two physically accessible quantities are the static correlations $\chi_{ij}^{\Lambda ab}(\omega = 0)$ and the equal-time correlations

$$\langle T_\tau S_i^a(\tau = 0) S_j^b(\tau = 0) \rangle = \int d\omega \chi_{ij}^{ab}(\omega). \quad (4.135)$$

The additional integration over continuous frequencies has to be performed numerically, which introduces additional errors and thus most pf-FRG studies focus on static correlations when analyzing physical properties of a system. Exceptions include Refs. [130, 138], which showed that violations of the single-occupation constraint discussed in Sec. 4.2.8 can be quantified through equal-time correlators.

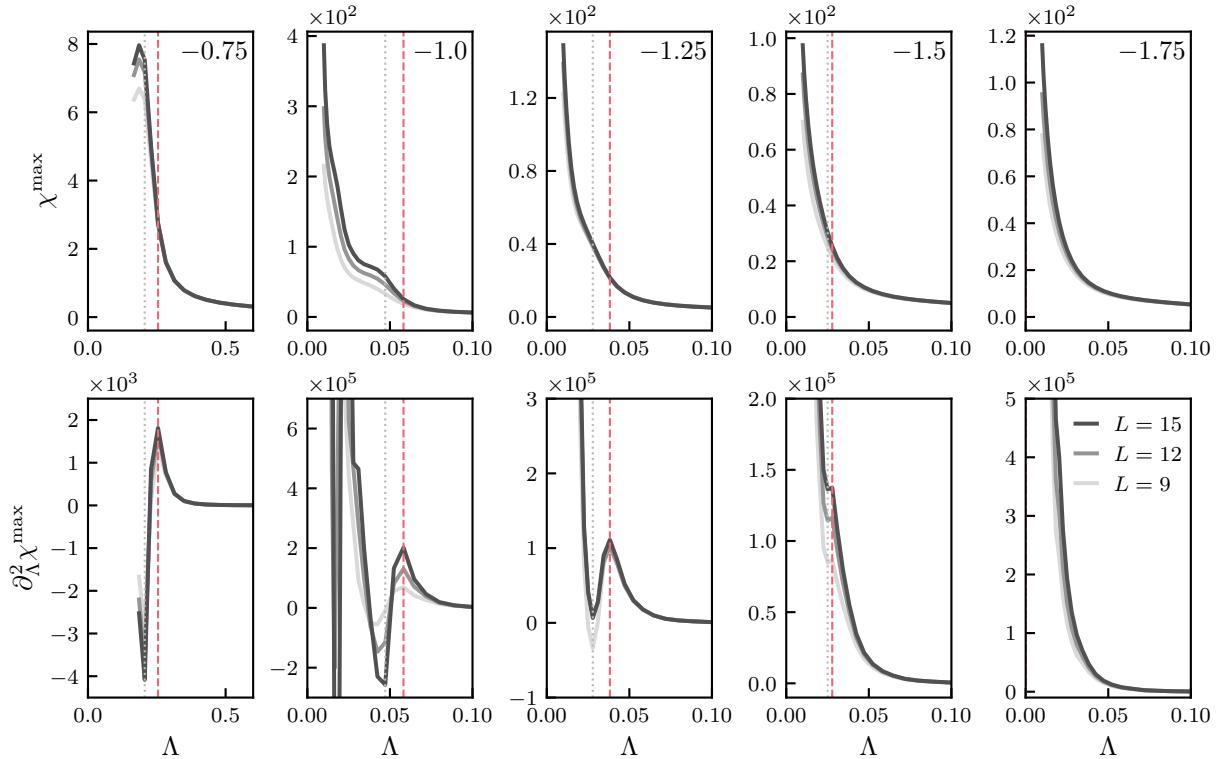


Figure 4.2 – RG flows of the structure factor across a Néel–PM phase transition. Each column shows the flow of the structure factor at its maximum wave vector, $\chi^{\max} = \chi^{zz}(\mathbf{k}^{\max})$ (top), and the corresponding second derivative (bottom). Data are for the nearest-neighbor model on the maple-leaf lattice (discussed in detail in Chapter 6) at fixed $J_t/J_h = -1$, with varying J_d/J_h indicated in the top-right corner of each panel. With increasing J_d/J_h , the system evolves from a Néel phase (left) to a paramagnetic phase (right). The dashed red line marks the critical scale Λ_c , where a flow breakdown is detected, signaling the onset of dipolar magnetic order. This is accompanied by a characteristic nonmonotonic feature in the second derivative. The dotted gray line indicates the scale at which the second derivative begins to increase again (Λ_2 used in our flow breakdown criterion described in the main text). In the rightmost column, no flow breakdown occurs, consistent with a paramagnetic ground state.

In the following, we describe how phase diagrams can be constructed from the pf-FRG flow of spin–spin correlations. The first step is to distinguish between magnetically ordered and nonmagnetic states. Here, we restrict *magnetically ordered* phases to those with dipolar order, i.e. characterized by a finite order parameter $m = \sum_{i,a} \lambda_{ia} S_i^a$ that is linear in the spin operators. Such states exhibit finite local magnetizations $\langle \mathbf{S}_i \rangle$, thereby breaking time-reversal symmetry and usually also lattice symmetries of the Hamiltonian. In the context of quantum spin models, we refer to all other states as *paramagnetic* (PM). This class includes symmetric spin liquids and quantum disordered states, but also phases that break lattice symmetries (e.g. dimer or nematic order) or spin symmetries (e.g. spin-nematic order). We will first explain how pf-FRG distinguishes between PM and magnetically ordered states, then discuss how to further characterize the states in each case, and finally outline how pf-FRG results can be directly compared with neutron-scattering experiments.

4.5.1 Distinguishing magnetically ordered and paramagnetic states

As a first step, we use the pf-FRG to detect whether the Hamiltonian has a magnetically ordered or paramagnetic ground state. As already explained, the preservation of symmetry obstructs the pf-FRG flow to truly entering the symmetry broken phase—the corresponding order parameters will always stay zero. Instead, it can be shown that the flow will exhibit an instability at a finite critical scale Λ_c that, in theory, manifests in a divergence of the susceptibility associated to the order parameter [68].

Most conveniently, this susceptibility is defined as a component of the static *structure factor*

$$\chi^{\Lambda ab}(\mathbf{k}) = \frac{1}{N} \sum_{ij} e^{i\mathbf{k} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \chi_{ij}^{\Lambda ab}(\omega = 0), \quad (4.136)$$

which is the Fourier transform of the spin-spin correlations. For Hamiltonians with a magnetically ordered ground state, the FRG flow of $\chi^{\Lambda ab}(\mathbf{k})$ exhibits a divergence—referred to as *flow breakdown*—at a finite scale Λ_c and at momenta \mathbf{k}_{\max} characterizing the ordering pattern. Conversely, the absence of a flow breakdown signals the *absence* of any dipolar order, i.e. a paramagnetic ground state.

In practice, the approximations we apply to the flow equations to numerically solve them will soften the divergence at the flow breakdown, often to much more subtle features. Most notably, any divergence will eventually plateau to a peak due to the finite correlation length imposed by the lattice truncation L . Moreover, truncating the hierarchy of flow equations may suppress divergences altogether, even at large L . As an illustration, Fig. 4.2 shows the typical behavior of the structure factor flow when traversing from an ordered phase to a paramagnetic phase (data shown are for the nearest-neighbor model on the maple-leaf lattice, discussed in detail in Chapter 6 below). Deep in the ordered phase (leftmost panel), the flow develops a sharp, L -dependent peak—a clear flow breakdown implying long-range order. Deep in the paramagnetic regime (rightmost panel), the flow remains convex, featureless, and only weakly L -dependent. In the intermediate region near the phase boundary (middle panels), however, the behavior is less clear-cut.

For these cases, there is no unambiguous way to distinguish a true flow breakdown from other correlation-induced effects. It is therefore not surprising that a variety of flow-breakdown criteria have been proposed ranging from the rather subjective approach of “carefully inspecting the flow by eye” (used in many early pf-FRG studies) to more systematic techniques such as scaling analyses of onsite correlations [120], or thresholds on the “nonconvexity” of the flow based on its second derivative [146]. None of these approaches can be *rigorously* justified, since the precise impact of the truncation of the flow equations and other numerical approximations on the exact flows is extremely difficult to quantify. The encouraging observation, however, is that the choice of criterion rarely changes the qualitative outcome: the same phases are typically identified, with only the location of the phase boundaries shifting (which may, however, not be true for very small paramagnetic phases). Since pf-FRG phase diagrams generally have a tendency to overestimate the extent of paramagnetic regimes, we have, building on previous work and our own experience, devised a flow-breakdown criterion that flags even subtle nonmonotonic features in the flow as instabilities. This criterion is described in the following.

A key observation motivating our criterion is that, even in clearly ordered phases, peaks disappear when the lattice truncation L is too small. Nevertheless, the flow still exhibits non-convex, L -dependent features. These are most clearly seen in the second derivative of the flow, $\partial_\Lambda^2 \chi_{\mathbf{k}_{\max}}^\Lambda$, shown in the lower panel of Fig. 4.2: with decreasing cutoff, the curve first dips down and then rises again, signaling nonmonotonic behavior. Importantly, this feature becomes more

pronounced with increasing L , suggesting that it may still grow into a divergence in the thermodynamic limit. By contrast, in regimes we identify as paramagnetic, the second derivative shows no such nonmonotonic behavior and only a weak L -dependence. These observations form the basis of our reproducible flow-breakdown criterion. Our procedure is as follows:

1. Determine the momentum \mathbf{k}^{\max} where the structure factor is maximal at low Λ , and extract the corresponding flow $\chi^{\max}(\Lambda) = \chi^{\Lambda z z}(\mathbf{k}^{\max})$.
2. Scan $\chi^{\max}(\Lambda)$ from large to small Λ and identify the first nonmonotonicity in its second derivative $\partial_{\Lambda}^2 \chi^{\max}$. Mark this point as Λ_1 (red dashed lines in Fig. 4.2).
3. Locate $\Lambda_2 < \Lambda_1$ where $\partial_{\Lambda}^2 \chi^{\max}$ turns upward again (gray dotted lines in Fig. 4.2).
4. Evaluate $\chi^{\max}(\Lambda_2)$ for increasing lattice truncations L (here $L = 9, 12, 15$). If the value increases by more than a relative threshold ϵ between successive L , we identify a genuine flow breakdown at $\Lambda_c \equiv \Lambda_1$.
5. If not, repeat the procedure for any subsequent nonmonotonicities at $\Lambda < \Lambda_2$. If none satisfy the criterion, set $\Lambda_c = 0$ and classify the state as paramagnetic.

Requiring an increase of at least ϵ (we typically choose $\epsilon = 3\%$) in $\chi^{\max}(\Lambda_2)$ with growing L helps us exclude nonmonotonic features in $\partial_{\Lambda}^2 \chi^{\max}$ that do not scale with L and therefore likely do not correspond to true divergences in the thermodynamic limit. Instead, such features likely originate from short-range correlation effects, which are absent in a fully uncorrelated paramagnetic phase where the susceptibility follows $\chi \sim 1/\Lambda$. Slight variations in ϵ do not affect the qualitative phase structure, though they may shift the precise location of phase boundaries.

We note that, beyond the absence of a flow breakdown, the momentum-resolved structure factor can provide additional qualitative guidance: paramagnetic states typically exhibit broad features, whereas ordered phases display sharp, well-defined peaks. This contrast can aid in distinguishing the two. We note however, that relatively sharp peaks in the structure factor do not generally contradict a spin liquid ground state. Indeed, several spin liquids emerge as molten versions of classical parent states and retain maxima at the same momenta as the corresponding Bragg peaks. For example, the quantum spin liquid in the triangular lattice J_1 - J_2 Heisenberg antiferromagnet shows soft but clearly pronounced maxima at the \mathbf{K} -points of the Brillouin zone, precisely where peaks of 120° order would be expected [147].

4.5.2 Characterizing ordered states

Let us now, assume we have identified a flow breakdown and thus predict a magnetically ordered ground state. We can then quantify the order by determining the position of all peaks \mathbf{k}^* in the momentum resolved structure factor at cutoffs just above the critical scale Λ_c . To infer the corresponding real-space ordering pattern, the simplest approach is to compare the peak positions with structure factors of ordered phases in the corresponding classical spin model. In the following we therefore shortly examine the properties of the structure factor for classically ordered state.

To this end, consider a classical spin-spiral state with wave vector \mathbf{q}^* (introduced in Sec. 3.1) of the form

$$\mathbf{S}_i^{\mathbf{q}^*} = \mathbf{S}_{m\alpha}^{\mathbf{q}^*} = \begin{pmatrix} \cos(\mathbf{q}^* \cdot \mathbf{R}_m + \phi_{\alpha}) \\ \sin(\mathbf{q}^* \cdot \mathbf{R}_m + \phi_{\alpha}) \\ 0 \end{pmatrix}, \quad (4.137)$$

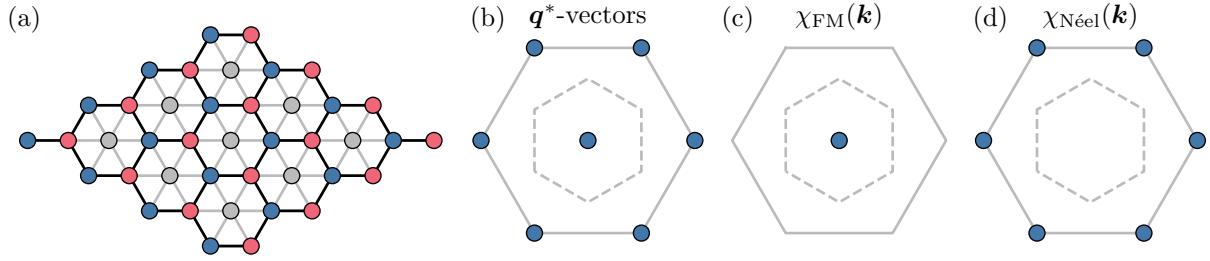


Figure 4.3 – Relation between q^* -vectors and structure factor peaks. (a) The honeycomb lattice with two sublattices (blue and red); including the gray sites reveals an underlying triangular lattice. (b) The q^* -vectors for the FM and Néel states, which are both $\mathbf{q} = 0$ states. The dashed line indicates the honeycomb Brillouin zone, while the solid line shows the *extended* Brillouin zone, which is the first Brillouin zone of the underlying triangular lattice. The corners of the extended Brillouin zone ($\mathbf{K}_{\text{ex}}, \mathbf{K}'_{\text{ex}}$) are related to $\mathbf{q} = 0$ by a reciprocal lattice vector and thus also represent valid q^* -vectors. (c,d) The structure factors of the FM and Néel states, which are periodic only in the extended Brillouin zone and show peaks at different $\mathbf{k} = \mathbf{q}^*$ vectors depending on the sublattice structure.

where the site position is written as $\mathbf{r}_i = \mathbf{R}_m + \mathbf{b}_\alpha$, with \mathbf{R}_m being the position of the unit cell and \mathbf{b}_α the position of the basis site within the unit cell. For such a state, the (summed) structure factor evaluates to

$$\begin{aligned} \sum_a \chi^{aa}(\mathbf{k}) &= \frac{1}{N} \sum_{ij} e^{i\mathbf{k}(\mathbf{r}_i - \mathbf{r}_j)} \mathbf{S}_i \cdot \mathbf{S}_j \\ &= \frac{N_{\text{UC}}}{2} \sum_{\mathbf{G}} \left(\delta_{\mathbf{k}, -\mathbf{q}^* + \mathbf{G}} \sum_{\alpha\beta} e^{i\mathbf{k}(\mathbf{b}_\alpha - \mathbf{b}_\beta) + i(\phi_\alpha - \phi_\beta)} + \delta_{\mathbf{k}, +\mathbf{q}^* + \mathbf{G}} \sum_{\alpha\beta} e^{-i\mathbf{k}(\mathbf{b}_\alpha - \mathbf{b}_\beta) - i(\phi_\alpha - \phi_\beta)} \right), \end{aligned} \quad (4.138)$$

where N_{UC} is the number of unit cells, and we used $\cos(x)\cos(y) + \sin(x)\sin(y) = \cos(x - y)$ and $\delta_{\mathbf{k}, \mathbf{k}'} = \frac{1}{N} \sum_m \exp(i(\mathbf{k} - \mathbf{k}') \mathbf{R}_m)$. This shows that the structure factor will always show peaks at wave-vectors \mathbf{k}^* that are related to \mathbf{q}^* by a shift of reciprocal lattice vector \mathbf{G} . More complicated multi- \mathbf{q} -states (see, e.g. the noncoplanar phase on the maple-leaf lattice illustrated in Table 3.1) will similarly show peaks for each \mathbf{q}^* . The intensity of these peaks (which may also vanish) is set by the sublattice structure—with geometric contributions from \mathbf{b}_α and magnetic contributions from ϕ_α .

Unlike the spiral wave vectors \mathbf{q}^* , which are always periodic within the first Brillouin zone, the periodicity of the structure factor also depends on the position of the basis sites. For nonBravais lattices that can be mapped onto an underlying Bravais lattice by adding sites at certain positions, the structure factor will always be periodic with the reciprocal lattice vector of this underlying lattice. For example, the honeycomb lattice turns into a triangular lattice by adding sites at the center of the hexagons, as illustrated in Fig. 4.3. The structure factor will then be periodic with the first Brillouin zone of the triangular lattice, which in this context we refer to as the *extended Brillouin zone*. When analyzing the pf-FRG flow, we thus typically search for divergence of the structure factor inside this extended zone.

From the peak positions in the structure factor, we can therefore not only extract the spiral wave vectors \mathbf{q}^* but also infer the sublattice structure. By comparison with classical results (e.g., Luttinger–Tisza), this typically allows us to identify the ordered ground state directly from the pf-FRG structure factor. This is illustrated in Fig. 4.3 for the FM and AFM state on the honeycomb lattice. Both are $\mathbf{q}^* = 0$ states, as the order is repeated on each unit cell, but the

FM state shows structure factor peaks at $\mathbf{k} = \Gamma$ and the AFM shows a peak at $\mathbf{k} = \mathbf{K}_{\text{ex}}, \mathbf{K}'_{\text{ex}}$, which are the corners of the extended Brillouin zone.

We emphasize that, because pf-FRG effectively simulates an infinite lattice, this procedure applies equally well to incommensurate phases—an important advantage over methods restricted to finite lattices, where boundary effects are often much stronger.

4.5.3 Characterizing paramagnetic states

If our flow-breakdown criterion detects no instability, indicating a paramagnetic ground state, we can still attempt to characterize the nature of this phase. As a first step, we can again calculate the momentum-resolved structure factor and compare with other known results. For instance, we can compare with known structure factors from classical spin liquids, which often exhibit so-called *pinch-point* singularities—signatures of algebraically decaying correlations and an emergent gauge structure—as we shortly hinted in the introduction and will discuss in more detail in Chapter 5³. Another option is to compare with structure factors obtained from the projective symmetry group (PSG) classification, which systematically enumerates all symmetry-allowed spin-liquid phases within fermionic parton constructions [7, 151]. Agreement with these structure factors indicates that the pf-FRG results are at least consistent with a corresponding QSL ground state.

We can additionally probe for lattice symmetry breaking states, such as dimer or nematic phases. These phases lack a finite local magnetization $\langle \mathbf{S}_i \rangle$ but break lattice symmetries and are characterized by order parameters quadratic in the spin operators. The corresponding susceptibility is quartic in spins and of the general form [148] $D_{ij,kl} = \langle (\mathbf{S}_i \cdot \mathbf{S}_j)(\mathbf{S}_k \cdot \mathbf{S}_l) \rangle - \langle \mathbf{S}_i \cdot \mathbf{S}_j \rangle \langle \mathbf{S}_k \cdot \mathbf{S}_l \rangle$. In the pseudo-fermion representation, this requires the calculation of four-particle correlation functions, which is not possible within the current pf-FRG framework. We can, however, probe the systems tendency towards a corresponding symmetry breaking pattern.

For dimer states, symmetry breaking typically separates bonds into those with strong correlations (S) and those with weaker correlations (W). If we change the original couplings in the Hamiltonian by adding a corresponding symmetry breaking perturbation $J_{ij} \pm \epsilon$ (with \pm on S/W bonds), we can calculate the dimer response of the system as [14]

$$\chi_{\text{dimer}}^{\Lambda} = \frac{J}{\epsilon} \frac{\chi_S^{\Lambda} - \chi_W^{\Lambda}}{\chi_S^{\Lambda} + \chi_W^{\Lambda}}, \quad (4.139)$$

where $\chi_{S/W}^{\Lambda}$ are the spin-spin correlator on strong/weak bonds. The dimer response is normalized to $\chi_{\text{dimer}}^{\Lambda \rightarrow \infty} = 1$ at high cutoffs, and will grow larger with decreasing cutoff if the systems “accepts” the symmetry breaking, or will get weaker if the it is “rejected”. Similarly, we can probe for spin-nematic order, which instead breaks the SU(2) of a Heisenberg model (or similarly a U(1) symmetry of less symmetric models) by adding an SU(2) symmetry breaking term to Heisenberg interactions of the form

$$J \mathbf{S}_i \cdot \mathbf{S}_j \rightarrow (J - \epsilon)(S_i^x S_j^x + S_i^y S_j^y) + (J + \epsilon)(S_i^z S_j^z), \quad (4.140)$$

³We note, however, that pinch-point singularities alone, do not guarantee a spin liquid ground state. The Heisenberg antiferromagnet on the pyrochlore lattice, for example, shows pinch points in its pf-FRG structure factor [14, 121, 148]. While this model was long believed to realize a quantum spin liquid ground state [149, 150], more recent studies indicate symmetry-breaking tendencies consistent with a dimer phase [36, 37].

and calculating the spin-nematic response as

$$\chi_{\text{SN}}^{\Lambda} = \frac{J \chi_{ij}^{xx} - \chi_{ij}^{zz}}{\epsilon \chi_{ij}^{xx} + \chi_{ij}^{zz}}. \quad (4.141)$$

A large dimer or nematic response indicates that the system has a strong tendency towards the corresponding symmetry breaking. However, it does not guarantee a corresponding ground-state—it is not clear that the pf-FRG in the Katanin truncation can capture such phases at all. For any *quantitative* analysis of paramagnetic phases, complementary methods are therefore required.

4.5.4 Comparison with neutron-scattering experiments

The structure factor is of particular physical relevance, as it can be directly measured in neutron-scattering experiments. Using unpolarized neutrons, inelastic neutron scattering probes the dynamical structure factor [152, 153]

$$S_{\perp}(\mathbf{k}, \omega) = \sum_{a,b} \left(\delta_{ab} - \frac{k^a k^b}{\mathbf{k}^2} \right) \chi^{ab}(\mathbf{k}, \omega), \quad (4.142)$$

where ω is now a *real* frequency. The prefactor reflects that neutrons only couple to spin components transverse to the momentum transfer \mathbf{k} . Measurements at low energy and temperature can then be qualitatively compared to the static structure factors obtained within pf-FRG, and such comparisons have been successfully carried out in several works [154–156].

Even more information can be obtained by experiment with polarized neutrons, where the structure factors are defined in terms of the incident neutrons' polarization \mathbf{n} [152, 153], effectively separating the unpolarized neutron structure factor into two channels, namely the non-spin-flip (NSF) channel

$$S_{\perp}^{\text{NSF}}(\mathbf{k}, \omega) = \sum_{a,b} n^a \chi^{ab}(\mathbf{k}, \omega) n^b \quad (4.143)$$

and the spin-flip channel

$$S_{\perp}^{\text{SF}}(\mathbf{k}, \omega) = \sum_{a,b} (\mathbf{n} \times \mathbf{k})^a \chi^{ab}(\mathbf{k}, \omega) (\mathbf{n} \times \mathbf{k})^b. \quad (4.144)$$

For simplicity, we have assumed here that the local magnetic moments $\mathbf{m}_i = \mathbf{g}_i \mathbf{S}_i$ are related to the spin operators by an isotropic g -tensor ($\mathbf{g}_i = \mathbb{1}$). For spin-orbit coupled ions in an anisotropic crystal field, as discussed in Sec. 2.1.3, the correct expression is obtained by replacing $\mathbf{S}_i \rightarrow \mathbf{m}_i$ in the formulas above.

4.6 Discussion and outlook

The main strength of the pf-FRG is definitely its broad applicability—as we have shown in can be applied to highly frustrated models in both two and even three dimensions—a regime where most other numerical methods struggle. At the same time, the current implementation also comes with important limitations. In the following, we discuss these limitations and review different approaches that are currently being developed, or have already been proposed, to overcome them.

Magnetic fields Our current pf-FRG implementation cannot include magnetic fields, since their presence breaks time-reversal symmetry, which is heavily exploited in deriving the symmetry relations for the self-energy and vertex (see Table 4.2). It was long assumed that this would render pf-FRG numerically infeasible. A recent study [157], however, demonstrated that by exploiting the remaining symmetries and the $SU(2)$ gauge redundancy of the parton construction, breaking only time-reversal symmetry increases the cost mainly through two effects: (i) multiple self-energy components Σ^μ must be computed (two for $U(1)$ -symmetric systems, four without spin symmetry), and (ii) the vertex must be evaluated at twice the number of frequencies. Overall, this leads to at most an eightfold increase in cost. Magnetic fields can then be used in two ways. First, infinitesimal seed fields can regularize the flow and allow access to ordered phases, enabling explicit calculation of order parameters such as the magnetization—although Ref. [157] showed that pf-FRG strongly overestimated zero-field magnetizations. A key drawback, however, is that any symmetry-breaking pattern must be specified *a priori* to initialize the flow with the appropriate seed fields. For complex orders with large unit cells and many broken lattice symmetries, this requirement is the main driver of the increase in computational cost and can quickly render calculations prohibitive. Second, finite external fields can be studied directly, where results compare reasonably well with QMC, enabling, for instance, the study of magnetization plateaus.

Real frequencies Our pf-FRG implementation works on the Matsubara axis, giving access only to imaginary-frequency structure factors, while experiments typically probe the real-frequency dynamical structure factor. Access to real-frequency data would also be highly valuable for characterizing paramagnetic states—for example, to determine whether they are gapped or gapless.

The most straightforward approach to obtaining real-frequency data from our existing pf-FRG implementation would be to perform analytical continuation. This is, however, a numerically ill-posed problem and so far unsuccessful for pf-FRG. A recent study [158] instead implemented pf-FRG directly on the real-frequency axis within the Keldysh formalism. Although they showed that this is numerically feasible, the resulting spectra failed to reproduce any sharp magnon branches (e.g. in the 1D Heisenberg chain) and any gapped excitations, likely because the pf-FRG flow remained confined to the symmetric state. Including magnetic fields, which allow the flow to enter the symmetry broken regime, could therefore be a promising direction. Another important improvement would be a more rigorous enforcement of the single-occupation constraint, for example via a Popov–Fedotov potential (see Sec. 4.2.8), which Ref. [158] omits although working at finite temperature. So while the Keldysh approach has not yet been fully successful, it remains a promising direction for extending pf-FRG towards real-frequency observables.

Enforcing the single-occupation constraint and finite temperatures As discussed in Sec. 4.2.8, our pf-FRG implementation enforces the single-occupation constraint only on average, which restricts calculations to zero temperature—and even there, its violation may still affect results [138]. One remedy is the Popov–Fedotov potential, which, however, has been shown to only work for larger temperatures $T \gtrsim 0.3J$ [138]. A more powerful approach is to adopt a different parton representation: expressing spins through real Majorana instead of complex Abrikosov fermions yields an *exact* mapping free of unphysical states, giving rise to the pseudo-Majorana FRG (pm-FRG) [14, 139–141]. Early pm-FRG studies [140, 141] demonstrated that this can even provide quantitatively accurate results for high temperatures ($T \sim J$), but somehow failed to detect phase transitions into ordered states that occur below $T \gtrsim 0.3J$. This was later improved with the implementation of the temperature-cutoff scheme, where T itself serves as the RG scale [139].

A very recent application of the pm-FRG to the XXZ model on the pyrochlore lattice successfully detected phase transitions down to $T \sim 0.01J$, comparable to the lowest cutoff accessible in our pf-FRG implementation [159]. In addition to the absence of unphysical state enabling finite-temperature calculations, pm-FRG offers a further advantage: phase transitions can be unambiguously identified via finite-size scaling of a susceptibility [140], yielding a robust and unambiguous criterion, unlike the more subjective identification of flow breakdowns in pf-FRG (see Sec. 4.5). Whenever both pm- and pf-FRG are available, pm-FRG with temperature flow is therefore expected to provide more accurate and reliable results—except perhaps at very low temperatures, where challenges may still persist for certain models. At the time of writing, however, pm-FRG had not yet been implemented for the general anisotropic spin models considered here. Extending it in this direction should, however, be conceptually straightforward and thus represents a promising avenue for future work.

It was recently shown that pseudo-Majorana representations can also be constructed for spin magnitudes $S > 1/2$ [160]. An interesting open question is whether analogous representations exist for $SU(4)$ spin models, potentially at different fermionic fillings corresponding to distinct representations of the Lie algebra. In particular, quarter filling—corresponding to the fundamental four-dimensional representation of $\mathfrak{su}(4)$ —remains inaccessible within the current pf-FRG framework. Developing such a representation would open the door to studying a much broader class of materials, including, for example, twisted bilayer graphene [48].

Truncation of the flow equations and higher-order correlation functions All truncation schemes employed within pf-FRG (or pm-FRG) so far exclude the explicit calculation of vertices beyond the 2-particle vertex. This is, very likely, the most drastic approximation in the pf-FRG, as this can not be rigorously justified for the pseudo-fermion Hamiltonian which effectively sits in the infinite coupling regime (as there is not kinetic term). This is also true for the more advanced multiloop truncation scheme [129, 130]. As outlined in Chapter 1, the argument of why the pf-FRG still provides an unbiased probe to magnetically ordered and disordered states, was the realization that it reduces to the exact mean-field treatment in the $S \rightarrow \infty$ limit, which favors classically ordered states [67], and in the $S \rightarrow \infty$ limit [68], which favors quantum paramagnetic states. At finite $S = 1/2$, $N = 2, 4$, one can thus hope that the interaction between the different mean-field channels is taken into account in an unbiased way.

As mentioned above, however, finite temperature pf-FRG calculations on small spin clusters in Ref. [138], show that results deviate from quantitatively exact results for low temperatures, where the parameter J/T can no longer act as an effective small control parameter. This indicates that the truncation of the flow equation may still strongly affect the results. To remedy this, one approach is to employ nonperturbative methods. Promising candidates are found in the field of embedding methods, where the key idea is to treat the system nonperturbatively only locally, and then glue the local constituents to an approximation of the global theory (similar to our CMFT approximations discussed in Sec. 3.3). A prominent example is dynamical mean-field theory [161], where the self-energy is self-consistently approximated by a local version obtained by an effective impurity model. In order to capture effects from nonlocal fluctuations, diagrammatic extensions of DMFT have been developed [162], where not the self-energy, but parts of the two-particle vertex are approximated by their local counterpart (an example is the dynamical vertex approximation [163]). Similar concepts could be combined with the pf-FRG, e.g. by including nonlocal vertex quantities in the initial condition of the flow (similar DMF²RG methods where results of the DMFT are used as initial conditions for the FRG flow [164]).

Another promising route is to bypass the parton decomposition entirely and work directly with

spin operators. Due to their nontrivial commutation relations, an unconstrained path-integral representation of the partition function is not available, which prohibits the conventional derivation of flow equations. However, an alternative formulation based on a generating functional at the *operator level* is possible, leading to the development of the *spin-FRG* by the group of Peter Kopietz [133]. Common truncations of FRG flow equations excluding vertices with more than four legs (such as the Katanin truncation), would—for spin-FRG—already include four-spin correlation functions, in contrast to only two-spin correlators available in pf- and pm-FRG. This likely not only improves on the approximation due to the truncation, but also enables the direct calculation of dimer and nematic susceptibilities (see Sec. 4.5.3). Current implementations already show promising results in terms of accuracy [165–168]. However, numerically more advanced implementations comparable to state-of-the-art pf- and pm-FRG are currently still missing and thus a fair comparison remains to be made.

Part II

Applications

Chapter 5

The role of quantum fluctuations in pyrochlore rare-earth oxides

When searching for experimental realizations of frustrated quantum magnetism, pyrochlore rare-earth oxides have proven especially fruitful [24]. These compounds have the chemical formula $R_2M_2O_7$, where R is a magnetic rare-earth ion, M a nonmagnetic transition metal, and O is oxygen. The rare-earth ions form a pyrochlore lattice—a network of corner-sharing tetrahedra illustrated in Fig. 5.1(a). This geometry is inherently frustrated already at the nearest-neighbor level. Rare-earth ions have partially filled $4f$ shells, which are strongly localized orbitals, giving rise to well-defined local magnetic moments. Compared to typical $3d$ transition metals, rare-earth ions are much heavier and therefore subject to strong spin-orbit coupling. As discussed in Chapter 2.2.4, this may result in (i) highly anisotropic g -tensors that constrain magnetic moments to specific directions or planes, and (ii) strongly anisotropic, bond-dependent interactions between them. The combination of geometric frustration and these anisotropies can, even at the classical level, produce an extensively degenerate ground-state manifold that is *accidental*, i.e., not protected by symmetry. Such a degeneracy is the basis of many interesting phenomena associated with frustrated magnets, and thus many rare-earth oxides realize exotic magnetic behavior.

One possibility is that the degeneracy is simply lifted by fluctuations. Thermal fluctuations may lift the degeneracy by favoring states with larger entropy ('thermal order-by-disorder' [16]), while quantum fluctuations can select states via corrections to the classical ground-state energy ('quantum order-by-disorder' [24]). In real magnets, both effects will typically act together (and can either select the same or different states). One of the few materials where order-by-disorder is relatively well established is the pyrochlore rare-earth oxide $Er_2Ti_2O_7$, where a sixfold degenerate ground state is stabilized from an extensive classical manifold, likely through a combination of thermal and quantum fluctuations [24].

The arguably even more interesting situations arises when fluctuations fail to lift the degeneracy. In this case, the system shows *no* long-range magnetic order and retains a finite residual entropy at zero temperature—the defining properties of a *classical spin liquid*. The most prominent realizations of classical spin-liquid phases in real materials are also found among the pyrochlore rare-earth oxides: the *dipolar spin-ice* compounds $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_6$ [20]. Strong spin-orbit coupling and crystal fields constrain their large $\sim 10\mu_B$ moments to point almost perfectly along the local $\langle 111 \rangle$ axes, defined as the axis connecting the center of the tetrahedra with its vertex [see Fig. 5.1(a)]. Dominant dipolar interactions lead to an effective AFM Ising model on the pyrochlore lattice—a model for which Anderson already in 1956 predicted the extensive spin-ice ground-state manifold [169]. As discussed in detail in Sec. 1.1, this manifold is defined

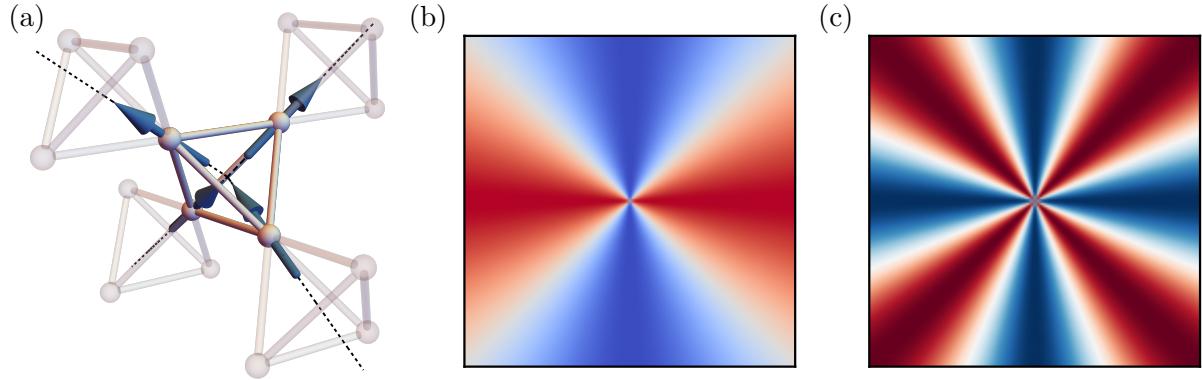


Figure 5.1 – Pinch points on the pyrochlore lattice. (a) Pyrochlore lattice with a classical spin configuration obeying the “two-in, two-out” spin-ice constraint. The dotted lines indicate the local \hat{z} -axes along which the spins are constrained in spin-ice compounds. (b) Twofold pinch point in momentum space, originating from dipolar correlations in spin ice and characteristic of an emergent gauge theory governed by Gauss’ law. (c) Example of a fourfold pinch point, characteristic of states described by higher-rank emergent gauge theories that obey generalized Gauss’ laws.

by the local constraint

$$\left| \sum_{i \in t} S_i^{\hat{z}} \right|^2 = 0, \quad (5.1)$$

where \hat{z} denotes the local $\langle 111 \rangle$ axis and the sum runs over the four spins of a tetrahedron. The phase becomes particularly interesting when viewed in its effective low-energy continuum description. For a continuous field \mathbf{B} microscopically representing the local magnetization, the constraint takes the form of Gauss’ law

$$\sum_a \partial_a B^a \equiv \nabla \cdot \mathbf{B} = 0, \quad (5.2)$$

which can be resolved by introducing a gauge field $\mathbf{B} = \nabla \times \mathbf{A}$ —making spin ice a textbook example of an *emergent gauge theory*, remarkably realized in $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Ho}_2\text{Ti}_2\text{O}_6$ as confirmed by experiment [20].

Spin ice is just one example of a classical spin liquid state. In fact, according to a recent classification scheme [27, 28], it belongs to a broader family of *algebraic* spin liquids—gapless classical liquid phases with algebraically decaying correlations. Such phases are further distinguished by the conservation laws that define them: like spin ice, they can be described by local constraints that lead generalizations of Gauss’ law and give rise to emergent $\text{U}(1)$ gauge structures. Of particular interest are cases where the conserved quantity is not a vector field but a rank-2 tensor (i.e. a matrix), leading to constraints of the form

$$\sum_a \partial_a B^{ab} = 0 \quad \text{or} \quad \sum_{ab} \partial_a \partial_b B^{ab} = 0, \quad (5.3)$$

and to emergent *higher-rank* gauge theories. These theories are associated with *fracton* phases—a recently discovered novel phase of matter where excitations are severely restricted in their mobility, being either completely immobile or confined to move only along certain directions [170].

The fingerprints of such gauge theories appear directly in spin–spin correlations. For instance,

in conventional U(1) gauge theory the effective field shows dipolar correlations of the form

$$\langle S_i^a S_j^b \rangle \sim \langle B^a(\mathbf{r}_i) B^b(\mathbf{r}_j) \rangle \sim \frac{1}{r_{ij}^3} \left(\delta_{ab} - 3 \frac{r_{ij}^a r_{ij}^b}{r_{ij}^2} \right), \quad (5.4)$$

which in momentum space gives¹

$$\langle B^a(\mathbf{q}) B^b(-\mathbf{q}) \rangle \sim \left(\delta_{ab} - \frac{q^a q^b}{q^2} \right). \quad (5.5)$$

The discontinuity at $\mathbf{q} = 0$ produces the well-known twofold pinch point [illustrated in Fig. 5.1(b)]. By contrast, higher-rank gauge theories yield *multifold* pinch points. For example, fourfold pinch-point are generated by terms such as as

$$\langle B^{ab}(\mathbf{q}) B^{cd}(-\mathbf{q}) \rangle \sim \frac{q^a q^b q^c q^d}{q^4}, \quad (5.6)$$

which is illustrated in Fig. 5.1(c). Recent work has shown that a wide range of such higher-rank classical spin liquids naturally arise in the most general symmetry-allowed nearest-neighbor Hamiltonians on the pyrochlore lattice, making them potentially relevant for many pyrochlore rare-earth oxides [25, 26]. Since the corresponding higher-rank pinch-point singularities should also appear in neutron-scattering experiments, this opens the possibility of detecting fracton phases—and, more broadly, other exotic spin-liquid states—in real materials.

What remains challenging, however, is that although the physics of classical spins on the pyrochlore lattice—both ordered and liquid—is fairly well understood, much less is known about the role of quantum fluctuations. These are particularly important in rare-earth oxides with effective $S = 1/2$ moments. Among quantum spin-liquid phases, the best-studied is probably again quantum spin ice: adding small transverse exchange terms $\sim (\hat{S}_i^+ \hat{S}_j^- + \hat{S}_i^- \hat{S}_j^+)$ to the classical spin-ice Hamiltonian promotes the classical spin liquid into a U(1) quantum spin liquid. This phase hosts emergent photon-like excitations and provides a condensed-matter realization of physics closely analogous to quantum electrodynamics [171]. Much less is known about the fate of higher-rank spin liquids under quantum fluctuations. A recent study, using the pf- and pm-FRG, suggests that effects are much more drastic compared to spin ice [172]. While in spin-ice, pinch-point singularities are merely washed out akin to a thermal broadening [173], in phases described by higher-rank gauge theories the quantum structure factor showed a signal significantly modified from the multifold pinch-points of the classical theory, indicating a fragility of such phases under quantum fluctuations.

In addition to the more theoretical concept of higher-rank spin liquids, the role of quantum fluctuations in real pyrochlore materials is still far from settled. A good example is $\text{Yb}_2\text{Ti}_2\text{O}_7$, long considered a promising quantum spin-ice candidate [174, 175]. Early experiments reported broad continua in the dynamical structure factor, suggestive of fractionalization [176–178], and the nature of its ground state was contested. More recent studies on cleaner samples with improved resolution, however, established a ferromagnetic ground state with well-defined magnon branches [179, 180]. Intriguingly, the shape of the spectrum is best captured by a combination of spin-wave dispersions from both ferro- and antiferromagnetic states [179, 180]. The estimated parameters indeed lie close to the classical phase boundary between two such ordered phases.

¹The Fourier transform can be best carried out by using $(\delta_{ab} - 3r_{ij}^a r_{ij}^b / r_{ij}^2) / r_{ij}^3 = (\delta_{ab} \nabla^2 - \partial_a \partial_b) G(\mathbf{r})$ with $G(\mathbf{r}) = 1/r$ and using the fact that derivatives turn into simple factor of $\partial_a \rightarrow q^a$ under the Fourier transform.

In a purely classical description, however, it is not clear how these mixed correlations would be realized, opening the possibility that quantum fluctuations may lead to tunneling between the two competing orders. Current theoretical approaches, largely based on classical Monte Carlo and linear spin-wave theory, are not able to resolve this issue. Beyond $\text{Yb}_2\text{Ti}_2\text{O}_7$, several compounds exhibit no magnetic order down to the lowest temperatures and remain poorly understood, including, for example, the spin-liquid candidates $\text{Tb}_2\text{Ti}_2\text{O}_7$ and $\text{Pr}_2\text{Zr}_2\text{O}_7$ [24].

Motivated by both the theoretical interest in higher-rank gauge theories and the many open questions surrounding real materials, in this chapter we employ the pf-FRG—one of the few methods capable of treating highly frustrated three-dimensional magnets—to shed at least some light on the role of quantum fluctuations in these systems. Specifically, we investigate the phase diagram of the most general symmetry-allowed nearest-neighbor Hamiltonian on the pyrochlore lattice for Kramers' ions with effective $S = 1/2$ moments. We focus on a region of parameter space near a particular classical higher-rank spin liquid known as the *pinch-line* spin liquid, which occurs at a *classical triple point* where three ordered phases meet [181]. Of additional interest is a spin-nematic state predicted to lie along one of the phase boundaries terminating at this triple point. The parameter region is further motivated by its vicinity to the parameter estimates of several Yb^{3+} and Er^{3+} based pyrochlore rare-earth oxides [152]. Our pf-FRG calculations reveal a pronounced shift of the phase boundaries relative to the classical model. We examine the implications of these results for $\text{Yb}_2\text{Ti}_2\text{O}_7$, which sits directly at this boundary, by computing quantum phase diagrams around different experimental parameter estimates for this compound.

To characterize the different paramagnetic regimes revealed by our pf-FRG calculations, we compare the quantum correlations with classical results obtained from the self-consistent Gaussian approximation (SCGA)[182, 183]. The SCGA extends the unconstrained Luttinger–Tisza approach introduced in Sec. 3.1 to finite temperatures by enforcing only the weak spin constraint(3.8) through a self-consistently determined Lagrange multiplier at each T . This provides an approximation to the classical spin-spin correlations, $\chi^{ab}(\mathbf{q}, T) = \langle S^a(\mathbf{q})S^b(-\mathbf{q}) \rangle$, which can be directly compared to the pf-FRG structure factor defined in Eq. (4.136). The method is equivalent to a large- N expansion in the number of spin components and is often referred to as the large- N approximation. The SCGA calculations were performed by Daniel Lozano-Gómez and are described in detail in our joint publication Ref. [P4], which he co-authored.

The remainder of this chapter is organized as follows. We begin by introducing the nearest-neighbor model on the pyrochlore lattice and reviewing its classical phase diagram (largely following Refs.[24, 152]), highlighting the spin-nematic and pinch-line spin liquid phases. We then present the corresponding quantum phase diagram and classify the emerging paramagnetic regimes by comparing their structure factors with classical results. Finally, we focus on the parameter regime relevant to $\text{Yb}_2\text{Ti}_2\text{O}_7$ and discuss the implications of our findings for this compound.

This chapter is closely based on Ref. [P4], and most figures are redrawn from that work. The data for all quantum calculations were obtained by the author of this thesis, while the SCGA data were provided by Daniel Lozano-Gómez.

5.1 Spin Hamiltonian

Since the $4f$ electrons in pyrochlore rare-earth oxides are highly localized, the single-ion physics dominates over the two-ion exchange interactions and determines the effective magnetic moments in a given material. The hierarchy of energy scales is such that the Coulomb interaction is strongest, followed by spin-orbit coupling, and finally the crystal field. The coulomb and spin-

orbit energy can be minimized by following Hund's rules discussed in Chapter 2.1 leading to a ground state manifold with total angular momentum J and a $(2J+1)$ -fold degeneracy [24]. This degeneracy is lifted by the crystal field generated by the surrounding oxygen ions. Here we focus on Kramers ions with an odd number of f electrons, where Kramers' theorem ensures that the lowest-energy state is always a doublet associated to an odd J . Using these two states as the basis of the ground-state manifold and labeling them as $|\uparrow\rangle$ and $|\downarrow\rangle$, we can define an effective pseudo-spin operator $\tilde{\mathbf{S}}$ in this subspace as [24]

$$\tilde{S}^z \equiv \frac{|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|}{2}, \quad \tilde{S}^+ \equiv |\uparrow\rangle\langle\downarrow|, \quad \tilde{S}^- = |\downarrow\rangle\langle\uparrow|. \quad (5.7)$$

This operator transforms under symmetry just like a conventional $S = 1/2$ spin for the Yb^{3+} and Er^{3+} ions relevant to this work. As a concrete example, Yb_3^{+} ions have a $4f^{13}$ electron configuration, which corresponds to a single hole in the $4f$ manifold. The ground-state manifold has total angular momentum $J = |L+S| = |3+1/2| = 7/2$ consistent with Hund's rules described in Sec. 2.1. A strong crystal field splits the energy of the $(2J+1)$ degenerate manifold to a Kramers doublet comprised primarily of $m_J = \pm 1/2$ states that are well separated in energy from the excited states [184]. A description in terms of the corresponding effective $S = 1/2$ operators should thus be a good description of the material [24]. For other Kramers ions, however, the \tilde{S}^\pm components may correspond to parts of higher multipole moments (e.g., magnetic quadropoles) rather than conventional dipoles, a case we do not pursue here [24].

The hat in above definition indicates that the spin operator is defined in a *local* frame, where \tilde{S}^z points along the local $\langle 111 \rangle$ axis, denoted \tilde{z} . For the transverse axes \tilde{x} and \tilde{y} we follow the convention of Ref. [152], with explicit definitions provided in Appendix B. In this basis the g -tensor is typically diagonal and, when projecting onto the ground-state doublet, the magnetic moment takes the form

$$\boldsymbol{\mu}_i = -\mu_B \left[g_\perp (\tilde{x}_i \tilde{S}_i^x + \tilde{y}_i \tilde{S}_i^y) + g_\parallel \tilde{z}_i \tilde{S}_i^z \right]. \quad (5.8)$$

For $g_\parallel \gg g_\pm$ the moments are constrained along the local \tilde{z} direction, as in spin ice, while for $g_\parallel \ll g_\perp$ they take on XY character. For generality, in the following we set $g_\parallel = g_\perp = 1$, since the materials of interest lie between these extremes and anisotropic g -tensors only introduce qualitative modifications to the correlations [152].

Although, in principle, interaction within the full J manifold can be relevant for rare-earth oxides, when the ground-state doublet is well separated in energy from the first excited state it is a reasonable approximation to project the exchange interactions into this subspace. This yields an effective spin Hamiltonian directly for the pseudospins $\tilde{\mathbf{S}}$. Owing to the strong localization of the $4f$ orbitals, it is often sufficient to restrict to nearest-neighbor interactions, for which the most general symmetry-allowed Hamiltonian reads [24]

$$H = \sum_{\langle ij \rangle} \left[J_{zz} \tilde{S}_i^z \tilde{S}_j^z - J_\pm (\tilde{S}_i^+ \tilde{S}_j^- + \tilde{S}_i^- \tilde{S}_j^+) \right. \\ \left. + J_{\pm\pm} (\gamma_{ij} \tilde{S}_i^+ \tilde{S}_j^+ + \gamma_{ij}^* \tilde{S}_i^- \tilde{S}_j^-) - J_{z\pm} (\gamma_{ij}^* \tilde{S}_i^z \tilde{S}_j^+ + \gamma_{ij} \tilde{S}_i^z \tilde{S}_j^- + (i \leftrightarrow j)) \right]. \quad (5.9)$$

In the first line, the term $\sim J_{zz}$ corresponds to the classical spin-ice Hamiltonian, while $\sim J_\pm$ represents XXZ-type corrections. The second line introduces *bond-dependent* phase factors γ_{ij} (defined in Appendix B). For generic couplings, this Hamiltonian has *no* residual $\text{SU}(2)$ spin symmetry. The classical spin-ice phase is recovered in the limit $J_{zz} > 0$ with all other couplings set to zero. For non-Kramers ions, the pseudospin transforms differently from a conventional

$S = 1/2$ [24], enforcing $J_{z\pm} = 0$ by symmetry. In this case, the spin-ice state is stable over a finite region with $|J_{\pm}|, |J_{\pm\pm}| \ll J_{zz}$ [26]. In this chapter, however, we focus on the Kramers case, where $J_{z\pm} \neq 0$ and classical spin ice is absent, but a variety of other interesting phases emerge as we will explore in the following sections.

Although the local $\langle 111 \rangle$ frame is convenient for writing down the Hamiltonian, we carry out our calculations in the global frame, which corresponds to the lab frame. We denote the spin operators in this frame by \mathbf{S}_i , related to the local operators $\hat{\mathbf{S}}_i$ by a sublattice-dependent basis transformation (see Appendix B for details). In the global frame the Hamiltonian takes the form

$$H = \sum_{\langle ij \rangle} \mathbf{S}_i^T \mathbf{J}_{ij} \mathbf{S}_j, \quad (5.10)$$

where the coupling matrix \mathbf{J}_{ij} depends on the relative positions of sites i and j within the tetrahedral unit cell. For example,

$$\mathbf{J}_{01} = \begin{pmatrix} J_2 & J_4 & J_4 \\ -J_4 & J_1 & J_3 \\ -J_4 & J_3 & J_1 \end{pmatrix} = \begin{pmatrix} J_H + K & D/\sqrt{2} & D/\sqrt{2} \\ -D/\sqrt{2} & J_H & \Gamma \\ -D/\sqrt{2} & \Gamma & J_H \end{pmatrix}. \quad (5.11)$$

We use the parametrization by the four couplings J_1, J_2, J_3, J_4 . In the literature, it is also common to use an equivalent parametrization in terms of a Heisenberg coupling J_H , a Kitaev coupling K , a Dzyaloshinskii–Moriya coupling D , and a symmetric off-diagonal exchange Γ . The matrices on all other bonds are related to \mathbf{J}_{01} by simple rotations that effectively only permute the rows and columns.²

There exists a duality in this global parameterization: the sign of $J_{z\pm}$ in the local Hamiltonian is not fixed, since a C_2 rotation by π around the local \tilde{z} axis maps $J_{z\pm} \rightarrow -J_{z\pm}$ but otherwise leaves the Hamiltonian invariant. In the global frame, however, this transformation mixes the couplings and gives a new set $\bar{J}_1, \bar{J}_2, \bar{J}_3, \bar{J}_4$ (or equivalently $\bar{J}_H, \bar{K}, \bar{D}, \bar{\Gamma}$) [185], which will be relevant when we compare our results with literature later. We list the explicit transformation matrices between the local and global couplings in Appendix B.

5.2 Classical phase diagram

As a first step toward understanding the Hamiltonian (5.10), we outline how to determine the ground-state phase diagram for *classical* spins \mathbf{S}_i , following Ref. [152]. We then focus on the regime $J_4 = 0$ and $J_3 < 0$, where the phase diagram hosts both the pinch-line spin liquid and a spin-nematic phase, which we discuss in more detail.

5.2.1 Deriving the classical phase diagram

Every nearest-neighbor bond in the pyrochlore Hamiltonian belongs to a single tetrahedron. This allows us to rewrite the Hamiltonian as a sum over tetrahedra as

$$H = \sum_t H_t, \quad (5.12)$$

²We note that for a pf-FRG implementation, where we typically specify all bonds J_{ij} and J_{ji} explicitly, J_4 (or equivalently D) has to change sign under site exchange

with the single-tetrahedron Hamiltonian

$$H_t = \sum_{\mu, \nu=0}^3 \mathbf{S}_\mu \mathbf{J}_{\mu\nu} \mathbf{S}_\nu, \quad (5.13)$$

where $\mathbf{J}_{\mu\nu}$ are coupling matrices as defined in Eq. (5.11). It can be shown that, for classical spins, any state that minimizes the single-tetrahedron Hamiltonian H_t can be extended to the entire lattice and is therefore also a ground state of the full Hamiltonian [152]. This also implies that, for any set of couplings, a $\mathbf{q} = 0$ state belongs to the ground-state manifold—a four-sublattice configuration that repeats across all tetrahedral unit cells.

The task then reduces to finding ordered ground states of H_t , which can be classified by how they break the point-group symmetries of a single tetrahedron T_d . Introducing order parameters \mathbf{m}_λ that transform according to the irreducible representations $\lambda = A_2, E, T_{1-}, T_{1+}, T_2$ of T_d , the Hamiltonian takes the diagonal form [152]

$$H_t = \frac{1}{2} \left[a_{A_2} m_{A_2}^2 + a_E m_E^2 + a_{T_2} m_{T_2}^2 + a_{T_{1-}} m_{T_{1-}}^2 + a_{T_{1+}} m_{T_{1+}}^2 \right], \quad (5.14)$$

where $a_\lambda(J_\mu)$ are scalar functions depending only on the couplings, and $\mathbf{m}_\lambda(J_\mu, S_\mu)$ additionally depend on the four classical spins in a tetrahedron (explicit forms are given in Appendix B). Imposing the constraint of constant spin length $|\mathbf{S}_i|^2 = 1/4$ implies [152]

$$\sum_\lambda \mathbf{m}_\lambda^2 = 1 \quad \text{and} \quad \max_\lambda \mathbf{m}_\lambda^2 = 1. \quad (5.15)$$

The classical ground state can therefore be determined by identifying the irrep λ^* with a minimal prefactor $a_{\lambda^*}^*$ and then choosing a spin configuration that maximizes $m_{\lambda^*}^2 = 1$.

For example, states with $m_{A_2}^2 = 1$ correspond to the all-in-all-out configuration, where all spins point either into or out of the tetrahedron. The E and T_2 irreps describe different types of antiferromagnets, while the $T_{1\pm}$ irreps correspond to splayed ferromagnets, with spins canted away from a common axis by an angle θ . Importantly, the E states form a one-dimensional $U(1)$ manifold (although in practice fluctuations usually select a discrete subset), whereas all other irreps have discrete ground-state manifolds.

If only one a_λ is minimal, the ground state is a unique ordered $\mathbf{q} = 0$ phase. More interesting physics arises when two or more a_λ are degenerate: the resulting manifold can be subextensive [the number of states scaling as $O(2^L)$ or $O(2^{L^2})$] or even extensive [$O(2^{L^3})$ as in spin ice], depending on which irreps are involved. Such degeneracies occur only at the phase boundaries between ordered states.

5.2.2 Magnetically ordered phases

The above analysis applies throughout the entire parameter space, and the full classical phase diagram is presented in Refs. [25, 26]. In the following we focus on the plane with $J_3 < 0$ and $J_4 = 0$, which was analyzed in detail in Ref. [152]. Notably, the compounds $Yb_2Ti_2O_7$, $Er_2Sn_2O_7$ (both located near phase boundaries), and $Er_2Ti_2O_7$ are all expected to lie approximately within this parameter plane [152].

The corresponding classical phase diagram is shown in Fig. 5.2(a). It consists of three ordered phases: T_{1-} (FM), T_2 (AFM), and E (AFM). Within the E manifold, thermal fluctuations select a discrete set of states, either coplanar (ψ_3) or noncoplanar (ψ_2) [152], not explicitly shown here. The chance for more interesting physics occurs only at the phase boundaries. At the $T_{1-} - E$

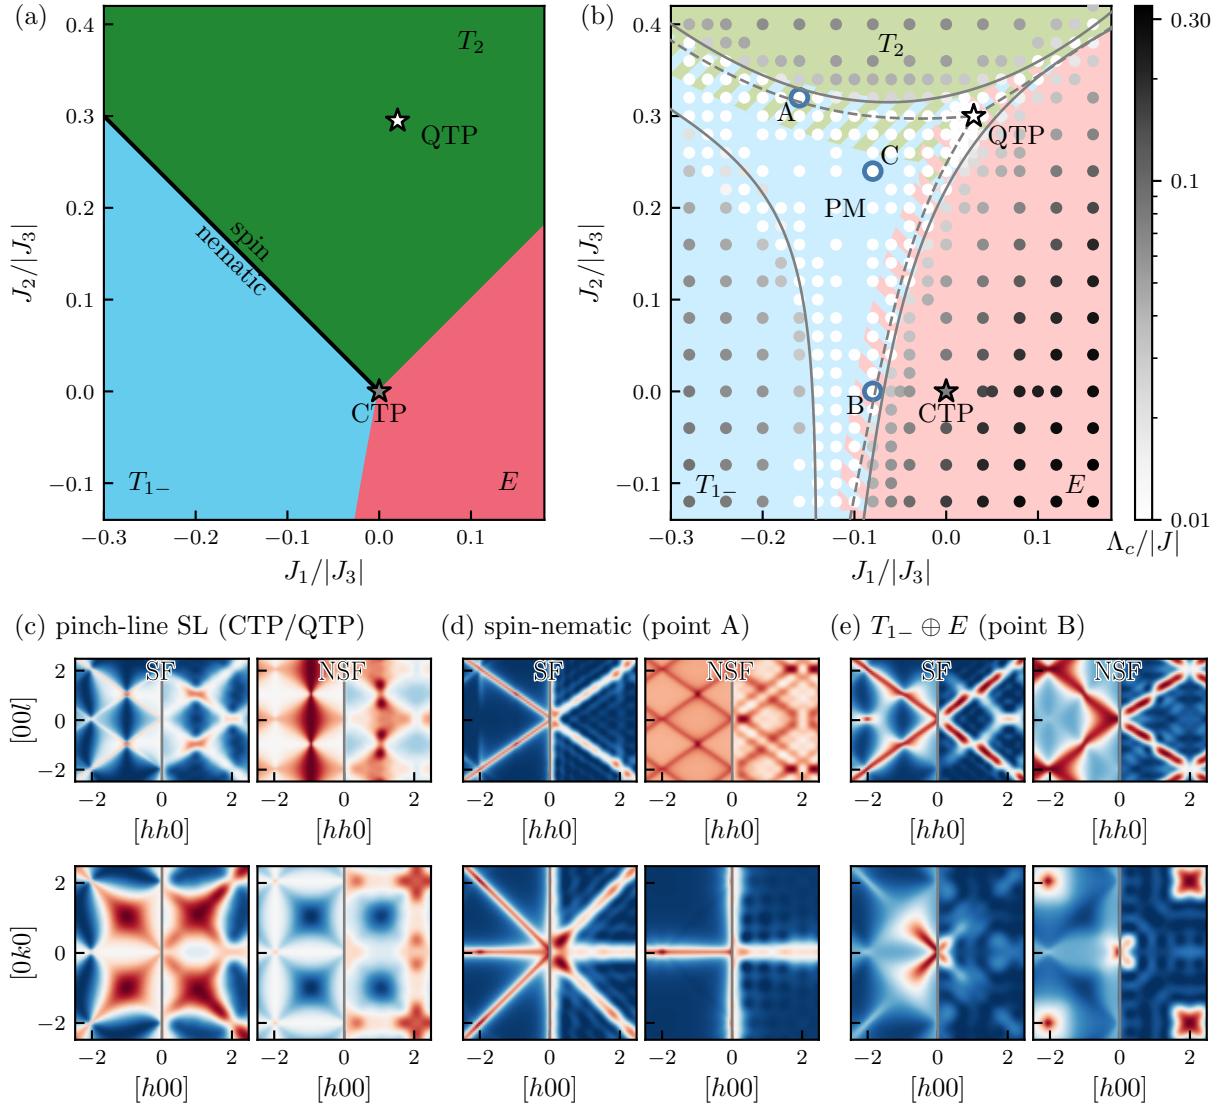
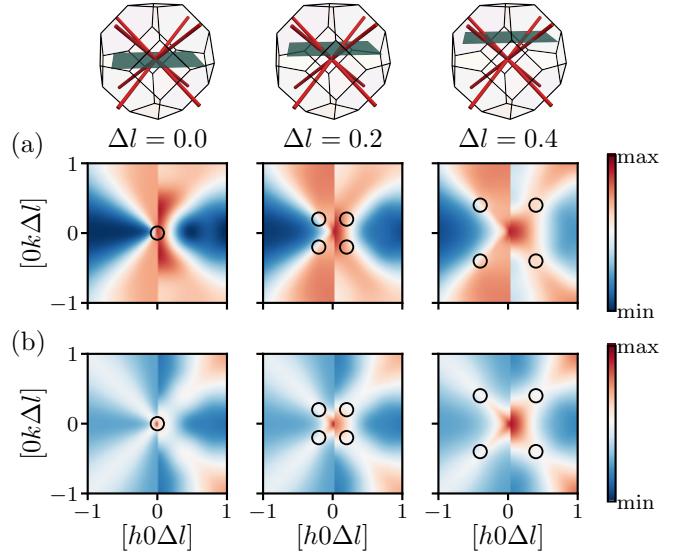


Figure 5.2 – Classical and quantum phase diagrams and neutron-scattering structure factors. (a) Exact classical phase diagram for $J_4 = 0$ and $J_3 < 0$. Three $\mathbf{q} = 0$ phases meet at the classical triple point (CTP), where the ground state is the pinch-line spin liquid. (b) Quantum phase diagram from pf-FRG. Colors indicate the dominant order-parameter susceptibility; hatched regions mark areas where multiple susceptibilities gain significant weight during the flow. Dashed lines indicate where the dominant susceptibility changes, meeting at the quantum triple point (QTP) at $J_1 \approx 0.03$, $J_2 \approx 0.3$ where all three susceptibilities are maximally degenerate. The solid lines mark the approximate phase boundaries of a PM region with no dipolar magnetic order (indicated by white markers). Outside this region, the marker saturation reflects the critical scale Λ_c at which a flow breakdown occurs, signaling the onset of dipolar magnetic order. (c–e) Polarized neutron-scattering structure factors in the spin-flip (SF, left) and non-spin-flip (NSF, right) channels, shown in the $[hhl]$ plane (top) and $[hk0]$ plane (bottom), at representative points. In each panel, the left half shows the classical SCGA result, and the right half the $S = 1/2$ pf-FRG result. The underlying parameters $(J_1/|J_3|, J_2/|J_3|)$ for SCGA/pf-FRG are: (c) $(0.0, 0.0)/(0.03, 0.3)$, (d) $(-0.28, 0.28)/(-0.28, 0.36)$, and (e) $(-0.0189, -0.1)/(-0.08, 0.0)$.

Figure 5.3 –

Pinch-line singularities and fourfold pinch-point in order-parameter susceptibilities. (a) and (b) show the susceptibilities $\langle m_E^1(\mathbf{q})m_E^1(-\mathbf{q}) \rangle$ and $\langle m_{T_2}^z(\mathbf{q})m_{T_2}^z(-\mathbf{q}) \rangle$, respectively. For each, the left side display classical results from SCGA at the CTP, and the right side quantum $S = 1/2$ results from pf-FRG at the QTP. Each column corresponds to a different horizontal cut through the first Brillouin zone of the pyrochlore lattice, parameterized by $q_z = \Delta l$ as indicated in the schematic above. The correlations show broadened pinch-points where the horizontal cuts intersect with the $\langle 111 \rangle$ directions, which are highlighted by black circles in (a, b) and red rods in the Brillouin zone schematic. For $m_{T_2}^z$ these merge at $[000]$ in a fourfold pinch-point indicative of a higher-rank emergent gauge theory.



boundary, thermal fluctuations select a finite set of states within the E manifold [121, 186]. At the $E-T_2$ boundary, a similar scenario is likely, although we have not found conclusive results in the literature. At the T_1-T_2 boundary, classical Monte Carlo simulations show the emergence of a spin-nematic state, characterized by a quadrupolar order parameter (quadratic in the spin operators) [187]. Finally, at the point $J_1 = J_2 = J_4 = 0$, where all three phases meet—which we refer to as the *classical triple point* (CTP)—thermal fluctuations fail to stabilize order, and instead the system realizes a higher-rank pinch-line spin liquid [181], which we discuss in the following.

5.2.3 Pinch-line spin liquid

At the CTP, the prefactors a_λ of the single-tetrahedron Hamiltonian introduced in Eq. (5.14) satisfy $a_E = a_{T_{1-}} = a_{T_2} < m_{A_2}, a_{T_{1+}}$, indicating that the ground state is determined by the constraints

$$m_{A_2} = 0, \quad \mathbf{m}_{T_{1+}} = 0 \quad (5.16)$$

on every tetrahedron. These constraints do not fully fix an ordered ground state but instead leave the fields $\mathbf{m}_{T_{1-}}, \mathbf{m}_{T_2}, \mathbf{m}_E$ free to fluctuate, leading to an extensive ground-state degeneracy.

Upon coarse-graining, these constraints—combined with the strong spin-constraint of fixed spin length—can be recast as conservation laws for a matrix \mathbf{B} constructed from the E , T_{1-} and T_2 irrep fields as [26, 181]

$$\mathbf{B} = \begin{pmatrix} 2m_E^1 & \sqrt{3}m_{T_2}^z & -\sqrt{3}m_{T_2}^y \\ -\sqrt{3}m_{T_2}^z & -m_E^1 + \sqrt{3}m_E^2 & \sqrt{3}m_{T_2}^x \\ \sqrt{3}m_{T_2}^y & -\sqrt{3}m_{T_2}^x & -m_E^1 - \sqrt{3}m_E^2 \end{pmatrix} - 3 \sin \theta \begin{pmatrix} 0 & m_{T_{1-}}^z & m_{T_{1-}}^y \\ m_{T_{1-}}^z & 0 & m_{T_{1-}}^x \\ m_{T_{1-}}^y & m_{T_{1-}}^x & 0 \end{pmatrix}, \quad (5.17)$$

where θ is a function of the coupling parameters J_μ , and m_λ^a are the components of the order parameters. For more details we refer the reader to Appendix B and Ref. [26]. In terms of this

\mathbf{B} -field, the constraints can be expressed as two generalized Gauss's laws [26, 181], namely

$$|\epsilon_{abc}| \partial_a B^{bc} = 0 \quad \text{and} \quad \partial_a B^{ab} = 0, \quad (5.18)$$

where ϵ_{abc} is the fully antisymmetric tensor. The ground state of the CTP is thus an example of a realization of an emergent higher-rank gauge theory.

Following the classification of spin liquids in Ref. [27, 188], this state can be further identified as an algebraic pinch-line spin liquid. The eigenenergies obtained by Fourier transforming the interaction matrix yield a flat band at the bottom of the spectrum, representing the extensive degeneracy. More importantly, the band structure is gapless in an unusual way: the first dispersive band touches the flat band not only at isolated points but along a one-dimensional line in momentum space—a nodal line. This nodal line leads to a pinch-line singularity—a line in reciprocal space along which pinch-point features appear in different spin-spin correlations [181]. The gapless nature implies that all spin-spin correlations decay algebraically.

The pinch-point singularities can be seen in the polarized neutron scattering structure factors [defined in Eqs. (4.143, 4.144)] obtained from the SCGA shown on the left-hand-side of Fig. 5.2(c), which exhibit twofold pinch points and pinch lines along the $\langle 111 \rangle$ and symmetry-related directions. The pinch-line singularity appears more clearly in the order-parameter susceptibilities, defined as

$$\chi_{\lambda}^a(\mathbf{q}) = \langle m_{\lambda}^a(\mathbf{q}) m_{\lambda}^a(-\mathbf{q}) \rangle = \frac{1}{N_{\text{uc}}} \sum_{t,t'} \exp[-i\mathbf{q}(\mathbf{r}_t - \mathbf{r}_{t'})] \langle m_{\lambda}^a(\mathbf{r}_t) m_{\lambda}^a(\mathbf{r}_{t'}) \rangle, \quad (5.19)$$

where the sum runs over the N_{uc} tetrahedron unit cells of the pyrochlore lattice and \mathbf{r}_t denotes the positions of the tetrahedra's centers. The classical correlation functions of m_E^1 and $m_{T_2}^z$, shown on the left side in Fig. 5.3, display twofold pinch points where the scattering planes intersect the $\langle 111 \rangle$ directions, merging into a fourfold pinch point at the Γ -point—consistent with Ref. [181]. This fourfold pinch point is a direct hallmark of the higher-rank nature of the emergent gauge theory. Because of the pinch-line singularity, this state is referred to as a *pinch-line spin liquid*.

5.2.4 Spin nematic state

On the $T_{1-} - T_2$ boundary, the spectrum of the interaction matrix again shows a flat band, indicating a large degeneracy and possibly classical spin liquid behavior. In fact, the ground state exhibits an accidental U(1) symmetry, spanned by single-tetrahedron configurations that are mixtures of T_{1-} and T_2 states parametrized as $m^a(\theta) = m_{T_2}^a \cos(\varphi) + m_{T_{1-}}^a \sin(\varphi)$, which can be consistently tiled across the full lattice by varying φ on each tetrahedron [121, 187]. Extensive Monte Carlo simulations show that thermal fluctuations select a subset of states from this accidental U(1) manifold, accompanied by cubic-symmetry breaking manifested in the suppression of one of the three order-parameter components a [187]. All dipolar order parameters remain zero, confirming the absence of conventional magnetic order. Instead, spin-nematic order parameters of the form

$$Q_{\perp}^{\text{site}} = \left\langle \left| \frac{1}{N} \sum_i \left(\frac{(\tilde{S}_i^x)^2 - (\tilde{S}_i^y)^2}{2\tilde{S}_i^x \tilde{S}_i^y} \right) \right| \right\rangle, \quad (5.20)$$

and³

$$Q_{\perp}^{\text{bond}} = \left\langle \left| \frac{1}{3N} \sum_{\langle ij \rangle} \begin{pmatrix} \tilde{S}_i^x \tilde{S}_j^x - \tilde{S}_i^y \tilde{S}_j^y \\ \tilde{S}_i^x \tilde{S}_j^y + \tilde{S}_i^y \tilde{S}_j^x \end{pmatrix} \right| \right\rangle, \quad (5.21)$$

become finite at a critical temperature, most likely through a first-order phase transition. On this basis, Ref. [187] identified the resulting phase as a *spin-nematic* state.

The spin structure factor in this phase is shown on the left side of Fig. 5.2(d) and exhibits distinctive rods of scattering. These originate from low-energy bands of the interaction matrix, which are flat along the $\langle 111 \rangle$, $\langle 001 \rangle$, and symmetry-related directions in momentum space—precisely where the rods appear.

5.3 Effects of quantum fluctuations in the $S = 1/2$ model

We now turn to the role of quantum fluctuations in the $S = 1/2$ model. Using pf-FRG, we calculate the quantum phase diagram in the same parameter region as in the classical case, with particular focus to the fate of the pinch-line spin liquid and the spin-nematic state.

5.3.1 Quantum phase diagram

To obtain the quantum phase diagram, we use the pf-FRG introduced in detail in Chapter 4. The main output of this method is the RG flow of the spin-spin susceptibility $\chi_{ij}^{\Lambda ab}$ defined in (4.134). Our calculations are typically carried out on lattice truncations $L = 3, 5, 7$ (with $L = 7$ including 864 correlated sites) and a frequency grid of $35 \times 30 \times 30$ Matsubara frequencies.

We first distinguish dipolar ordered phases from paramagnetic (PM) ones using the flow-breakdown criterion described in Sec. 4.5.1: a divergence (or numerical precursor thereof) in the flow of the structure factor [Eq. (4.136)] at a finite critical scale Λ_c signals the onset of dipolar magnetic order, while its absence identifies a PM state (representative flows of the structure factor are shown in Appendix B). To further classify both ordered and disordered phases, we track the flow of the summed order-parameter susceptibilities defined in Eq. (5.19) at $\mathbf{q} = 0$, which evaluate to⁴

$$\chi_{\lambda} \equiv \sum_a \chi_{\lambda}^a(\mathbf{q} = 0) = \frac{1}{N_{\text{uc}}} \sum_{t, t'} \langle \mathbf{m}_{\lambda}(\mathbf{r}_t) \cdot \mathbf{m}_{\lambda}(\mathbf{r}_{t'}) \rangle, \quad (5.22)$$

so that, when $\mathbf{q} = 0$ order in irrep λ^* is present, $\chi_{\lambda^*}^{\Lambda}$ dominates and should show a flow breakdown. The resulting quantum phase diagram is shown in Fig. 5.2(b) and cuts through this phase diagram are provided in Appendix B. Two key differences compared to the classical diagram stand out:

- (i) the E -ordered phase expands noticeably, leading to a substantial shift of the phase boundaries. This is indicated by the background color, which highlights the dominant susceptibility, and the dashed lines, which mark where the dominant susceptibility χ_{λ} changes (this is more clearly illustrated in Fig. 5.6(a), which shows the maximal order-parameter susceptibility at each point).

³This is only explicitly shown in Ref. [187] for Q_{\perp}^{site} but the authors confirmed that Q_{\perp}^{bond} also becomes finite, which is essential for the realization of a quantum analog of the state as quantum spin $S = 1/2$ obey $(S_i^a)^2 \equiv 1/4$ and can thus not realize a quadrupole moment on a single site.)

⁴Since the order-parameter fields \mathbf{m}_{λ} are linear combinations of the spin operators (defined in Appendix B), the product $\langle \mathbf{m}_{\lambda}(\mathbf{r}_t) \cdot \mathbf{m}_{\lambda}(\mathbf{r}_{t'}) \rangle$ can be directly obtained from the spin-spin susceptibilities χ_{ij}^{ab} computed in the pf-FRG.

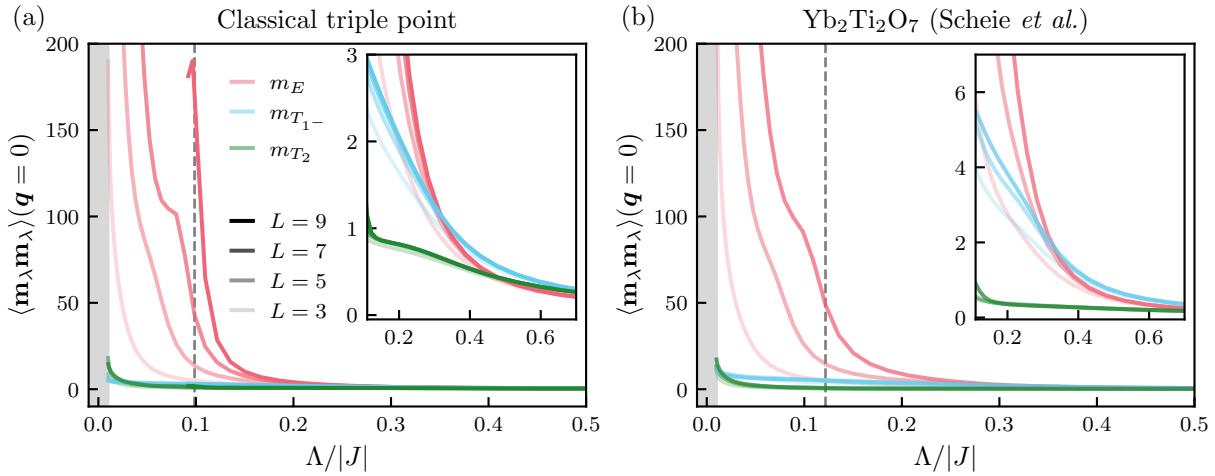


Figure 5.4 – Order-parameter susceptibility flows at the CTP (top) and for the Yb₂Ti₂O₇ parameters from Scheie *et al.* [179] (bottom) (numerical values in Table 5.1). At both points, flows show clear signatures of a flow breakdown at finite critical scale Λ_c (dashed gray lines) indicating the onset of magnetic order. Insets show zooms at larger cutoffs, highlighting that clear E -order dominance appears only in the low-cutoff limit.

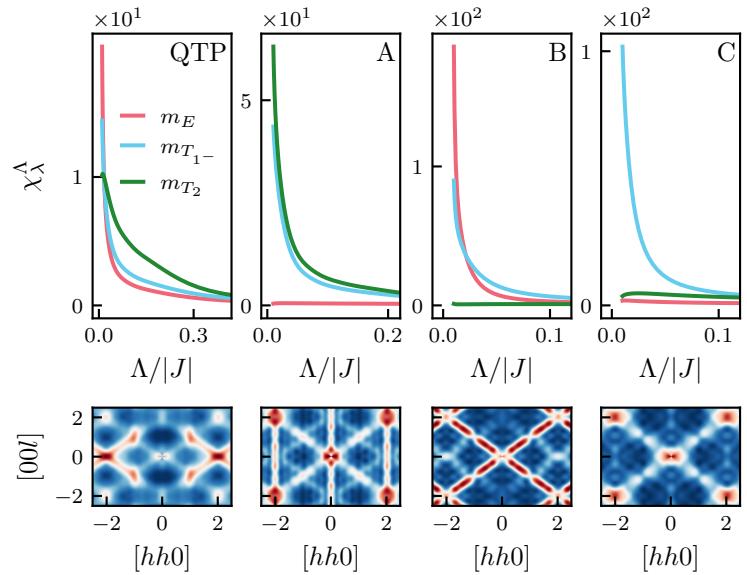
- (ii) a large paramagnetic region emerges that shows *no* sign of a flow breakdown, whose approximate phase boundaries are indicated by the solid lines in Fig. 5.2(b). Interestingly, this region is not centered on the point of maximal classical phase competition (the CTP), contrary to conventional expectations.

Concerning point (i), a shift of comparable size was also found in nonlinear spin-wave theory [189] and exact diagonalization [186] for the E - T_{1-} boundary (albeit at slightly different parameter values). Surprisingly, the classical triple point (CTP), where the classical ground state is the pinch-line spin liquid, now lies well inside the E -ordered phase and exhibits a clear flow breakdown [see Fig. 5.4(a)]. This shows that quantum fluctuations stabilize an ordered state out of a classically degenerate manifold, whereas thermal fluctuations do not [181]. This effect is beyond the reach of linear spin-wave theory, which also predicts a disordered state at the CTP [152]. The nature of the paramagnetic regime mentioned in point (ii) is discussed in the following.

5.3.2 Paramagnetic regime

In much of the paramagnetic regime, the T_{1-} order-parameter susceptibility is dominant, as indicated by the blue background color in Fig. 5.2(b). Closer to the boundaries, however, we find finite regions where multiple susceptibilities gain significant weight, as illustrated for the representative points QTP, A, B, and C in Fig. 5.5. The corresponding spin structure factors also differ qualitatively depending on which irreps are involved, as shown in the same figure. This suggests two possible scenarios: (i) the PM regime is a single phase in which short-range correlations evolve continuously with parameters, or (ii) it consists of multiple distinct phases, possibly including symmetry-breaking states such as spin nematics, as well as symmetric quantum spin liquids. As discussed in Sec. 4.5.3, our pf-FRG framework cannot unambiguously distinguish these cases, since detecting nematic or dimer order requires access to four-spin susceptibilities, which are currently out of reach. Nonetheless, the strongly varying correlation patterns observed in pf-FRG point toward scenario (ii).

Figure 5.5 – Susceptibility flows and structure factors in the paramagnetic regime. Shown are results for points labeled QTP, A, B, and C in Fig. 5.2(b). The top row shows the RG flow of the order-parameter susceptibilities $\langle m_\lambda \cdot m_\lambda \rangle$ for $\lambda \in E, T_{1-}, T_2$. Smooth flows down to $\Lambda \rightarrow 0$ indicate the absence of dipolar magnetic order, while the relative magnitudes highlight the dominant competing correlations. The bottom row shows the total neutron-scattering structure factor S_\perp . Parameters $(J_1/|J_3|, J_2/|J_3|)$ are $(0.03, 0.3)$, $(-0.28, 0.36)$, $(-0.08, 0.0)$ and $(-0.08, 0.24)$ for QTP, A, B, and C, respectively, with $J_3 < 0$ and $J_4 = 0$.



Motivated by this, we classify the different paramagnetic regimes by examining which order-parameter correlations remain degenerate in the low-cutoff limit⁵. This analysis yields four regimes, labeled $E \oplus T_{1-} \oplus T_2$, $T_{1-} \oplus T_2$, $E \oplus T_{1-}$, and T_{1-} -only, which exhibit the same degeneracies and similar correlations as points QTP, A, B, and C, respectively. To estimate the extent of these regimes, we define the normalized relative susceptibility

$$\chi_\lambda^{\text{rel}} = \frac{\chi_\lambda}{\max_{\lambda^*} \chi_{\lambda^*}}, \quad (5.23)$$

which is expected to be large for the irreps that are degenerate within each regime. This is illustrated in Fig. 5.6(b–d), from which the approximate extent of each regime can be inferred. Regions where more than one relative irrep susceptibility exceeds 20% are additionally highlighted with hatched backgrounds in Fig. 5.2(b).

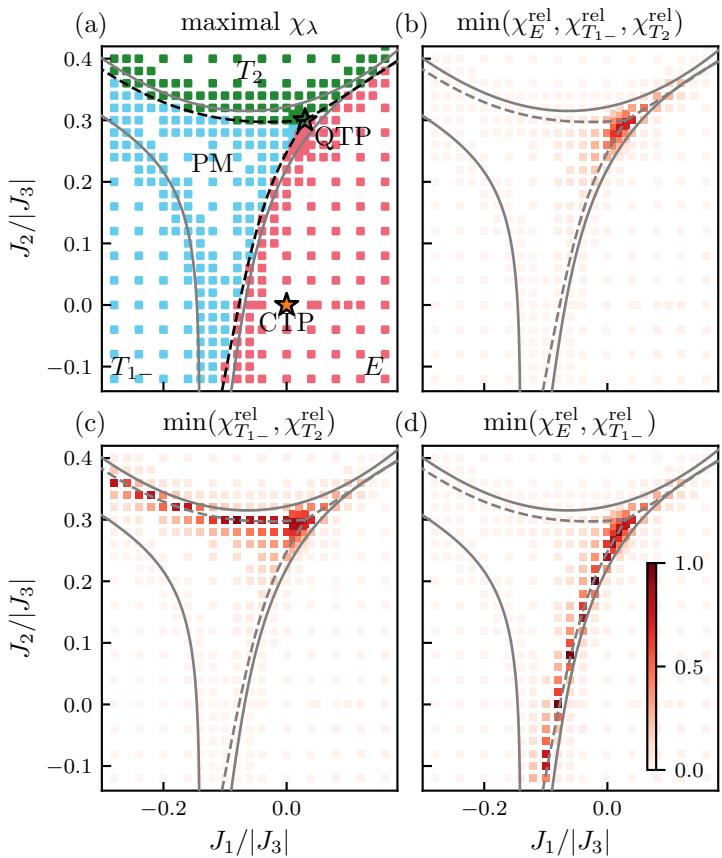
In the following, we analyze the correlations and structure factors in each of these paramagnetic regimes and compare them to classical results. This reveals signatures consistent with quantum analogues of both the pinch-line spin liquid and the spin-nematic state.

Pinch-line spin liquid in the $E \oplus T_{1-} \oplus T_2$ regime The most exotic state in the classical phase diagram is the higher-rank pinch-line spin liquid at the CTP. While pf-FRG predicts E order at the CTP, the large shift in phase boundaries suggests that a state similar to the pinch-line spin liquid may still occur, albeit at a shifted location in the phase diagram. The natural candidate is the region where the order-parameter susceptibilities of E , T_{1-} , and T_2 become maximally degenerate, illustrated in Fig. 5.6(b). Inside this region we determine the *quantum triple point* (QTP) by locating the parameters where $\min_{\lambda \in (E, T_{1-}, T_2)} \chi_\lambda^{\text{rel}}$ is maximized in the low-cutoff limit. This occurs at $J_1/|J_3| \approx 0.03$, $J_2/|J_3| \approx 0.3$ at the top right corner of the paramagnetic regime, precisely where the dashed lines marking changes in the dominant susceptibility meet in Fig. 5.2(b). The RG flows at the QTP, shown in Fig. 5.5, confirms that all three susceptibilities grow to a comparable magnitude.

⁵In our calculations, the low-cutoff limit corresponds to $\Lambda = 0.02/|J|$, with the normalization $|J|^2 = J_1^2 + J_2^2 + J_3^2 + J_4^2$.

Figure 5.6 –

Maximal order-parameter susceptibilities and degeneracies. (a) Irrep λ^* with maximal susceptibility in the low-cutoff limit. The dashed lines, also shown in Fig. 5.2(b), indicate where λ^* changes. These lines converge at the QTP, where the irreps become maximally degenerate. (b-d) Minimal relative susceptibilities between different irreps. Here we plot the minimal normalized susceptibility among the relevant sets to highlight regions of degeneracy. From these plots, one can identify the approximate boundaries of distinct paramagnetic regimes [hatched backgrounds in Fig. 5.2(b)], corresponding to (b) $E \oplus T_{1-} \oplus T_2$, (c) $T_{1-} \oplus T_2$, and (d) $E \oplus T_{1-}$, as discussed in the main text.



And indeed, the correlations at the QTP strongly resemble those of the classical pinch-line spin liquid. The polarized neutron-scattering structure factors in Fig. 5.2(c) exhibit pinch-points and lines along the $\langle 111 \rangle$ directions, resembling the classical case. Likewise, the order-parameter correlations in Fig. 5.3 reproduce the pinch-line and fourfold pinch-points reported in Ref. [181]. We find similar correlations throughout the surrounding $E \oplus T_{1-} \oplus T_2$ region [the extent is visible in Fig. 5.6(b)], where all three order-parameter susceptibilities remain significant.

These findings suggest that a quantum analog of the classical pinch-line spin liquid—together with its higher-rank gauge structure—may survive in the quantum model and even extend into a finite phase around the QTP. We note that the pinch-point singularities are noticeably broadened compared to the classical case, a feature previously linked to violations of the gauge constraint. This implies that finite gauge-charge fluctuations may be induced here by quantum fluctuations [148, 172].

Spin-nematics in the $T_{1-} \oplus T_2$ regime Away from the quantum triple point, but close to the boundary of the paramagnetic regime with the T_2 phase, we observe a finite region where both T_{1-} and T_2 correlations have similar magnitude [shown in Fig. 5.6(c)]. Comparing the polarized neutron scattering structure factors of the quantum model in this region, with the classical correlations of the spin-nematic state exactly at the classical T_{1-} - E phase boundary, we again find striking similarity, as illustrated in Fig. 5.2(d). Both cases show clear rods of scattering along $\langle 111 \rangle$ and $\langle 100 \rangle$ directions. These findings suggest that, unlike in the classical case where the nematic order is only present strictly at the phase boundary, the quantum model may host a small but *extended* quantum nematic phase.

We note however, that the rods of scattering are mostly an indication of the near degeneracy of the T_{1-} and T_2 correlations. An actual proof of nematic order is, as mentioned above, beyond the scope of the pf-FRG due the quadrupolar nature of the order-parameter. Moreover, the method to probe for such symmetry breaking via response functions discussed in Sec. 4.5.3 is unfortunately also not applicable, as the $U(1)$ “symmetry” broken by the order parameter is only accidental and not a true symmetry of the Hamiltonian. Although our results are not conclusive, the possible realization of a spin-nematic state in this model is still fascinating, as no such phase has yet been established for $S = 1/2$ systems in three dimensions without an applied magnetic field. This makes it a particularly promising avenue for further exploration with complementary theoretical and numerical methods.

Correlations in the $E \oplus T_{1-}$ regime and their connection to $\text{Yb}_2\text{Ti}_2\text{O}_7$ Near the phase boundary between the PM and E phases, the E and T_{1-} order-parameter susceptibilities become nearly degenerate (see Fig. 5.6(d)). The absence of a flow breakdown points to either a spin liquid or an unconventional type of order, which we cannot further specify. The quantum structure factor [shown in Fig. 5.5 and Fig. 5.2(e)] in this regime shows rods of scattering *only* along the $\langle 111 \rangle$ directions, with additional peaks at $[220]$. This closely resembles the classical structure factor at the $E-T_{1-}$ phase boundary at intermediate temperatures [also shown in Fig. 5.2(e)]. Strikingly, the same pattern has been observed in neutron-scattering experiments on $\text{Yb}_2\text{Ti}_2\text{O}_7$ just above its ordering transition, where the material is believed to realize a correlated paramagnet [180, 190]. Moreover, Ref. [180] argued that the dynamics of $\text{Yb}_2\text{Ti}_2\text{O}_7$ are best captured by a coexistence of E and T_{1-} correlations, precisely the type of behavior found in this regime. We return to the implications for $\text{Yb}_2\text{Ti}_2\text{O}_7$ in the next section.

Extended paramagnetic region dominated by T_{1-} correlations Finally, we identify a relatively large regime that shows no clear signs of a flow breakdown, with only the T_{1-} susceptibility growing significantly [blue background in the PM region of Fig. 5.2(b)]. The corresponding structure factor displays broad features along the $\langle 111 \rangle$ directions [see Fig. 5.5 point C], although these are less pronounced than in the other paramagnetic regimes. Two interpretations seem possible: either this is indeed a genuine paramagnetic region dominated by short-range T_{1-} correlations, or pf-FRG overestimates the extent of the paramagnetic regime and this region should instead be part of the ordered T_{1-} phase. Interestingly, an independent pf-FRG study [122] also reported an absence of ordering over a wide portion of the T_{1-} phase near the E boundary (in a different parameter regime), suggesting that this behavior is not merely an artifact of our specific implementation. In any case, the true nature of this regime remains an open question.

5.4 Implications for $\text{Yb}_2\text{Ti}_2\text{O}_7$

The large shift of phase boundaries has strong implications for materials located near such boundaries. A prime example is $\text{Yb}_2\text{Ti}_2\text{O}_7$, which, as discussed in the introduction, sits close to the classical $T_{1-}-E$ boundary, supposedly has a ferromagnetic ground state, but displays signatures of both orders in its dynamical neutron-scattering spectra [179, 180]. The mechanism resulting in this spectrum is still not clear. A possible scenario proposed in Ref. [179] is that quantum fluctuations enable tunneling between the two orders in time. Our finding of a mixed $T_{1-} \oplus E$ phase at $T = 0$, where exactly these correlations appear to coexist, indicates that this is indeed possible. Proximity to this regime may thus account for the observed spectra.

Table 5.1 – Exchange constants for $\text{Yb}_2\text{Ti}_2\text{O}_7$ from different experimental studies. The couplings were converted to the global frame (J_1, J_2, J_3, J_4) (see Appendix B for details) if not directly stated in the references, without accounting for uncertainties.

| Label | Reference (<i>et al.</i>) | J_1 (meV) | J_2 (meV) | J_3 (meV) | J_4 (meV) |
|-------|-----------------------------|-------------|-------------|-------------|-------------|
| (a) | Scheie [179] | -0.026 | -0.307 | -0.323 | 0.028 |
| (b) | Robert [191] | -0.03 | -0.32 | -0.28 | 0.02 |
| (c) | Thompson [178] | -0.028 | -0.326 | -0.272 | 0.049 |
| (d) | Ross [174] | -0.09 | -0.22 | -0.29 | 0.01 |

To make this connection more concrete, we now examine the quantum phase diagram in the vicinity of actual experimental parameter estimates for $\text{Yb}_2\text{Ti}_2\text{O}_7$. We have summarized the most relevant ones in Table 5.1. All agree on relatively small J_4 and negative J_3 , consistent with the regime considered above, but they also report large $J_2/|J_3|$ of order one, which places the material at the far south of the phase diagrams in Fig. 5.2, outside the region shown. To explore precisely this region, we calculate quantum phase diagrams near two parameter sets. We begin with the most recent and presumably most reliable estimates from Scheie *et al.* [179]. To validate our results, we then consider older estimates from Thompson *et al.* [178], in a regime where nonlinear spin-wave theory results already exist [189], with which we can directly compare. In both cases, we find that quantum fluctuations shift the phase boundaries such that $\text{Yb}_2\text{Ti}_2\text{O}_7$ falls into the AFM E phase, at odds with the experimentally observed FM T_{1-} ground state. We argue that this discrepancy could stem from the fitting procedure used to extract the parameters, which may need to be revisited in light of our results.

5.4.1 Phase diagram around parameters from Scheie *et al.*

To study the vicinity of the Scheie *et al.* parameters more directly, we fix $J_3 = -0.323$ meV and $J_4 = 0.028$ meV to their reported values and calculate a quantum phase diagram as a function of J_1 and J_2 , using the same procedure as in the previous section. The result is shown in Fig. 5.7(a). Consistent with our earlier findings, we observe a substantial shift of the classical T_{1-} – E boundary (also shown in this figure) by about $J_1/|J_3| \approx -0.1$, which significantly enlarges the E phase compared to the classical case. Shifts of comparable magnitude were also seen in the spin-1/2 model with exact diagonalization at $T = 0$, as well as in numerical linked-cluster (NLC) and high-temperature expansion (HTE) calculations [186]. Notably, this places the Scheie *et al.* parameter estimates for $\text{Yb}_2\text{Ti}_2\text{O}_7$ inside the E phase, inconsistent with the experimentally observed T_{1-} ground state.

Between the two ordered phases we again find a narrow paramagnetic region. The RG flows and polarized neutron-scattering structure factors along a cut from the T_{1-} through the PM to the E phase are shown in Fig. 5.7(b,c). In the ordered phases, either the T_{1-} or E susceptibility dominates and exhibits signs of a flow breakdown, while the structure factor shows sharp peaks at momenta consistent with the expected classical order. Inside the PM region, however, we identify a domain with dominant T_{1-} correlations but no flow breakdown, corresponding to the T_{1-} -only regime discussed above. Here the structure factor already develops faint rods of scattering along the $\langle 111 \rangle$ directions in addition Bragg peaks associated to T_{1-} correlations. For slightly larger J_1 , we encounter points (marked by a hatched background) where E and T_{1-} susceptibilities become degenerate, forming a $T_{1-} \oplus E$ regime. At these points, the structure

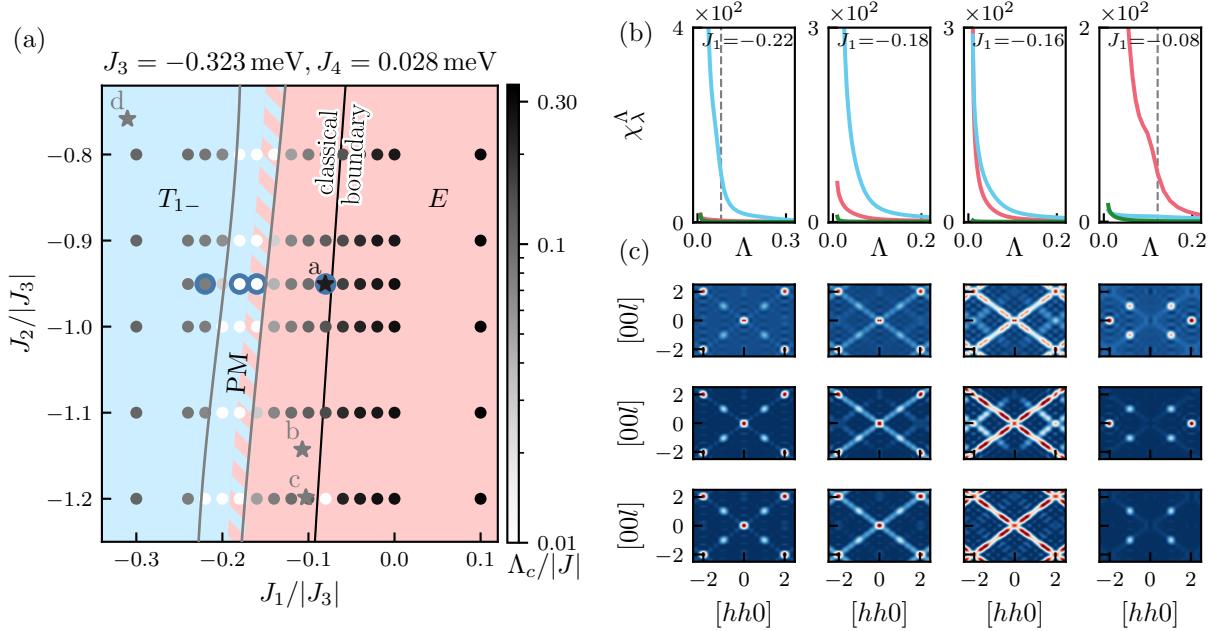


Figure 5.7 – Quantum phase diagram and correlations near $\text{Yb}_2\text{Ti}_2\text{O}_7$ parameters from Scheie *et al.* (a) Quantum phase diagram from pf-FRG with J_3 and J_4 fixed to the estimates of Scheie *et al.* [179]. Stars labeled ‘a’–‘d’ mark literature estimates of J_1 and J_2 for $\text{Yb}_2\text{Ti}_2\text{O}_7$ from different studies (numerical values in Table 5.1). The associated J_3 and J_4 values for ‘b’–‘d’ differ from those of ‘a’ (Scheie *et al.*), so these points do not lie exactly in the plane shown. Background colors follow the scheme in Fig. 5.2. (b,c) Results for the four parameter sets highlighted by blue circles in (a), chosen along the transition from the T_{1-} to E phase. Here J_2, J_3 , and J_4 are fixed to the exact Scheie *et al.* values, while J_1 is varied. (b) RG flows of the order-parameter susceptibilities for the relevant irreps, with dashed gray lines marking the critical scale Λ_c indicating the onset of magnetic order. (c) Corresponding neutron-scattering structure factors in the total (top), spin-flip (middle), and non-spin-flip (bottom) channel.

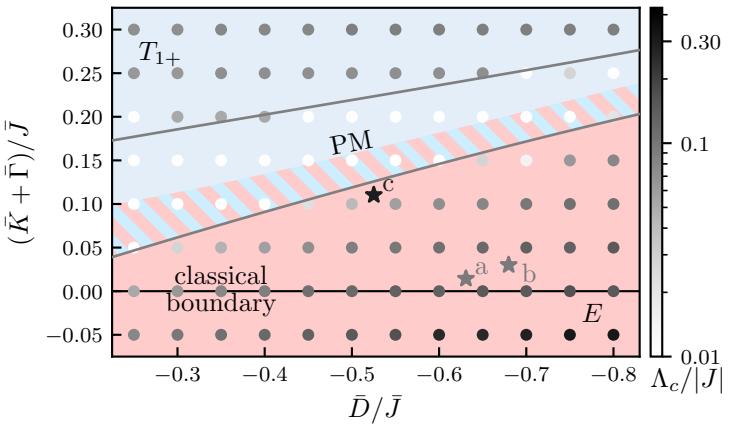
factor shows continuous rods of scattering along the $\langle 111 \rangle$ directions with peaks at $[220]$ closely resembling the experimental structure factor of $\text{Yb}_2\text{Ti}_2\text{O}_7$ in its correlated paramagnetic phase just above the ordering transition. This suggests that the rods seen experimentally may originate from proximity to this nonmagnetic $T_{1-} \oplus E$ phase. Ref. [180] proposed a similar scenario, but linked it instead to proximity to the classical pinch-line spin liquid at the CTP.

5.4.2 Phase diagram around parameters from Thompson *et al.* and comparison with nonlinear spin-wave theory

The significant enlargement of the E phase is absent in a linear spin-wave treatment [152]. However, a study using nonlinear spin-wave theory (NLSWT), which accounts for magnon interactions, did observe a similar effect [189]. They examined the phase diagram near the $\text{Yb}_2\text{Ti}_2\text{O}_7$ parameters reported by Thompson *et al.* [178], finding that NLSWT breaks down in a broad region of the T_{1-} phase close to the E boundary. While they could not definitively identify the ground state in this unstable region, they reported indications of E order—consistent with our pf-FRG results.

To allow a direct comparison, we repeated our pf-FRG analysis for the exact parameter set considered in Ref. [189]. These couplings are not expressed in the J_1, \dots, J_4 parametrization, but instead in the *dual* global frame with $\bar{J}_H, \bar{K}, \bar{D}, \bar{\Gamma}$, as introduced after Eq. (5.11). Details of the conversion are provided in Appendix B. In this frame, the classical $T_{1-} \rightarrow E$ phase boundary

Figure 5.8 –
Quantum phase diagram near $\text{Yb}_2\text{Ti}_2\text{O}_7$ parameters from Thompson *et al.* obtained from pf-FRG. The couplings are given in terms of the *dual* parameters introduced after (5.11). The region shown corresponds exactly to the one studied with NLSWT in Ref. [189], with $(\bar{K} - \bar{\Gamma})/\bar{J}_1 = -0.096$ fixed so that the $\text{Yb}_2\text{Ti}_2\text{O}_7$ parameters from Thompson *et al.* [178] fall in the plane (star ‘a’). Stars ‘b’–‘d’ mark parameter sets from other experimental studies, which do not lie exactly within the plane shown (see Table 5.1).



lies at $(\bar{K} + \bar{\Gamma})/J_H = 0$, while the Thompson *et al.* parameters are $(\bar{K} + \bar{\Gamma})/J_H = 0.11$, $(\bar{K} - \bar{\Gamma})/J_H = -0.096$, and $D/J_H = -0.525$. Our pf-FRG phase diagram for fixed $(\bar{K} - \bar{\Gamma})/J_H$ in this region is shown in Fig. 5.8. We observe the same key features as before: the *E* phase expands significantly, placing the $\text{Yb}_2\text{Ti}_2\text{O}_7$ parameters inside this phase, and a small paramagnetic region emerges. Importantly, the area where pf-FRG predicts *E* order but classical theory favors T_{1-} aligns almost perfectly with the unstable region identified by NLSWT (see Fig. 3 in Ref. [189]). This close agreement supports that the boundary shift is a genuine quantum effect rather than an artifact of the pf-FRG.

5.4.3 Challenges of parameter fitting in $\text{Yb}_2\text{Ti}_2\text{O}_7$

There have been several studies determining the coupling constants that best describe $\text{Yb}_2\text{Ti}_2\text{O}_7$ (see Table 5.1), which is mostly achieved by fitting linear spin-wave theory spectra to neutron scattering data of the dynamical structure factor at high magnetic fields, where the sample is polarized, magnon branches are most sharp, and classical spin-wave theory is predicted to be accurate. Notably, the four examples for the so obtained couplings given in Table 5.1 all place the compound in the FM T_{1-} phase of the classical phase diagram, but in close vicinity to the boundary to the *E* phase. We now argue, however, that this is also somewhat by design.

In Ref. [179], it was shown that the fit to high-field spin-wave spectra is underconstraint—equally good fits can be obtained along a one-dimensional line in parameter space. To further constrain the fit, the authors additionally matched a gap in the spin-wave spectrum *at zero magnetic field* to classical calculations. They found that, classically, this gap closes precisely at the FM-AFM transition. Since the experimentally observed gap is very small (around 0.11 meV), this procedure necessarily forces the parameters to lie near that phase boundary.

A similar strategy was employed in Ref. [191], which also reported an underconstraint fit when relying only on the inelastic high-field spectrum. There, refinement came from matching rods of scattering in the structure factor that appear just above the ordering transition, again *at zero magnetic field*, with results from classical Monte Carlo. In the classical model, these rods appear exactly at the AFM–FM phase boundary [see Fig. 5.2(e)], again constraining the parameters to this region.

If, as our calculations suggest, quantum fluctuations significantly affect the zero-field phase diagram and strongly shift the phase boundary, then these parameter estimates will no longer be accurate, since comparisons with classical zero-field calculations are not appropriate. Only

methods that explicitly incorporate quantum fluctuations are capable of reliably calculating correlations in the zero-field limit. At present, however, pf-FRG cannot access real-frequency spectra for anisotropic models (see discussion in Sec. 4.6), and to our knowledge no existing method can yet do so reliably for frustrated three-dimensional quantum magnets. A resolution of this problem must therefore be postponed to future developments.

Another possible explanation for the disagreement with experiment is that E order only dominates at very low temperatures. At small but finite temperatures, the combination of quantum and thermal fluctuations may instead favor the T_{1-} phase. An observation consistent with this scenario appears in Fig. 5.4(b), which shows the flow of the order-parameter susceptibility for the parameters from Scheie *et al.*: at higher cutoffs $\Lambda/|J| > 0.3$, the T_{1-} susceptibility dominates, while dominant E order occurs only at lower cutoffs (and within a mean-field interpretation, the RG cutoff can be regarded as an effective temperature [147].)

5.5 Discussion

Our pf-FRG study demonstrates that quantum fluctuations strongly reshape the zero-temperature phase diagram of Kramers rare-earth pyrochlores, beyond effects captured by linear spin-wave theory. Most notably, we find a pronounced enlargement of the E -phase, leading to substantial shifts of the phase boundaries consistent with other numerical approaches [186, 189], as well as the emergence of extended paramagnetic regimes whose precise nature could not yet be conclusively determined.

On the theoretical side, this has important consequences for the fate of classical higher-rank spin liquids, which always occur at points of maximal degeneracy. At the classical triple point, where the pinch-line spin liquid was previously identified [181], quantum order-by-disorder selects E order. Yet, at a shifted “quantum triple point” inside an emergent paramagnetic regime, we observe correlations that closely mirror those of the classical pinch-line spin liquid, suggesting that remnants of the higher-rank gauge structure may persist in the quantum model, possibly even forming an extended phase. It would be worthwhile to extend our analysis to the full *Atlas of Classical Pyrochlore Spin Liquids* derived in Ref. [26] for the model also considered here, which features a wide variety of higher-rank spin liquids. In addition, a related classical state has been predicted on the breathing pyrochlore lattice with DM interactions [192], a system that could be directly studied with our existing pf-FRG implementation.

We also find evidence for a small but finite quantum spin-nematic regime near the $T_{1-} - T_2$ boundary—a rare example of nematic correlations in a three-dimensional $S = 1/2$ system without an applied field. Clarifying the true nature of these exotic phases is an exciting direction for future work.

On the materials side, the boundary shifts have direct consequences for compounds close to competing orders. For $\text{Yb}_2\text{Ti}_2\text{O}_7$, our analysis places existing parameter estimates deep in the E phase, inconsistent with the experimentally established T_{1-} ground state. This suggests that parameter fitting procedures, which currently rely on classical modeling, may need to be revisited. Alternatively, the interplay of thermal and quantum fluctuations at small but finite temperatures could stabilize the observed order at small but finite temperatures. This could, e.g., be investigated with the recently developed pm-FRG discussed in Sec. 4.6. In addition, the mixed $T_{1-} \oplus E$ correlations we find may account for the “correlated paramagnet” observed experimentally just above the ordering transition, as well as the mixed dynamical correlations reported in the low-temperature phase.

It would be interesting to extend the analysis performed in this chapter to other materials

near phase boundaries. A natural candidate is $\text{Er}_2\text{Sn}_2\text{O}_7$, which has a T_2 -ordered ground state but is believed to lie very close to the boundary with the E phase [152, 193].

Chapter 6

Unconventional quantum states in maple-leaf magnets

Quantum spin models on two-dimensional lattices built from triangular motifs have long been regarded as prime candidates for realizing exotic quantum paramagnetic ground states. Two canonical examples are the triangular lattice, composed of corner-sharing triangles, and the kagome lattice, composed of edge-sharing triangles. The triangular J_1 – J_2 Heisenberg antiferromagnet is widely believed to host spin-liquid phases [147, 194], while the nearest-neighbor kagome Heisenberg antiferromagnet has emerged as the prototypical case where geometric frustration alone stabilizes a QSL ground state [195, 196].

This motivates the study of spin models on other two-dimensional lattices built from triangular motifs, such as the star, bounce, and trillium lattices. Realizations of these in actual materials, however, are exceedingly rare. The maple-leaf lattice, illustrated in Fig. 6.1, is more promising, as it is realized in several Cu^{2+} -based minerals with quantum spins. To name a few examples:

- spangolite $\text{Cu}_6\text{Al}(\text{SO}_4)(\text{OH})_{12}\text{Cl}\cdot 3\text{H}_2\text{O}$ [197, 198], bluebellite $\text{Cu}_6\text{IO}_3(\text{OH})_{10}\text{Cl}$ [199–202],

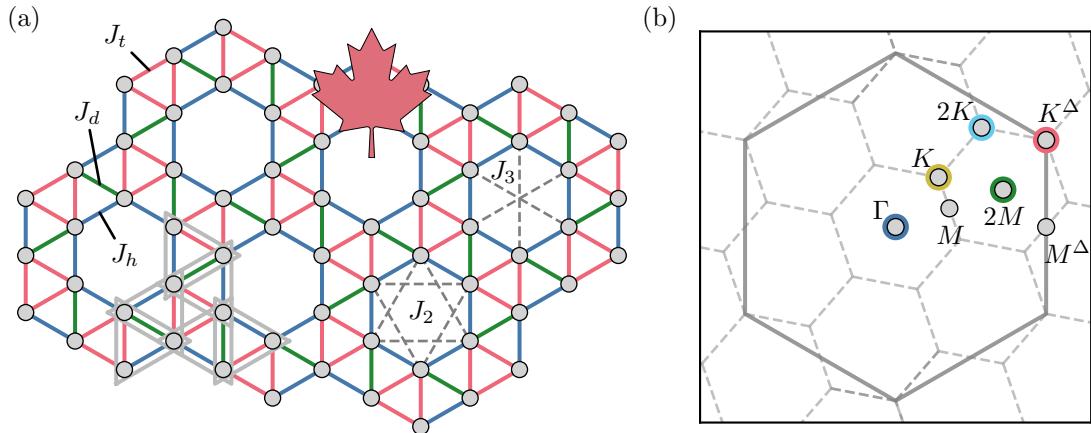


Figure 6.1 – The maple-leaf lattice in real and momentum space. (a) Real-space structure of the maple-leaf lattice, showing the three symmetry-inequivalent nearest-neighbor couplings J_t , J_d , and J_h in different colors. Dashed lines mark additional cross-plaquette interactions J_2 and J_3 that we study in Sec. 6.4. The gray triangular outlines show a triangle decomposition used for the derivation of the exact dimer singlet (DS) ground state. A full definition of the lattice given in Appendix C. (b) Momentum space structure. The dashed line outlines the first Brillouin zone of the maple-leaf lattice, periodically repeated. The solid line indicates the extended Brillouin zone, defined as the Brillouin zone of the underlying triangular lattice obtained by adding sites at the centers of the hexagonal plaquettes of the maple-leaf lattice. Dots show several high-symmetry points. The color-highlighted momenta are the allowed momenta for eighteen site clusters we employ in our CMFT calculation.

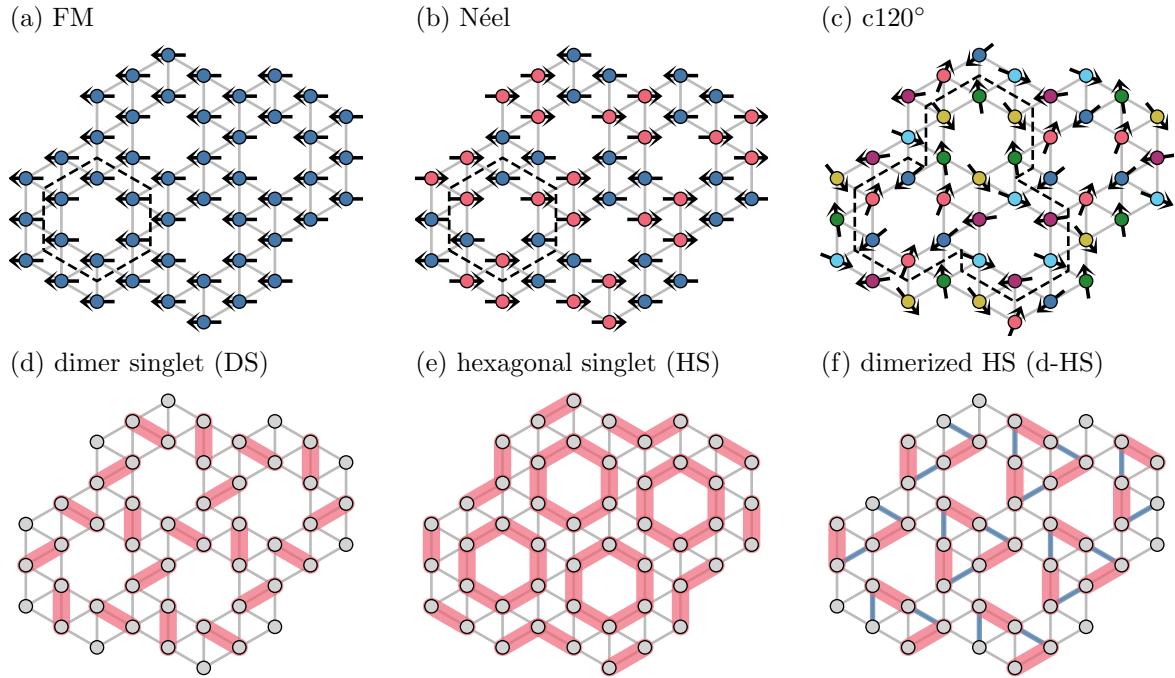


Figure 6.2 – Ground states of the nearest-neighbor model of the maple-leaf lattice (a)-(c) Magnetically ordered states. The magnetic sublattices are highlighted by different colors, and the magnetic unit cells are outlined by dashed black lines. The $c120^\circ$ has an eighteen-site magnetic unit cell and six magnetic sublattices. (d)-(f) Singlet ground states on the maple-leaf lattice. Red bonds highlight the singlet pattern. In the hexagonal singlet, all spins of the hexagon reside in the combined singlet sector. In the dimerized singlet, the strong ferromagnetic correlations are highlighted by blue bonds.

- mojaveite $\text{Cu}_6\text{TeO}_4(\text{OH})\text{Cl}$ [199], fuettererite $\text{Pb}_3\text{Cu}_6\text{TeO}_6(\text{OH})_7\text{Cl}_5$ [203], sabelliite $(\text{Cu},\text{Zn})_2\text{Zn}[(\text{As},\text{Sb})\text{O}_4](\text{OH})_3$ [204].

While the second group await detailed magnetic characterization, spangolite and bluebellite have already been shown to be well described by $S = 1/2$ Heisenberg models with mixed ferro- and antiferromagnetic interactions on five symmetry-inequivalent nearest-neighbor bonds of the crystal lattice [198, 202]. Among them, bluebellite exhibits a magnetically ordered ground state [200], whereas spangolite shows no long-range order down to $T \approx 8$ K and instead likely hosts a correlated dimer ground state [197, 198].

From a theoretical perspective, the maple-leaf lattice has a coordination number $z = 5$ and can be obtained from the triangular lattice by a $\frac{1}{7}$ site depletion. In this sense, it lies between the triangular lattice ($z = 6$, no depletion) and the kagome lattice ($z = 4$, $\frac{1}{4}$ depletion), and can therefore be expected to display strong frustration effects. The nearest-neighbor antiferromagnet on the maple-leaf lattice is believed to host a magnetically ordered $c120^\circ$ state [30, 205] [illustrated Fig. 6.2(c)]. However, the strongly reduced sublattice magnetization found in this state already points to significant quantum fluctuations. Together with the widely mixed couplings present in real material realizations, this motivates a broader exploration of the nearest-neighbor phase diagram in search of unconventional quantum paramagnetic phases. For simplicity, in the following we restrict ourselves to the three symmetry-inequivalent nearest-neighbor couplings intrinsic to the maple-leaf lattice: J_h (hexagon bonds), J_t (triangle bonds), and J_d (dimer bonds), as illustrated in Fig. 6.1(a). The corresponding Hamiltonian becomes

$$H = J_h \sum_{\langle ij \rangle_h} \mathbf{S}_i \mathbf{S}_j + J_d \sum_{\langle ij \rangle_d} \mathbf{S}_i \mathbf{S}_j + J_t \sum_{\langle ij \rangle_t} \mathbf{S}_i \mathbf{S}_j, \quad (6.1)$$

Indeed, moving away from the isotropic antiferromagnetic point $J_h = J_t = J_d$ by increasing J_d , it was recently discovered that the ground state becomes a nonmagnetic valence-bond solid (VBS), where all sites connected by the dimer bonds J_d form exact singlets, as illustrated in Fig. 6.2(d). Remarkably, this simple dimer product state can be shown to be the *exact* ground state for $J_d > 2$ [30]. Together with the well-known dimer state of the Shastry–Sutherland model [206], this is the only model on two-dimensional lattices with uniform tilings for which an exact dimer ground state is known [30]. Dimerized phases also appear in two other limiting cases of the nearest-neighbor model. First, for vanishing $J_t = J_d = 0$ but antiferromagnetic $J_h > 0$, the maple-leaf lattice reduces to decoupled hexagons, whose ground state lies in the singlet sector of the total spin of each hexagon, referred to as a *hexagonal singlet* and illustrated in Fig. 6.2(e). Second, for antiferromagnetic $J_h > |J_t|$ and large ferromagnetic $J_d \rightarrow -\infty$, the ground state is argued (see below) to be another distinct VBS[31], consisting of singlet dimers on half of the hexagonal bonds [see Fig. 6.2(f)] and strong ferromagnetic correlations along J_d . This state additionally breaks the rotational symmetry of the maple-leaf lattice. A triplon mean-field analysis predicts that both states appear in extended phases in the phase diagram for $J_h > 1$ and $J_d, J_t > 0$ [31]. We provide a more detailed explanation of the different ground states of the nearest-neighbor model in Sec. 6.1.

All of these VBS states are expected to occur in the vicinity of antiferromagnetic phases—either the Néel state or the $c120^\circ$ state [Fig. 6.2(b, c)]. The close competition between AFM and VBS states is known to potentially give rise to unconventional critical behavior. A prominent example is the Shastry–Sutherland model, where the continuous transition between a plaquette VBS and a Néel state has been argued to proceed via a *deconfined quantum critical point* (DQCP)[207]. Such continuous transitions between phases breaking different symmetries lie outside the Ginzburg–Landau paradigm. In the DQCP scenario, the critical point is described by a field theory with emergent gauge fluctuations and fractionalized quasiparticles [208, 209], similar to those in a quantum spin liquid. The precise phase diagram of the Shastry–Sutherland model remains debated: while some studies support a direct DQCP, more recent work suggests a narrow intervening spin-liquid regime that ends at a nearby DQCP accessed by tuning further-neighbor couplings [210, 211]. A closely related scenario has also been proposed for the square-lattice J_1 – J_2 Heisenberg model, where the transition from a Néel state to a VBS may involve an intermediate spin-liquid phase [111, 212], although the existence of this regime is still under active debate [110].

The phase boundaries between AFM and dimer states are therefore promising regions to search for exotic physics associated with deconfined fractionalized quasiparticles. This motivates a numerical investigation of the parameter regions that may host such states, which provides the main motivation for the first part of this chapter.

We begin in Sec. 6.1 by discussing the possible ordered states and, more importantly, the valence-bond solid (VBS) phases that arise in the nearest-neighbor model on the maple-leaf lattice.

In Sec. 6.2 (based on Ref. [P5]), we investigate the quantum phase diagram using the pf-FRG in a parameter region that contains both an AFM and the exact DS state, but the nature of the transition between these phases is not clear. To be precise, we consider $J_h = J_t = 1$ and $J_d \in [0, 4]$, where $c120^\circ$ order is expected at small J_d and the DS state becomes the exact ground state for $J_d > 2J_h$ [30]. Our results suggest the presence of a narrow quantum-paramagnetic

regime separating these two phases.

Motivated by the possible occurrence of the HS and d-HS states for antiferromagnetic $J_h > 0$ but ferromagnetic $J_d < 0$, in Sec. 6.3 (based on Ref. [U1]) we explore a broader phase diagram with mixed couplings: antiferromagnetic $J_h > 0$ and ferromagnetic $J_t, J_d < 0$. As mentioned above, such mixed couplings also appear in spangolite and bluebellite, albeit with different signs, raising the possibility that other materials may lie close to this parameter regime. To obtain a qualitative picture of the phase diagram, we combine three complementary methods: the Luttinger–Tisza approach (LT, Sec. 3.1), cluster mean-field theory (CMFT, Sec. 3.3), and pf-FRG. Our results provide evidence that the HS and d-HS states indeed form extended phases adjacent to the ordered $c120^\circ$ and N’eel states, consistent with Ref. [31]. Moreover, we identify a large quantum-paramagnetic region with correlations distinct from the VBS phases, which may host a variety of exotic states—including a regime with strong spin-nematic response.

In Sec. 6.4 (based on Ref. [P6]), we pursue a different route toward exotic quantum phases on the maple-leaf lattice—namely, putative chiral QSLs. A two-dimensional chiral spin liquid is a magnetically disordered phase that breaks time-reversal symmetry, exhibits topological order with fractionalized excitations, and supports chiral edge modes akin to those in fractional quantum Hall states [4, 213]. It has been shown that such states can emerge by “quantum melting” of classical *noncoplanar* spin orders, which are characterized by a finite scalar spin chirality $\mathbf{S}_1 \cdot (\mathbf{S}_2 \times \mathbf{S}_3)$ of three neighboring spins. If quantum fluctuations are enhanced—for example, by reducing the spin from large S to $S = 1/2$ or by tuning exchange couplings to more frustrated regimes—long-range magnetic order may be destroyed while the chiral symmetry breaking persists, giving rise to a chiral spin liquid [214, 215]. Spin models whose classical ($S \rightarrow \infty$) limit hosts noncoplanar order are thus promising candidates for such phases at $S = 1/2$. Historically, in conventional spin models such three-dimensional spin textures have been thought to require complex nondiagonal or multi-spin interactions, or the inclusion of magnetic fields [216–222]. More recently, it has been established that also conventional Heisenberg models can realize noncoplanar orders by including competing long-range interactions. On both the kagome lattice [223, 224] and the square-kagome lattice [225], further-neighbor cross-plaquette couplings have been shown to stabilize noncoplanar “cuboc” orders in classical Heisenberg models. Moreover, in the kagome lattice such interaction are believed to indeed stabilize a chiral spin liquid [226–229].

This motivates us to introduce analogous cross-plaquette interactions J_2 and J_3 on the maple-leaf lattice, as illustrated in Fig. 6.1. Exploring both the classical and quantum phase diagrams for $J_1 = \pm 1$ with varying J_2 and J_3 , we identify regions where the classical model stabilizes noncoplanar ground states and pf-FRG indicates a quantum-paramagnetic regime—making these regions promising candidates for chiral spin liquid ground states. Importantly, in the Kalmeyer–Laughlin paradigm, a chiral spin liquid breaks both time-reversal and reflection symmetries while preserving their product [213]. Since the maple-leaf lattice lacks reflection symmetry altogether, a chiral spin liquid on this lattice could lie outside the conventional Kalmeyer–Laughlin framework.

6.1 Ordered and valence-bond solid states in the nearest-neighbor model

Before analyzing the phase diagram of the nearest-neighbor Hamiltonian, we briefly summarize the ordered and quantum paramagnetic states that appear in the parameter regime of interest. Their real-space structures are illustrated in Fig. 6.2.

6.1.1 Ordered states

Canted 120° order. Around the isotropic AFM point $J_h = J_t = J_d > 0$, the classical ground state realizes a *canted* 120° configuration (c120°). In this state, spins on the red triangles coupled by J_t form an ideal 120° pattern, while the spins on neighboring triangles are canted by an angle θ that depends on the ratio of the competing couplings J_h and J_t [30]. This state has an eighteen-site magnetic unit cell comprising six distinct spin sublattices. At the fully isotropic point, it reduces to the conventional 120° state.

Néel order. For dominant antiferromagnetic couplings on the hexagons (J_h), the ground state adopts a Néel configuration: spins on each hexagon align antiparallel to their nearest neighbors, while the spins on the triangles are all parallel. The magnetic unit cell coincides with the six-site geometric unit cell, resulting in two magnetic sublattices.

6.1.2 Valence-bond solids

Dimer singlet For dominant $J_d > 0$, the ground state becomes a dimer product state in which each J_d bond forms an exact dimer singlet (DS)

$$|\psi_d\rangle = \bigotimes_{\langle ij \rangle_d} |S_{ij} = 0\rangle, \quad (6.2)$$

where $|S_{ij} = 0\rangle$ is the conventional singlet state on the bond between site i and j . Since neither the J_h nor the J_t terms renormalize this state, it is an eigenstate of the Hamiltonian and the exact ground state for $J_h = J_t > 0$ and $J_d/J_h > 2$. Following Ref. [30], this can be seen by rewriting the Hamiltonian (6.1) as a sum over single triangle Hamiltonians

$$H = \sum_t h_t \quad h_t = J_h \mathbf{S}_1 \mathbf{S}_2 + \frac{J_d}{2} \mathbf{S}_2 \mathbf{S}_3 + J_t \mathbf{S}_3 \mathbf{S}_1, \quad (6.3)$$

where each triangle contains exactly one bond of each type. The green J_d bonds appear in two neighboring triangles (yielding the factor 1/2), while all other bonds appear only once. An example for the unit cell of such a triangle decomposition is given by the gray outlines in Fig. 6.1(a). The ground-state energy of a single triangle h_t is a lower bound for the ground-state energy per site E/N of the full Hamiltonian, as it minimizes each component of the triangle decomposition. For the DS state, the energy per site is simply the singlet energy divided by two,

$$E_d/N = -\frac{J_d}{2} \frac{3}{4}, \quad (6.4)$$

since the number of J_d dimers is half the number of sites. For $J_h = J_t > 0$ and $J_d/J_h > 2$, this energy matches the lower bound $e_\Delta = E_d/N$ (which can be shown by exactly diagonalizing h_t), proving that $|\psi_d\rangle$ is indeed the exact ground state. This state does not break any lattice symmetries.

Hexagonal singlet. For $J_h > 0$ and $J_d = J_t = 0$, the model decouples into isolated hexagons with nearest neighbors coupled antiferromagnetically. The full ground state is therefore a product state,

$$|\psi_{\text{HS}}\rangle = \bigotimes_h |\psi_h^0\rangle, \quad (6.5)$$

where $|\psi_h^0\rangle$ denotes the unique ground state of the single-hexagon Hamiltonian

$$H_h = J_h \sum_{\langle ij \rangle \in \text{hex}} \mathbf{S}_i \cdot \mathbf{S}_j. \quad (6.6)$$

This ground state is a unique total singlet with eigenvalue $S_h(S_h + 1) = 0$ of the total spin operator

$$\mathbf{S}_h^2 = \left(\sum_{i \in \text{hex}} \mathbf{S}_i \right)^2. \quad (6.7)$$

It preserves all lattice symmetries. Using triplon mean-field theory, Ref. [31] further argued that this state (or a phase continuously connected to it) remains stable in a finite region for small ferromagnetic J_d and J_t , consistent with our later findings.

Dimerized hexagonal singlet. The dimerized hexagonal singlet (d-HS), illustrated in Fig. 6.2, is a state in which dimers form only on one subset of the J_h bonds (hexagonal or vertical), thereby breaking the C_6 symmetry of the maple-leaf lattice. Following Ref. [31], this state can be argued to be the ground state in the limit $J_d \rightarrow -\infty$ with $J_h > |J_t| > 0$. In this regime, the strong ferromagnetic J_d coupling projects the two spins on each J_d bond into the triplet sector with $S_d = 1$, so that the system reduces to effective spin-1 moments interacting antiferromagnetically via J_h on a kagome lattice. The spin-1 kagome antiferromagnet is known to form a trimerized singlet ground state that breaks the symmetry between up and down triangles [230, 231]. Analogous to the AKLT construction for the spin-1 chain [232], this trimerized state can be represented in terms of composite $S = 1/2$ moments forming singlets, which in the present case coincide with the original microscopic spins of the model—leading precisely to the d-HS pattern. Triplon mean-field theory further predicts that this state is stabilized over a finite parameter region with $J_d < J_t < 0$ and $J_h > 0$.

6.2 Candidate quantum disordered intermediate phase in the Heisenberg antiferromagnet

We now examine in detail the transition from the $c120^\circ$ state to the exact DS state. The results presented in this section are based on Ref. [P5], for which the author of this thesis generated all data. All figures shown here are redrawn from that work.

Throughout this section, we parametrize the Hamiltonian by a single anisotropy parameter α , defined as

$$J_d = 2\alpha J_h = 2\alpha J_t > 0. \quad (6.8)$$

As discussed above, Ref. [30] proved that the DS state is the exact ground state for $\alpha \geq 1$, and in fact remains an eigenstate of the Hamiltonian at all couplings. Using DMRG, they further argued that the DS is likely the ground state already for $\alpha > \alpha_c^2$ with $\alpha_c^2 \approx 0.675$. They additionally showed that the classical $c120^\circ$ state has lower energy than the DS for $\alpha_b = \frac{1+\sqrt{7}}{8} \approx 0.456$, which provides a lower bound for the DS regime. It has remained unclear, however, whether these two phases meet directly in a phase transition, or whether an intermediate quantum phase intervenes.

To address this question, we perform large-scale pf-FRG simulations with lattice truncations up to $L = 18$ (corresponding to 822 correlated sites, see Sec. 4.4.2) and a frequency grid of

$35 \times 40 \times 40$. The resulting phase diagram is shown schematically in Fig. 6.3(a). Our analysis shows that $c120^\circ$ order melts into a quantum paramagnetic (PM) regime at $\alpha_c^1 \approx 0.67$, while the DS phase emerges for $\alpha > \alpha_c^2 \approx 0.8$ —a larger value than the values reported by DMRG. This leaves an intermediate window $\alpha_c^1 < \alpha < \alpha_c^2$, where a quantum paramagnetic phase persists, but whose nature we cannot determine conclusively. In the discussion section we compare these findings with other advanced numerical approaches and consider possible interpretations of this regime.

Before this, we first describe how the quantum phase diagram and the two critical couplings are determined, starting with the $c120^\circ \rightarrow$ PM transition at α_c^1 and then the onset of the DS phase at α_c^2 .

6.2.1 Melting of magnetic order

In order to study the transition from the ordered $c120^\circ$ state to the PM regime, we analyze the flow of the structure factor at the dominant momentum $\mathbf{k}^{\max} = \mathbf{K}^\Delta$ [as defined in Fig. 6.1(b)]. Using the flow-breakdown criterion described in Sec. 4.5.1, we extract the critical scale as a function of α , shown in Fig. 6.4(a) together with examples of the flow in Fig. 6.4(d). For small $\alpha < 0.5$, the flow displays a hump that becomes more pronounced with L , a clear signature of $c120^\circ$ order. As α increases, this feature gradually weakens until the flow becomes smooth and nearly L -independent, indicating a paramagnetic regime [Fig. 6.3]. We find no evidence of a flow breakdown for $\alpha > \alpha_c^1 \approx 0.67$, indicating the onset of the paramagnetic regime.

Additional evidence comes from the momentum-resolved structure factor: its peak broadens as α increases, eventually becoming featureless deep in the DS phase [Fig. 6.3(c)]. To quantify this broadening and check that the extent of the PM regime is not overestimated, we follow Refs. [233–235] and compute the correlation ratio R defined as

$$R = 1 - \chi(\mathbf{K} + \boldsymbol{\delta})/\chi(\mathbf{K}), \quad (6.9)$$

where $\boldsymbol{\delta} = \frac{2\pi}{L} (1/\sqrt{3}, -1)^\top$ is the shortest reciprocal lattice vector of the triangular lattice underlying the maple-leaf lattice, scaled by the maximal correlation length \tilde{L} (in real-space units) permitted by a bond truncation length L . The result is shown in Fig. 6.3(b). The correlation ratio tends to $R = 1$ in the ordered phase, as the Bragg peak becomes increasingly sharp in the thermodynamic limit, and to $R \rightarrow 0$ in the PM phase. The initial decrease of R coincides reasonably well with α_{c1} obtained from the flow-breakdown analysis. Unlike the standard expectation for continuous transitions [233–235], we do not observe a clear crossing of R curves for different L . The $L = 12$ and $L = 15$ curves do intersect, but well below α_c^1 (for $L = 18$ we obtained data only very close to the estimated transition due to long computation times), which suggests that the extent of the PM regime has not been strongly overestimated. We note, however, that R is evaluated at fixed $\Lambda/|J| = 0.015$, a scale lying below the flow breakdown in the ordered phase ($\alpha < \alpha_c^1$), where the numerical data are unreliable—likely preventing the appearance of a clean crossing point anyway.

6.2.2 Onset of dimer singlet order

To further characterize paramagnetic regime, we calculate nearest-neighbor correlations in the low-cutoff limit on the dimer (J_d), triangle (J_t), and hexagon (J_h) bonds, denoted by $\chi_{d/t/h}$ [Fig. 6.4(a)]. Within the nonmagnetic regime, these correlations are essentially independent of system size, while in the ordered regime finite size effects are strong, as expected. For large α , χ_t and χ_h tend to zero, while χ_d saturates at a constant negative value. This behavior indicates

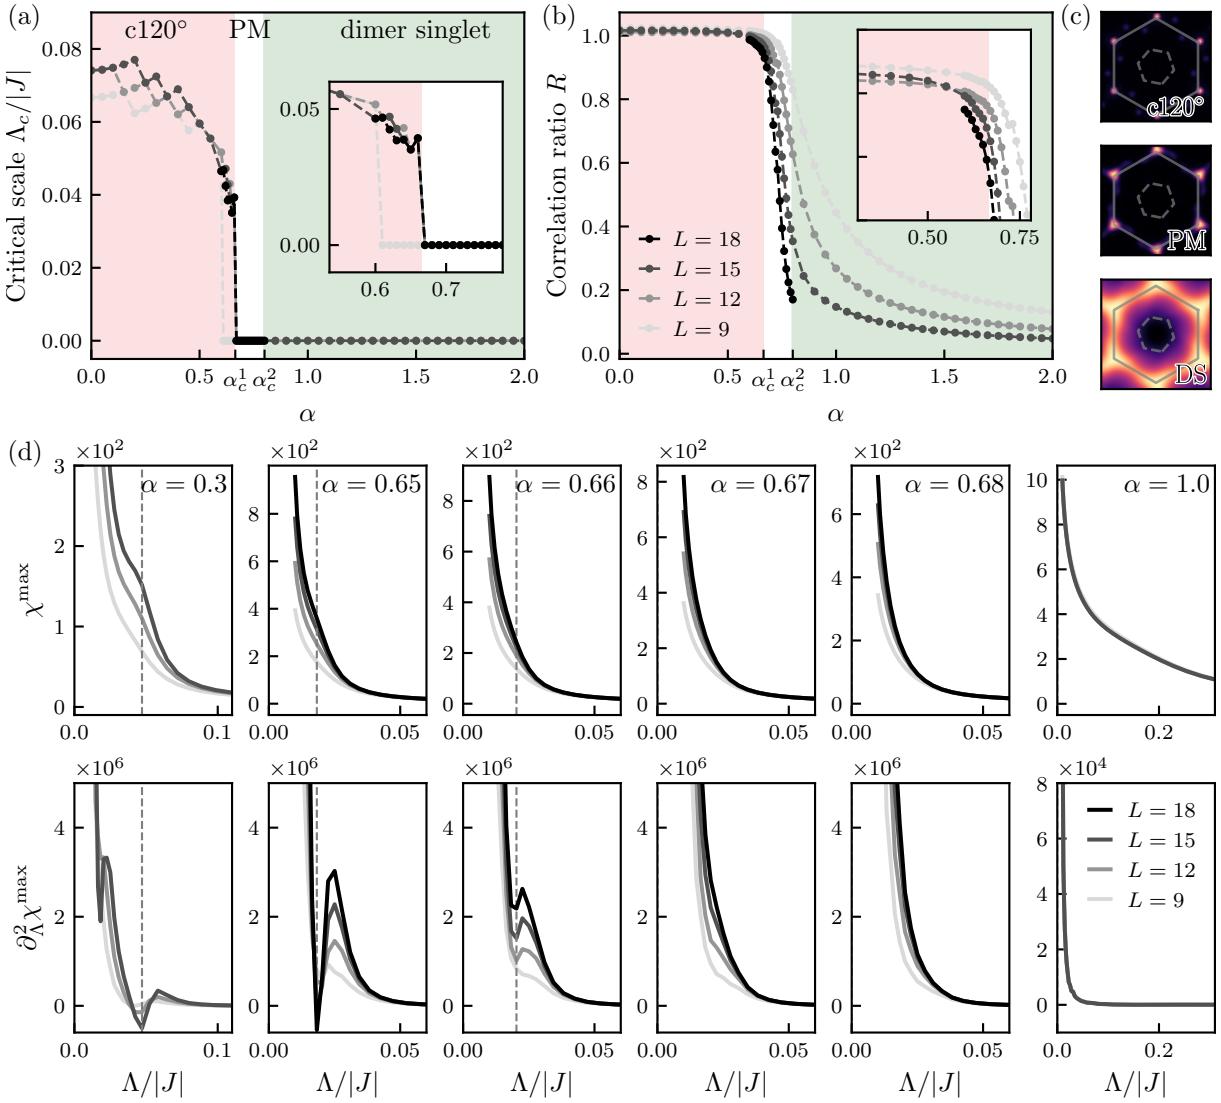


Figure 6.3 – Transition from the $c120^\circ$ phase to the paramagnetic regime. (a) Critical scale Λ_c as a function of α , obtained by identifying nonmonotonic features in the second derivative of the structure-factor flow at $\mathbf{k}^{\max} = \mathbf{K}^\Delta$ [illustrated in (d), where Λ_c is marked by dashed lines]. For $\alpha > \alpha_c^1 \approx 0.67$ the flow shows no breakdown, indicating the absence of conventional magnetic order, while for $\alpha > \alpha_c^2 \approx 0.80$ the system transitions into the DS state [see Fig. 6.4]. (b) Correlation ratio quantifying the broadening of the structure-factor peak. (c) Structure factors representative of the three regimes (shown for $\alpha = 0.3, 0.73$, and 2.0 from top to bottom).

the formation of the DS state. A similar evolution of spin correlations was previously observed in pf-FRG studies of the Shastry–Sutherland model [236], which also realizes an exact dimer ground state. Consistently, the momentum-resolved structure factor at large α shows only very broad features.

Fig. 6.4(b) illustrates the onset of dimer correlations χ_d as a function of Λ , which become increasingly steep with decreasing cutoff. Following Ref. [236], we assume that in the thermodynamic limit the transition into the DS phase involves the immediate formation of the exact product state, and therefore a rapid saturation of spin correlations on the J_d bonds. Under this assumption, we determine the transition point $\alpha_c^2(\Lambda)$ from the intersection of two linear fits [black

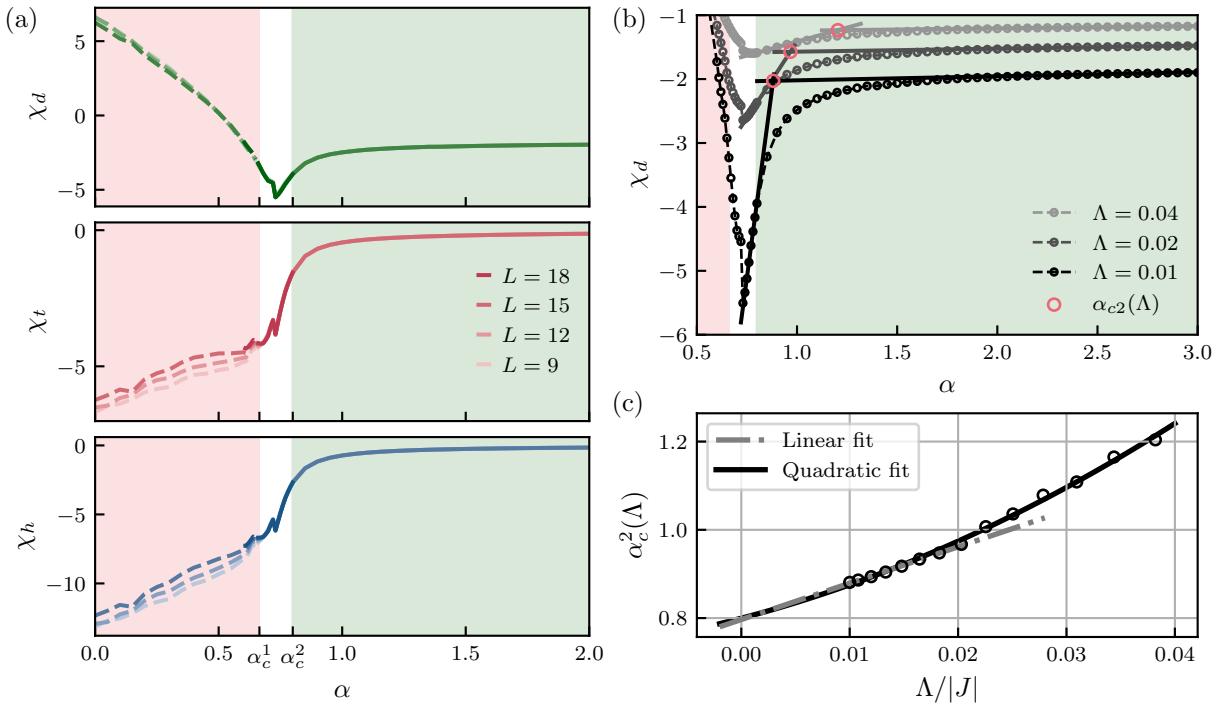


Figure 6.4 – Transition into the dimer singlet phase. (a) Nearest-neighbor correlations on the J_d , J_t , and J_h bonds at $\Lambda = 0.01|J|$. For large α , the correlations χ_d saturate to a constant value while χ_t and χ_h vanish, consistent with a dimer singlet (DS) ground state. A dependence on lattice size appears only in the long-range ordered c120° phase (dashed lines), where the RG flow breaks down at a finite $\Lambda_c > 0.01|J|$. (b) Evolution of χ_d at different RG scales Λ for fixed $L = 15$. The transition into the DS state is determined from the intersection of two linear regressions (solid black lines) at each Λ . Extrapolating the resulting $\alpha_c^2(\Lambda)$ to $\Lambda \rightarrow 0$ as in (c) gives a critical value of $\alpha_c^2 \approx 0.8$, above which the DS state is predicted to be the ground state.

lines in Fig. 6.4(a)] and extrapolate to $\Lambda \rightarrow 0$, obtaining $\alpha_c^2 \approx 0.80$ [Fig. 6.4(b, c)]. Our analysis therefore suggests that the DS crystal is the ground state for $\alpha > \alpha_c^2 \approx 0.8$. The assumption of instantaneous saturation of the correlations is reasonable since the DS is always an eigenstate of the Hamiltonian. However, if the DS belongs to a larger degenerate ground-state manifold, this assumption may no longer hold and a more gradual transition is possible.

6.2.3 Discussion

Our analysis predicts the suppression of magnetic order above $\alpha_c^1 \approx 0.67$ and the onset of the exact DS phase at $\alpha_c^2 \approx 0.8$, leaving open an intermediate paramagnetic regime $\alpha_c^1 \lesssim \alpha \lesssim \alpha_c^2$. In this window, the structure factor retains similarities to the c120° state, but with softened peaks [Fig. 6.3(c)], while the dimer correlations χ_d remain well below their saturated value in the DS phase. This strongly suggests that the system is not yet in the exact DS state. Since all couplings in this regime are of comparable strength, the HS and d-HS states introduced earlier are also unlikely candidates. This leaves three main possibilities: (i) the regime is not a distinct phase but a correlated version of the DS state smoothly connected to the exact product state, (ii) it hosts a symmetric quantum spin liquid (QSL), or (iii) it realizes another type of quantum paramagnet. In support of scenario (ii), Ref. [237] proposed a U(1) QSL with a qualitatively similar structure factor (albeit with broader peaks), identified as UC10 within a

projective symmetry group classification.

Although the phase diagram has been examined with a range of numerical methods, including more recent works appearing after Ref. [P5], the results are contradictory, and no consensus has yet emerged.

The earliest investigation [238], using the coupled cluster method and ED, reported magnetic order up to $\alpha_c^2 \approx 0.725$ and proposed a direct, likely first-order, transition into the DS phase. The DMRG study of Ref. [30] obtained a slightly lower transition point of $\alpha_c^2 \approx 0.675$, though it did not provide a value for α_c^1 or clarify the nature of the transition. More recently, Ref. [239] found $\alpha_c^1 \approx 0.7095$ using iDMRG and $\alpha_c^1 \approx 0.615$ using neural quantum states (NQS). For the DS transition they reported $\alpha_c^2 \approx 0.732$ (iDMRG) and $\alpha_c^2 \approx 0.7095$ (NQS). While an intermediate PM regime cannot be excluded, they argue for a direct transition, though without definitive evidence.

Adding yet another perspective, Ref. [240] employed large-scale tensor network simulations with iPEPS and proposed a completely different scenario: they reported no symmetry breaking for $0 < \alpha < 0.65$, suggesting a QSL ground state that undergoes a quantum phase transition from gapless to gapped at $\alpha \approx 0.25$. This would imply that even the isotropic AFM point of the maple-leaf lattice realizes a QSL ground state. In their study, the DS phase emerges at $\alpha \approx 0.725$, again lower than our pf-FRG estimate, with indications of a possible intervening magnetically ordered phase between the QSL and DS regimes.

In summary, the precise phase diagram of the maple-leaf model—and even the nature of the ground state at the isotropic AFM point $J_d = J_t = J_h > 0$ —remains unsettled. This underscores both the richness of the model and the formidable challenges it poses for current state-of-the-art numerical methods in two dimensions, making it a compelling problem for future investigations.

6.3 Quantum states in the ferro–antiferromagnetic Heisenberg model

Motivated by the possible emergence of extended VBS phases reported in Ref. [31], and by the presence of mixed interactions in the minerals spangolite [198] and blueelite [200, 202], we now turn to the extended phase diagram of the nearest-neighbor maple-leaf model with antiferromagnetic $J_h > 0$ but ferromagnetic $J_t, J_d > 0$. We explore this phase diagram using a combination of methods. We begin with a classical analysis via Luttinger–Tisza (LT), then study the quantum phase diagram and the emergence of VBS phases using CMFT on three distinct cluster geometries, and finally apply pf-FRG to examine both the global phase structure and the potential for spin-nematic order. Before discussing the results of each method in detail, we begin this section with a summary of the main findings obtained by combining them.

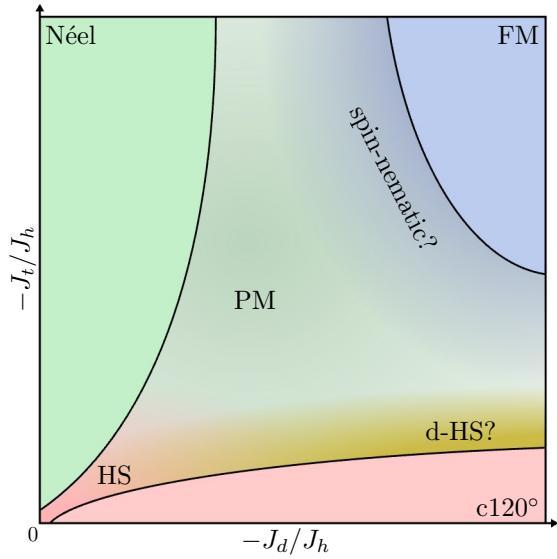
This section is based on Ref. [U1], from which all figures are redrawn; the numerical data underlying these figures was obtained by the author of this thesis.

6.3.1 Schematic phase diagram

Neither CMFT nor pf-FRG provides quantitative error bars, as their approximations do not allow for rigorous uncertainty estimates. As a result, precise phase boundaries cannot be determined, and parts of the phase diagram remain open for further study. Still, the two methods complement each other in a useful way: CMFT is expected to typically favor ordered states, while pf-FRG tends to overestimate paramagnetic regions. Where both approaches agree, the conclusions are therefore considerably more reliable. By combining consistent results from LT, CMFT, and pf-FRG, we construct a schematic composite phase diagram, shown in Fig. 6.5.

Figure 6.5 –

Schematic phase diagram from LT, CMFT, and pf-FRG. We consider antiferromagnetic $J_h > 0$ with ferromagnetic $J_d, J_t \leq 0$. For large negative J_d and/or J_t , conventional Néel, FM, and c120° orders are stabilized in all methods. Between them lies a broad paramagnetic (PM) region without dipolar order. Both CMFT and pf-FRG indicate an extended hexagonal singlet (HS) phase around $J_d = J_t = 0$, and a smaller dimerized HS (d-HS) phase near the c120° boundary. pf-FRG further reveals strong spin-nematic tendencies at the PM-FM boundary and identifies a distinct PM region (green) with correlations unlike HS or d-HS, suggestive of an additional nonmagnetic phase such as a QSL or VBS. Since precise phase boundaries within the PM regime cannot be resolved, the locations of different phases are indicated schematically using color gradients.



At large FM couplings we consistently identify three ordered phases with Néel, c120°, and FM order. All of these are exactly captured by LT, and reproduced by both CMFT and pf-FRG although with shifts of the phase boundaries. The ordered phases enclose a broad paramagnetic regime at intermediate competing values of $-J_d/J_h$ and $-J_t/J_h$, where no conventional magnetic order is observed. Correlations within this regime vary strongly across parameter space, indicating the presence of several distinct paramagnetic phases.

Both CMFT and pf-FRG point to an extended hexagonal singlet (HS) phase around $J_d = J_t = 0$, as well as a dimerized HS (d-HS) phase near the c120° boundary, consistent with triplon mean-field theory [31]. Interestingly, pf-FRG finds signatures of the d-HS phase all along the boundary to the c120° region, while CMFT detects it only in a smaller subregion. Moreover, the extend of the FM regime is substantially reduced in pf-FRG and replaced by a paramagnetic region dominated by ferromagnetic correlations and a strong spin-nematic response, indicating possible spin-nematic order. In addition, both pf-FRG and CMFT identify another potentially distinct PM region (approximately located in the green PM region in Fig. 6.5) with only a weak nematic response and correlations different from either HS or d-HS, suggestive of an additional nonmagnetic phase such as a QSL or VBS.

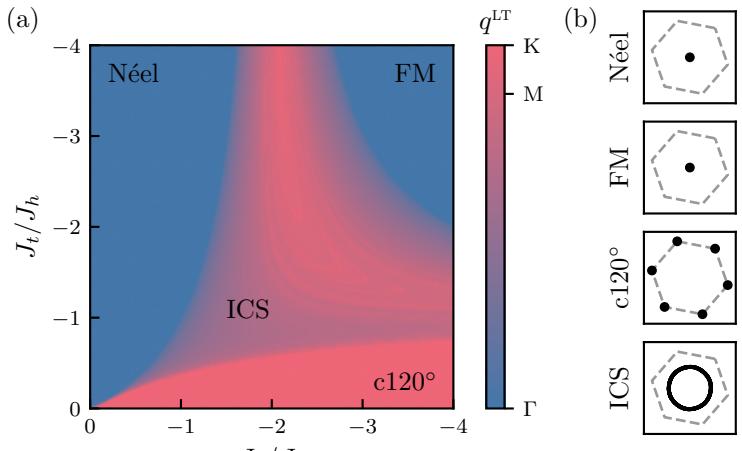
In the following, we present the results from LT, CMFT, and then pf-FRG in detail.

6.3.2 Classical phase diagram from Luttinger-Tisza

Using the Luttinger-Tisza method introduced in Sec. 3.1, we construct the classical phase diagram shown in Fig. 6.6. Here we display the magnitude of the \mathbf{q}^{LT} vectors. The Néel and FM phases have the same periodicity as the maple-leaf unit cell and thus correspond to $\mathbf{q}^{\text{LT}} = 0$ orders. The c120° order is characterized by $\mathbf{q}^{\text{LT}} = \mathbf{K}, \mathbf{K}'$ at the corners of the first Brillouin zone [see Fig. 6.1(b)]. All of these states satisfy the strong spin constraint within a single- \mathbf{q} spiral Ansatz and are therefore the exact classical ground states.

These ordered phases enclose a regime with a continuous manifold of incommensurate (ICS) wave vectors \mathbf{q}^{LT} , which shift smoothly across the ICS region [Fig. 6.6(a,b)]. In this regime, the strong spin constraint is not satisfied, and the exact classical ground state cannot be determined by LT. Possible candidates include incommensurate spiral order or, as suggested by the contin-

Figure 6.6 – Classical phase diagram from Luttinger–Tisza. (a) Magnitude $q^{\text{LT}} = |\mathbf{q}^{\text{LT}}|$ of the momenta with minimal LT eigenvalue. (b) Corresponding \mathbf{q}^{LT} vectors in the first Brillouin zone. In the Néel and FM phases $\mathbf{q}^{\text{min}} = \Gamma$, while in the $c120^\circ$ phase $\mathbf{q}^{\text{min}} = \mathbf{K}$. Between them, incommensurate (ICS) momenta interpolate continuously between Γ and \mathbf{K} . In all phases except ICS the hard spin-length constraint is fulfilled, whereas in the ICS phase only the soft constraint holds, leaving the classical ground state undetermined.



uous \mathbf{q}^{LT} manifold, a classical spin-liquid regime. As we show below, it is precisely in this ICS region that pf-FRG predicts a nonmagnetic phase, with structure factors displaying the same type of continuous features seen in LT.

6.3.3 Cluster mean-field theory

We now describe the application of CMFT, as formulated in Sec. 3.3, and discuss the resulting phase diagrams. The large unit cell of the maple-leaf lattice imposes strong constraints on the clusters that can be used, since periodic boundary conditions require cluster sizes that tile the full lattice (i.e., multiples of the unit cell). To capture the $c120^\circ$ order, cluster sizes of $N_C = 18$ are required, which also corresponds to the largest system size we can treat numerically. Moreover, recall that inside paramagnetic phases all mean fields vanish, so CMFT reduces to ED with periodic boundary conditions. To describe symmetry-breaking quantum paramagnetic phases, the chosen cluster must itself allow for the corresponding symmetry reduction. These considerations restrict us to the three clusters shown in Fig. 6.7, which accommodate distinct symmetry-breaking patterns of singlet states:

- Cluster (a) preserves all lattice symmetries and is therefore suited to describe the HS state.
- Cluster (b) exhibits the C_3 symmetry characteristic of the d-HS state, making it the natural choice for capturing this phase.
- Cluster (c) has only C_3 symmetry about the centers of triangles but treats a larger number of J_h couplings exactly, biasing it more strongly towards the HS state than cluster (a).

As a first step, we distinguish magnetically ordered from paramagnetic states by computing the average magnetization

$$m_{\text{avg}} = \frac{1}{N_C} \sum_{i \in C} |\langle \mathbf{S}_i \rangle|, \quad (6.10)$$

where the sum runs over all N_C sites of the cluster. This order parameter vanishes only when all local magnetizations are zero and saturates at $m_{\text{avg}} = 1/2$ for simple product states. The background color in Fig. 6.7 displays m_{avg} for the three clusters, and shows all putative phase boundaries or crossovers indicating qualitative changes in the ground state.

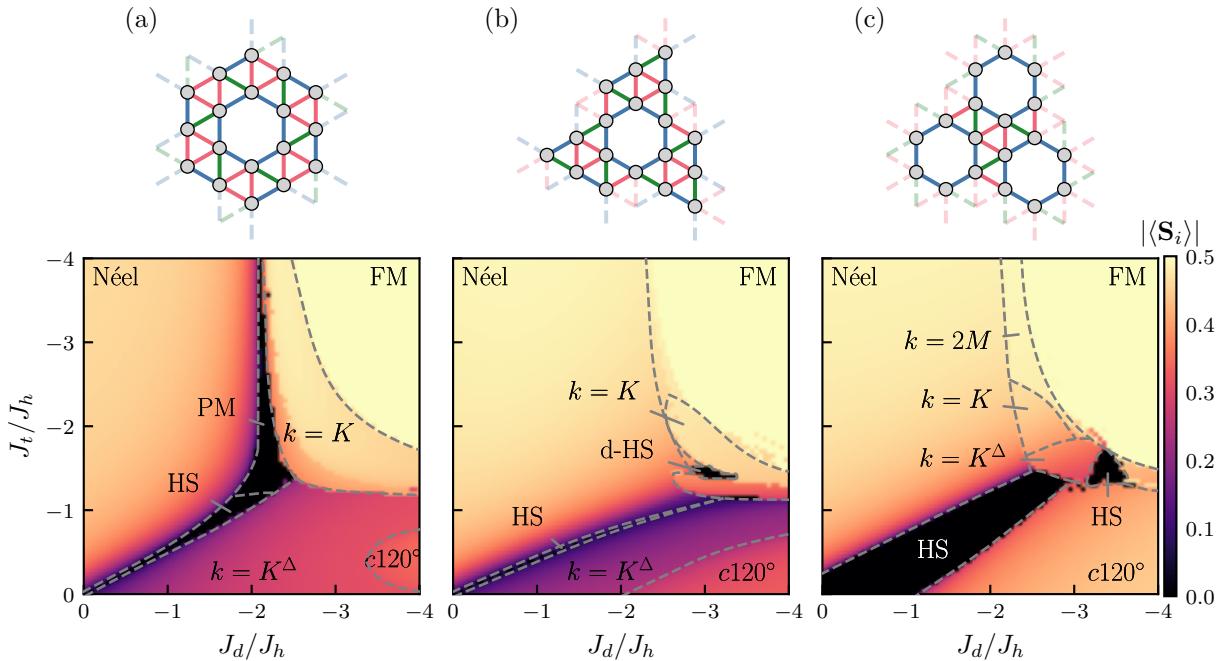


Figure 6.7 – Quantum phase diagrams from CMFT for three different clusters. Top row: the three eighteen-site clusters used in the analysis. Solid lines denote interactions treated exactly, while dashed lines represent interactions approximated by mean-field decoupling under periodic boundary conditions. Bottom row: corresponding phase diagrams (for $J_h < 0$). Dashed lines indicate qualitative changes in the ground state, corresponding either to crossovers or genuine phase transitions. The background color indicates the average local magnetization $|\langle \mathbf{S}_i \rangle|$; black regions indicate paramagnetic phases with $|\langle \mathbf{S}_i \rangle| = 0$. An extended HS phase appears in all clusters. Only cluster (b) preserves the symmetries necessary to host the d-HS state, which indeed emerges as a small but distinct region in its phase diagram. Cluster (a) further exhibits a paramagnetic (PM) region that resembles neither the HS nor d-HS states. In ordered regions where neither Néel, FM, or $c120^\circ$ order is not fully realized, regions are labeled by the momentum \mathbf{k}^{\max} at which the structure factor is maximal.

All clusters exhibit an extended PM region (black), though of different sizes, as well as large magnetically ordered regions. Within these, we also find states distinct from the three conventional ordered phases introduced above; these are labeled by the momentum at which the structure factor is maximal. Comparison with LT and pf-FRG suggests, however, that such additional ordered phases are likely finite-size artifacts, and ICS or disordered phases are more likely. In the following, we discuss in more detail how ordered and paramagnetic regions are characterized and how these tentative boundaries are identified using appropriate order parameters and observables.

Characterizing ordered states To characterize collinear FM and Néel phases, we define the order parameters

$$m_{\text{FM}} = \frac{1}{N_c} \left| \sum_{i \in \mathcal{C}} \langle \mathbf{S}_i \rangle \right|, \quad (6.11)$$

$$m_{\text{Néel}} = \frac{1}{N_c} \left| \sum_{i \in \mathcal{C}} (-1)^i \langle \mathbf{S}_i \rangle \right|, \quad (6.12)$$

where the site index i is chosen such that $(-1)^i$ reproduces the staggered spin pattern of the Néel state. For a pure product state of the corresponding order, both order parameters reach their maximum of 1/2.

Defining an order parameter for the $c120^\circ$ state in terms of staggered magnetization is more involved, since the relative angle between spins on distinct red triangles  depend on both the couplings and the cluster geometry. Instead, we use the fact that in the $c120^\circ$ phase each red triangle locally realizes 120° order, which is naturally captured by the vector chirality

$$\kappa^\Delta = \frac{4}{3\sqrt{3}N_t} \sum_{i,j,k \in \text{red triangles}} |\langle \mathbf{S}_i \times \mathbf{S}_j \rangle + \langle \mathbf{S}_j \times \mathbf{S}_k \rangle + \langle \mathbf{S}_k \times \mathbf{S}_i \rangle|, \quad (6.13)$$

where the sum runs over all red J_t triangles, (i, j, k) are the three sites in this triangle in a counterclockwise order, and N_t denotes their total number in the cluster. With this normalization, κ^Δ reaches 0.5 for perfect $c120^\circ$ order in a product state with $|\langle \mathbf{S}_i \rangle| = 1/2$.

In the CMFT phase diagrams we also encounter regions where none of the above order parameters is finite, or where several coexist, suggesting alternative types of magnetic order. To analyze these cases we compute the spin structure factor

$$S(\mathbf{k}) = \frac{1}{N} \sum_{i,j} e^{i\mathbf{k} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \langle \mathbf{S}_i \cdot \mathbf{S}_j \rangle, \quad (6.14)$$

and identify the momentum \mathbf{k}^{\max} where $S(\mathbf{k})$ is maximal. On the eighteen-site cluster only five symmetry-inequivalent momenta are allowed [highlighted in Fig. 6.1(b)].

The results for the order parameters and \mathbf{k}^{\max} for all three clusters are shown in Fig. 6.8. The three conventional magnetic orders—Néel, FM, and $c120^\circ$ —appear in every cluster, though their extent varies noticeably. This is expected, since different bonds are mean-field approximated in each cluster, biasing them toward different ordering patterns.

All clusters also exhibit regions where none or multiple order parameters are finite. Here, the ordering wave vector jump discontinuously between physically allowed momenta as $(J_d/J_h, J_t/J_h)$ are varied. Our Luttinger–Tisza analysis, on the other hand, shows \mathbf{q}^{sc} vectors at ICS momenta, suggesting that the finite-size cluster likely can't capture the correct order in this regime. This interpretation is supported by pf-FRG, which also reveals dominant ICS correlations in the same parameter regime.

Characterizing paramagnetic phases The d-HS state reduces the C_6 lattice symmetry to a C_3 rotation about the hexagon centers of the maple-leaf lattice. To quantify this symmetry breaking, we define the corresponding order parameter as

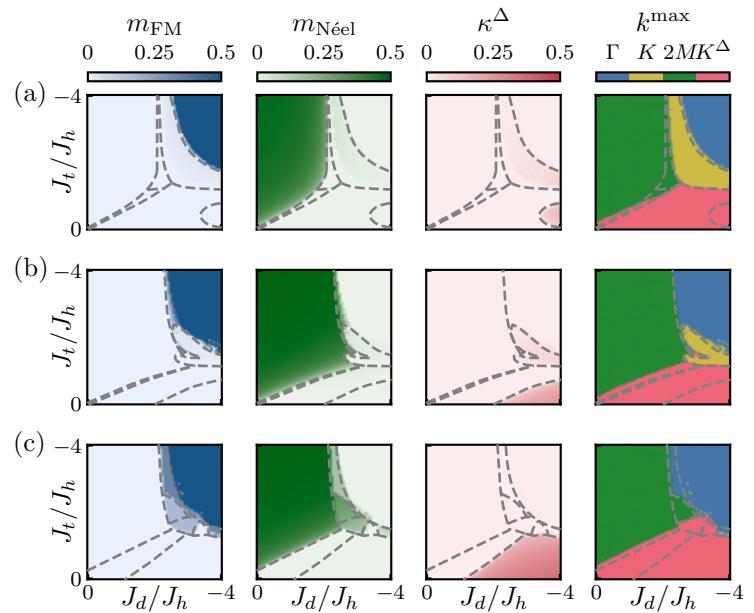
$$O_{\text{d-HS}} = \frac{1}{6N_h} \left| \sum_{\langle i,j \rangle \in \text{hexagons}} \mathbf{S}_i \cdot \mathbf{S}_j - \sum_{\langle i,j \rangle \in \text{hexagons}} \mathbf{S}_i \cdot \mathbf{S}_j \right|, \quad (6.15)$$

where the sums run over the bonds of all N_h fully connected hexagons in the clusters. A paramagnetic state with $m_{\text{avg}} = 0$ and finite $\langle O_{\text{d-HS}} \rangle$ is identified as a d-HS state. As explained above, only cluster (b) in Fig. 6.7 can capture a PM phase with this reduced C_3 symmetry.

By contrast, the HS state preserves all lattice symmetries and therefore cannot be captured by a conventional order parameter. Instead, we evaluate its presence through the overlap (fidelity) with the exact HS state, using the projector

$$P_{\text{HS}} = |\psi_{\text{HS}}\rangle \langle \psi_{\text{HS}}|, \quad (6.16)$$

Figure 6.8 –
CMFT observables characterizing ordered phases. Rows (a)–(c) correspond to the clusters shown in Fig. 6.7. The first three columns display the order parameters for the FM, Néel, and 120° phases. When no single order parameter is finite, the phase is instead identified by the momentum \mathbf{k}^{\max} at which the structure factor is maximal (last column). The symmetry-inequivalent momenta of the eighteen-site clusters are highlighted in Fig. 6.1(b), which also illustrates that both the HS and d-HS states correspond to $\mathbf{k}^{\max} = \mathbf{K}^\Delta$ (red).



where $|\psi_{\text{HS}}\rangle$ is the exact HS state defined above. A value of $\langle P_{\text{HS}} \rangle = 1$ signals that the ground state on the fully connected hexagons exactly corresponds to this state. More generally, $\langle P_{\text{HS}} \rangle$ measures the proximity to the HS state, though it does not allow a precise determination of phase boundaries.

For both HS and d-HS phases we can additionally compute nearest-neighbor correlations, which should reveal strong AFM amplitudes on certain bonds characteristic of the singlet pattern as illustrated in Fig. 6.2.

The evolution of these observables across the PM regime is shown in Fig. 6.9. The HS phase is consistently identified on all three clusters, marked by pronounced antiferromagnetic correlations on the hexagons and large $\langle P_{\text{HS}} \rangle$. The singlet formation observed on certain peripheral bonds in Fig. 6.9(e) is an artifact of CMFT: with all mean fields vanishing, each peripheral site couples via $J_h > 0$ to only one neighbor and thus forms a perfect singlet. Our analysis suggests the HS phase likely extends beyond the $J_t = J_d = 0$ point, consistent with the triplon mean-field analysis of Ref. [31].

The C_6 symmetry breaking in the d-HS phase can be resolved only on cluster (b), where it stabilizes over a relatively narrow region, approximately $J_d/J_h \in [-2.9, -3.3]$ and $J_t/J_h \in [-1.4, -1.5]$ as shown in Fig. 6.9(d). Its real-space spin-spin correlations, shown in Fig. 6.9(f), display the expected dimerized pattern: strong singlet amplitudes on ζ_3 bonds and strong ferromagnetic correlations on the J_d bonds. Within CMFT the HS and d-HS states are separated by magnetically ordered phases. Our pf-FRG calculations, presented in the next section, instead suggest that this separation may be a finite-size artifact, and that the two states are continuously connected in the thermodynamic limit.

Cluster (a) exhibits an additional PM regime that appears continuously connected to the HS phase, but with only small $\langle P_{\text{HS}} \rangle$, a characteristic momentum $\mathbf{k}^{\max} = 2\mathbf{M}$, and distinct real-space correlations [Fig. 6.9(g)]. The strong ferromagnetic (FM) correlations on the J_t bonds combined with the weaker antiferromagnetic (AFM) correlations on the J_h bonds set this regime apart from all other phases, suggesting that it either constitutes a distinct phase or, alternatively, a correlated extension of the HS state.

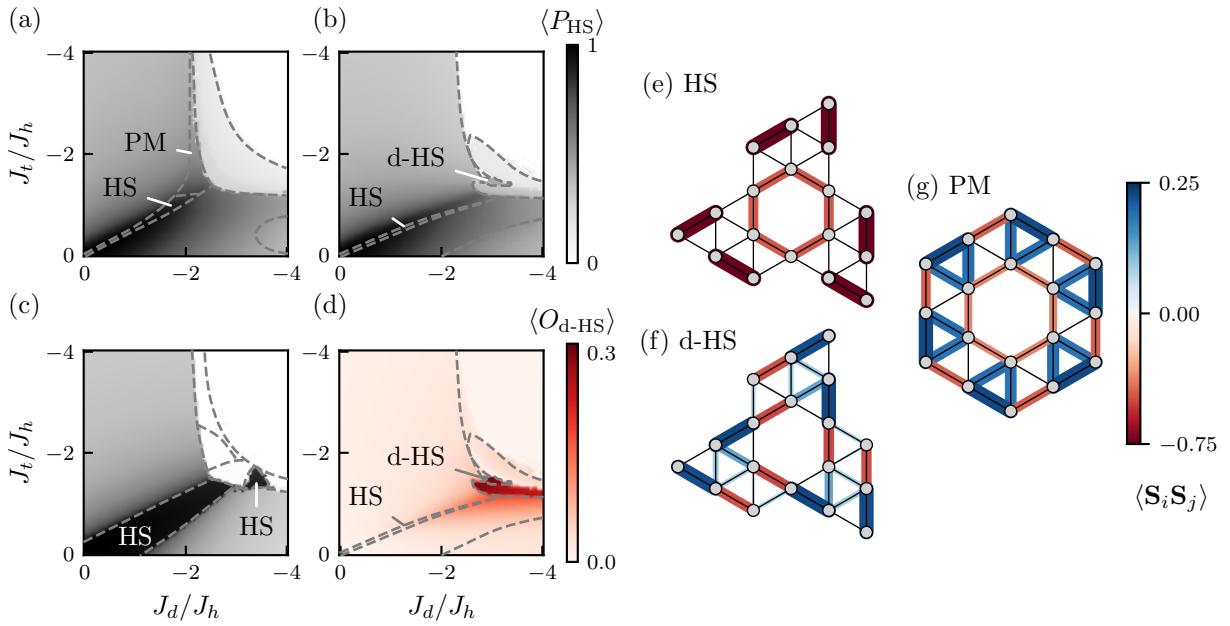


Figure 6.9 – CMFT observables for paramagnetic phases. (a)–(c) Expectation value of the projector P_{HS} onto the hexagonal singlet state for the three clusters shown in Fig. 6.7. (d) Order parameter of the d-HS state on cluster (b), which is the only cluster consistent with the required symmetries. In panel (a) the boundary between the region labeled “PM” and the HS phase is drawn at the points where $\langle P_{HS} \rangle \approx 0.8$, although the projector varies smoothly and never exhibits a sharp jump. This boundary should therefore be regarded only as a guide to the eye. (e)–(g) Nearest-neighbor isotropic equal-time spin correlations for the different paramagnetic regimes, represented simultaneously by color and line width.

6.3.4 Pseudo-fermion functional renormalization group

Finally, we employ the pf-FRG, which in principle provides the most reliable distinction between PM and ordered phases in the thermodynamic limit. Using the flow-breakdown criterion described in Sec. 4.5.1 (with representative flows shown in Fig. 4.2), we estimate the extent of the PM region and identify ordered phases through the momentum \mathbf{k}^{\max} of the structure-factor peaks. The resulting phase diagram, together with examples of structure factors and nearest-neighbor correlations, is shown in Fig. 6.10.

In addition to the three ordered phases (Néel, $c120^\circ$, and FM), the model hosts an extended paramagnetic regime where no flow breakdown is observed [indicated by circle markers and dashed lines in Fig. 6.10(a)]. Within this PM region, the spin structure factor typically shows broad intensity with soft maxima at wave vectors corresponding to nearby ordered states, and its profile evolves smoothly across parameter space [see Fig. 6.10(c)]. The nearest-neighbor spin correlations in Fig. 6.10(d) likewise exhibit qualitatively different patterns across the PM regime. Together, these findings suggest that the PM region is not uniform but instead consists of several distinct phases, which we analyze in detail below.

Spin nematic response A region of particular interest, distinct from both LT and pf-FRG expectations, is the PM sector close to the FM boundary. Here the PM regime arises because quantum fluctuations strongly suppress the extent of the classical FM phase, producing a “melted” FM state driven by competing AFM couplings. Such frustration of ferromagnetism has been argued to favor multipolar orders on a variety of lattices, including square [241–246], kagome [247],

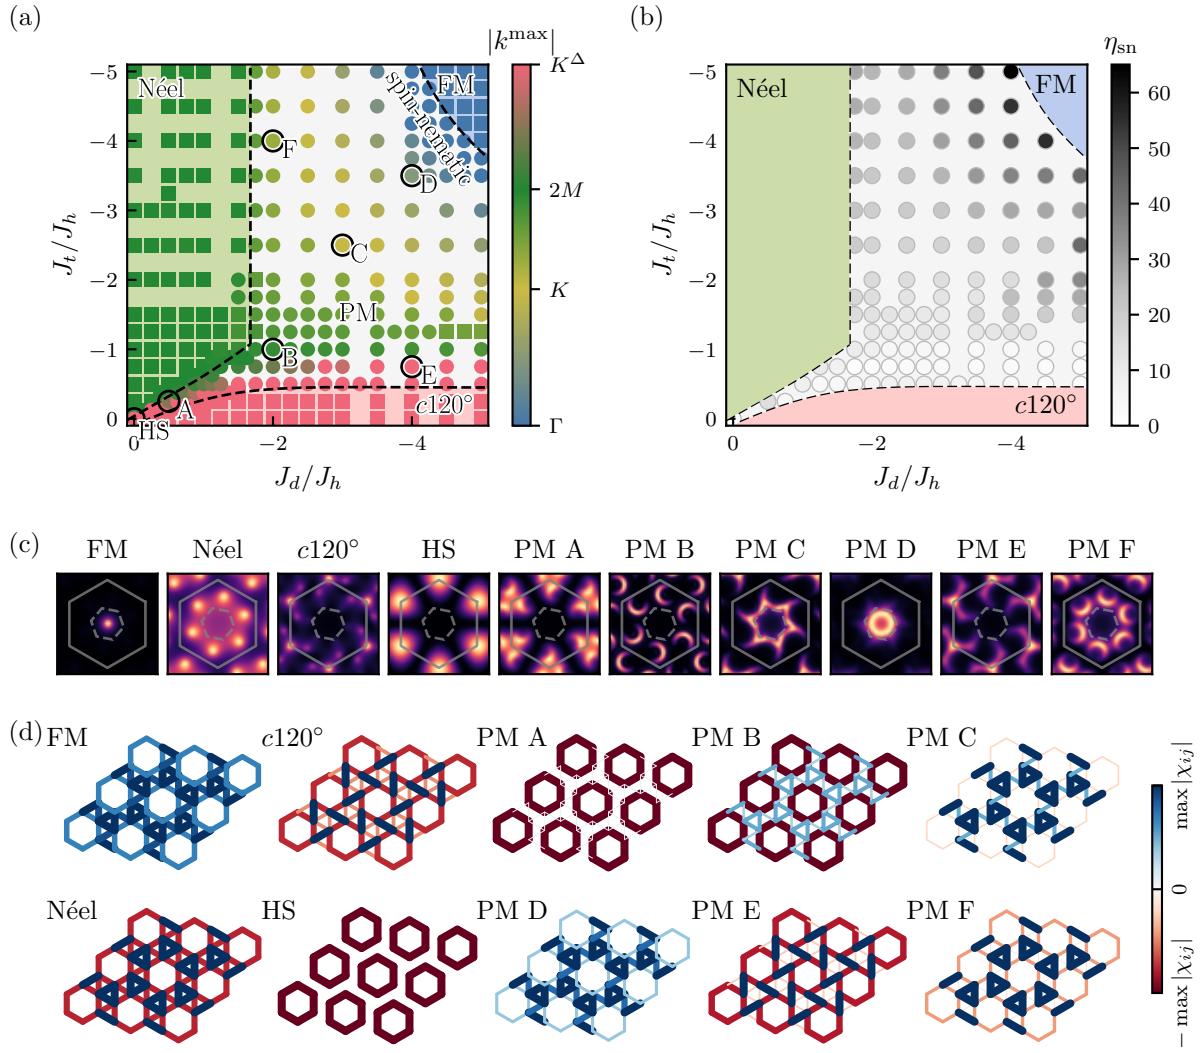


Figure 6.10 – Quantum phase diagram from pf-FRG. (a) Phase diagram as a function of the ferromagnetic couplings J_t and J_d (with $J_h > 0$). The color scale encodes the magnitude of the momentum \mathbf{k}^{\max} where the structure factor is maximal. Square markers denote couplings with a flow breakdown at a finite critical RG scale Λ_c , signaling conventional magnetic order, while circular markers indicate the absence of a breakdown, corresponding to paramagnetic (PM) states. Dashed lines serve as guides to the eye, separating ordered from paramagnetic regions. (b) Nematic response η_{SN} within the PM phase, highlighting the strongest tendency toward spin-nematic order near the FM boundary. (c) Structure factors in the ordered phases and at representative points within the PM regime: $J_d = J_t = 0.0$ (where the HS state is exact) and six example points (A–F) marked in (a). (d) Real-space nearest-neighbor spin correlations χ_{ij}^{zz} in the low-cutoff limit, normalized to the maximum value for each parameter point. Red (blue) bonds indicate AFM (FM) correlations.

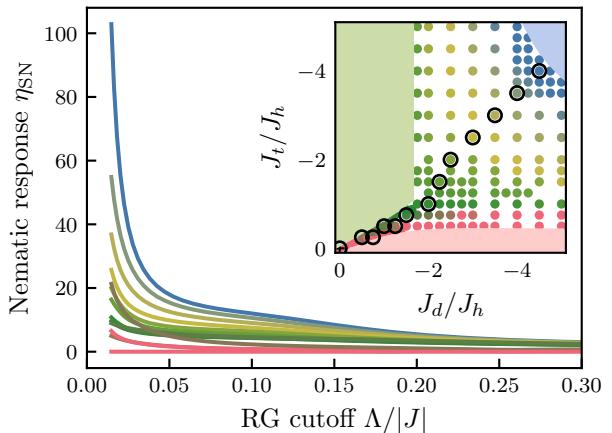
triangular [248], and body-centered cubic [249] lattices.

Our pf-FRG implementation allows us to probe for the tendency towards spin-nematic order by following the procedure outlined in Sec. 4.5.3. Concretely, we introduce a small anisotropy δ on the FM couplings J_d and J_t that breaks SU(2) symmetry down to U(1) as

$$\begin{aligned} J_d \mathbf{S}_i \cdot \mathbf{S}_j &\rightarrow (J_d - \delta)(S_i^x S_j^x + S_i^y S_j^y) + (J_d + \delta)S_i^z S_j^z, \\ J_t \mathbf{S}_i \cdot \mathbf{S}_j &\rightarrow (J_t - \delta)(S_i^x S_j^x + S_i^y S_j^y) + (J_t + \delta)S_i^z S_j^z. \end{aligned} \quad (6.17)$$

Figure 6.11 –

Nematic response from pf-FRG for selected points (marked by black circles in the inset) within the paramagnetic phase. The inset reproduces the phase diagram of Fig. 6.10(a). The color shows the magnitude of \mathbf{k}^{\max} where the structure factor is maximal. The response is strongest near the FM phase (blue) and weakest in the vicinity of the c120° phase (red).



To quantify the tendency towards this type of symmetry breaking—and hence towards spin-nematic order—we define the nearest neighbor correlations

$$\chi_{d/t}^{\mu\nu} = \chi_{ij}^{\mu\nu} \text{ with } (i, j) \in d/t, \quad (6.18)$$

as already studied in the previous section. From these, we construct the spin-nematic response

$$\eta_{\text{SN}}^{d/t} = \left| \frac{J_{d/t}}{\delta} \frac{\chi_{ij}^{xx} - \chi_{ij}^{zz}}{\chi_{ij}^{xx} + \chi_{ij}^{zz}} \right| \dots \quad (6.19)$$

For brevity, we suppress the explicit Λ -dependence. In practice, during the flow the denominator $\chi_{ij}^{xx} + \chi_{ij}^{zz}$ in the subdominant channel may cross zero, causing an artificial divergence of $\eta_{\text{SN}}^{d/t}$. To circumvent this, we consider the dominant nematic response defined as

$$\eta_{\text{SN}} = \begin{cases} \eta_{\text{SN}}^d & \text{if } \chi_d^{xx} > \chi_t^{xx}, \\ \eta_{\text{SN}}^t & \text{otherwise,} \end{cases} \quad (6.20)$$

which guarantees smooth behavior and avoids spurious singularities.

The evolution of η_{SN} in the low-cutoff limit is shown in Fig. 6.10(b), with representative flows provided in Fig. 6.11. As $\Lambda \rightarrow 0$, the nematic response peaks near the FM boundary and decreases steadily toward the Néel and c120° regions, indicating an enhanced tendency toward spin-nematic correlations in the vicinity of the FM phase. Although establishing true long-range nematic order would require four-spin susceptibilities—beyond the scope of the present pf-FRG framework—both the structure factor and the real-space correlations point to a distinct phase: the structure factor is confined to the first Brillouin zone and the correlations are purely ferromagnetic, in contrast to other PM regions that show additional spectral weight outside the first Brillouin zone together with AFM correlations on at least one bond type. This behavior is exemplified by point PM D in Fig. 6.10. Taken together, these observations suggest that PM D represents a distinct paramagnetic regime, plausibly of spin-nematic character.

VBS and putative QSL phases As with spin-nematic order, valence-bond solid (VBS) phases require access to four-spin susceptibilities, which lie beyond the scope of the current pf-FRG implementation. The HS state poses an additional challenge since it does not break any symmetries, rendering a response-function approach inapplicable. In principle, the d-HS state could be probed

by introducing a $C_6 \rightarrow C_3$ lattice anisotropy, but this would create two symmetry-inequivalent sites per unit cell and necessitate multiple self-energy components in the FRG flow—an extension that is computationally demanding and we have not yet implemented. Instead, we compare structure factors and real-space correlations with exact results and CMFT.

Both the HS and d-HS phases exhibit maximal spectral weight near the \mathbf{K}^Δ point in the structure factor. In pf-FRG this is observed for all red points in Fig. 6.10(a) (e.g., points labeled HS, PM A, and PM E). In this region the spin-nematic response is negligible, distinguishing it from the putative nematic regime. Still close to $J_d = J_t = 0$ (PM A), we find strong AFM correlations on J_h bonds and weak correlations elsewhere, consistent with an extended HS phase. Further away from this point but still near the $c120^\circ$ boundary (PM E), correlations show FM amplitudes on J_d and AFM on J_h , characteristic of the d-HS state, though without explicit dimerization (inaccessible to pf-FRG due to preserving C_6 symmetry). This suggests that the true ground state in this region is likely the d-HS. Interestingly, this would imply an extended d-HS regime directly connected to the HS phase, with no intervening ordered states as CMFT predicted. In this scenario, a direct transition from $c120^\circ$ order to d-HS order could occur, potentially giving rise to exotic deconfined critical behavior.

We further identify a PM region with weak spin-nematic response near the Néel boundary (green points in Fig. 6.10(a), e.g., PM B and PM F). Here the structure factor peaks near $2\mathbf{M}$, and real-space correlations show AFM J_h bonds alongside FM J_t and J_d bonds. This regime is distinct from both HS and d-HS. CMFT on cluster (a) finds a corresponding phase with weak HS overlap, consistent with our pf-FRG results. The associated structure factor exhibits half-moon patterns reminiscent of kagome VBS phases [250], suggesting either a larger-unit-cell VBS or, alternatively, a symmetric quantum spin liquid.

6.3.5 Discussion

By combining CMFT and pf-FRG we have established strong indications for a broad paramagnetic regime that most likely hosts several distinct phases, including putative VBS and spin-nematic ground states. The consistent identification of such a regime across both methods—despite their contrasting biases—provides compelling evidence that a genuine quantum paramagnet emerges in this part of the phase diagram. The precise nature of the underlying phases, as well as possible internal phase boundaries, however, cannot be resolved conclusively within our present approach. Of particular interest is the transition from the $c120^\circ$ to the putative symmetry-breaking d-HS phase, which raises the possibility of unconventional deconfined critical behavior. Determining whether this transition is direct and, if so, whether it is first-order or continuous, remains an open question of considerable interest.

This motivates further investigation with complementary techniques. Variational tensor-network methods such as iDMRG or iPEPS, or variational Monte Carlo with competing VBS and QSL Ansätze, would allow for a more direct energetic comparison between candidate states. Equally, computing dynamical quantities such as spin and dimer susceptibilities could help distinguish between spin-nematic and VBS tendencies. On the experimental side, characteristic fingerprints in the structure factor—such as broad continua or “half-moon” features—may serve as useful probes in candidate materials.

More broadly, our results show that competing ferro-antiferromagnetic interactions can give rise to unconventional quantum states beyond the standard paradigm of frustrated antiferromagnets. The maple-leaf lattice thus joins triangular, kagome, and Shastry–Sutherland systems as an arena where competing interactions stabilize exotic paramagnetic ground states. This highlights the potential of extending the search for novel spin-liquid and multipolar phases to other

frustrated lattices with mixed FM and AFM couplings, as well as to real material realizations.

It would be particularly interesting to extend the present CMFT and pf-FRG analysis to the parameter regimes relevant for the $S = 1/2$ models proposed for spangolite [197, 198] and bluebellite [199, 202] in recent studies, as well as, in the future, to the many not-yet-characterized candidate materials introduced in the introduction of this section.

6.4 Noncoplanar orders and putative chiral quantum spin liquids from cross-plaquette interactions

Motivated by the observation that long-range cross-plaquette interactions can stabilize noncoplanar orders in classical Heisenberg models on the kagome [223, 224] and square-kagome lattices [225]—both promising candidates for chiral QSLs in the $S = 1/2$ limit—we now extend our study beyond the nearest-neighbor Hamiltonian considered in the previous sections. Specifically, we take isotropic nearest-neighbor couplings $J_1 = J_d = J_h = J_t$ and add second- (J_2) and third-nearest neighbor (J_3) interactions across the hexagonal plaquettes of the maple-leaf lattice [see Fig. 6.1(a)]. The resulting Hamiltonian reads

$$H = J_1 \sum_{\langle ij \rangle_1} \mathbf{S}_i \cdot \mathbf{S}_j + J_2 \sum_{\langle ij \rangle_{h2}} \mathbf{S}_i \cdot \mathbf{S}_j + J_3 \sum_{\langle ij \rangle_{h3}} \mathbf{S}_i \cdot \mathbf{S}_j, \quad (6.21)$$

where $\langle ij \rangle_{h2}$ and $\langle ij \rangle_{h3}$ denote the cross-plaquette bonds.

This section is based on Ref. [P6], which maps out the classical and quantum ($S = 1/2$) phase diagrams for both AFM ($J_1 > 0$) and FM ($J_1 < 0$) cases. The classical phase diagram is obtained through a combination of Luttinger-Tisza, classical Monte Carlo, and a semianalytical approach introduced in Ref. [225], which allows for the determination of exact classical phase boundaries. These classical calculations were carried out by Martin Gembé and Heinz-Jürgen Schmidt. Here, we first briefly summarize the methods used to identify the classical phases and boundaries, with particular emphasis on the noncoplanar orders, before turning to the effects of quantum fluctuations for AFM $J_1 > 0$ and FM $J_1 < 0$. The quantum phase diagrams were obtained using pf-FRG by the author of this thesis, revealing several paramagnetic regimes—some overlapping with classical noncoplanar phases—that emerge as natural candidates for quantum spin liquid ground states.

All figures in this section are redrawn from Ref. [P6]. Most of the data shown was obtained by the author of this thesis. Exceptions are the classical phase boundaries, as well as all common origin plots and classical structure factors, which were created by Martin Gembé and Heinz-Jürgen Schmidt and reused here with permission. This section is supplemented by Appendix C, which presents phase diagrams obtained from unconstrained Luttinger-Tisza calculations, along with cuts through the quantum phase diagrams shown in the main text.

6.4.1 Methods for constructing the classical phase diagram

The starting point for constructing the classical phase diagram is the unconstrained Luttinger-Tisza (LT) method, where all \mathbf{q}^{LT} vectors minimizing the Hamiltonian under the weak spin-length constraint are calculated. The resulting phase diagrams are shown in Appendix C, already providing good indications of phase boundaries. Since we are particularly interested in noncoplanar ground states, a single- \mathbf{q} spiral Ansatz built from LT eigenvectors is insufficient. Instead, one must employ either the real-space version of LT (typically the simplest approach) or multi- \mathbf{q} Ansätze (both described in Sec. 3.1), which allow the identification of multiple coplanar and

noncoplanar phases. Nevertheless, there remain regions where LT does not produce spin configurations that satisfy the strong spin-length constraint, requiring alternative methods to determine the true ground state.

To this end, we employ classical Monte Carlo simulations (carried out by Martin Gembé) on finite lattices of $L \times L$ unit cells with periodic boundary conditions, with typical system sizes of $L = 12$ (864 sites). Proper thermalization at low temperatures is ensured by a parallel-tempering (replica-exchange) scheme [251, 252] with 192 logarithmically spaced temperatures between $T_{\min} = 10^{-4}$ and $T_{\max} = 10$. A detailed description of the Monte Carlo setup is given in the appendix of Ref. [P6].

The semianalytical method [225] refines classical ground states obtained from Monte Carlo simulations by reducing them to a small set of representative spin directions. Starting from the numerical spin configuration, one groups nearly parallel spins, identifies their underlying symmetry, and parametrizes the remaining independent spin directions by a few variables (at most $2K - 1$, where K is the number of symmetry inequivalent spin orientations). The classical energy is then expressed as a function of these parameters and numerically minimized, using the Monte Carlo result as an initial guess. This effectively yields an analytical expression of the energy, and thus allows the determination of exact phase boundaries. A more detailed explanation of the method is given in Refs. [P6, 225].

6.4.2 Phase diagram of the antiferromagnet

We begin with the case of AFM nearest-neighbor interactions $J_1 > 0$. The quantum phase diagram is obtained using pf-FRG in parameter regions where the classical analysis predicts phase transitions or noncoplanar orders. Our pf-FRG simulations employ a frequency grid of $40 \times 35 \times 35$ and lattice truncations up to $L = 15$. Ordered phases are identified by comparing spin structure factors with the corresponding Monte Carlo results, while paramagnetic—and thus putative QSL—regimes are detected as regions without a flow breakdown in the structure factor, as described in Sec. 4.5.1.

The combined classical and quantum phase diagrams, together with representative structure factors and common-origin plots, are shown in Fig. 6.12. Panel (a) displays the critical scale Λ_c , where regions with $\Lambda_c = 0$ (black) correspond to paramagnetic candidates for QSL phases. Panel (b) shows the evolution of the structure factor peaks, plotting the distance of \mathbf{k}^{\max} from \mathbf{K}^Δ (rather than \mathbf{k}^{\max} directly, as in previous sections) to better distinguish the observed phases. The solid gray lines denote classical phase boundaries obtained via Monte Carlo combined with the semianalytical method. In total, we find six distinct classical phases, labeled by Roman numerals: four noncoplanar (marked with an asterisk) and two coplanar. Below, we briefly outline the nature of these classical phases before turning to the role of quantum fluctuations. A detailed discussion of the real-space structures, ground-state energies, and symmetries of these classical phases is provided in Ref. [P6].

Classical states The coplanar phase I corresponds exactly to the $c120^\circ$ state already discussed in the previous two sections and illustrated in Fig. 6.2(c). Phase III is a distinct coplanar state with six spin sublattices, whose six-site magnetic unit cell coincides with the geometric unit cell of the lattice (i.e., it is a $\mathbf{q} = 0$ order). Both of these phases can be described exactly within LT using a single- \mathbf{q} Ansatz.

The noncoplanar phases II and VI can also be captured exactly by LT, but require a triple- \mathbf{q} Ansatz where the three wave vectors $\mathbf{q}^{\text{LT}} = (0, \pi)^T, (\pi, 0)^T, (\pi, \pi)^T$ (in the basis where the reciprocal lattice vectors are $\mathbf{G}_1 = (2\pi, 0)^T, \mathbf{G}_2 = (0, 2\pi)^T$) are assigned to the x , y , and z

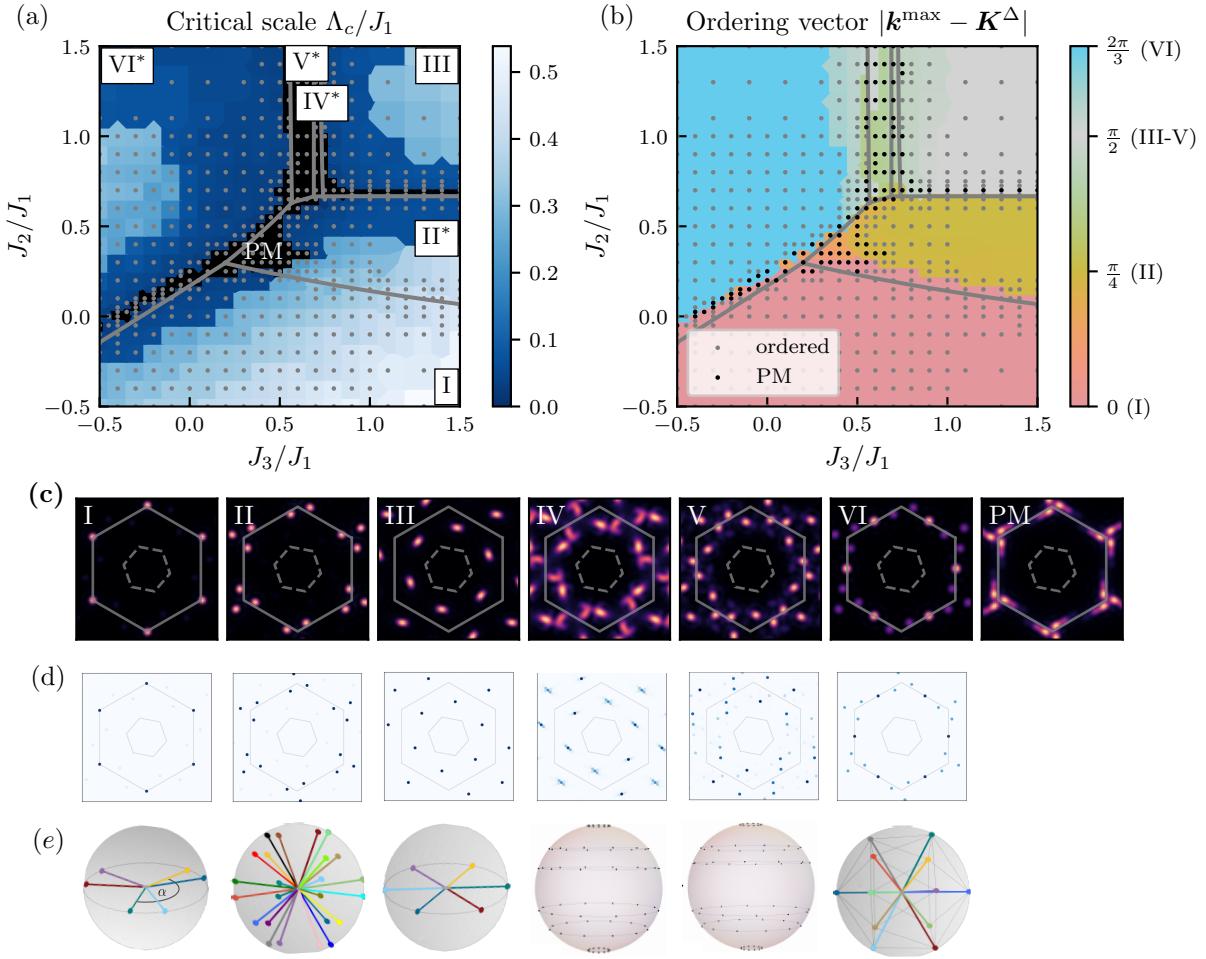


Figure 6.12 – Phase diagram of the antiferromagnet ($J_1 > 0$). (a) Critical scale Λ_c , at which the renormalization group flow develops an instability, signaling the onset of long-range order. Black regions indicate the absence of an instability ($\Lambda_c = 0$), corresponding to a PM ground state. Gray dots mark the coupling points where calculations were performed. Gray lines denote the classical phase boundaries, obtained from a semianalytical analysis. Roman numerals label the different ordered phases, with asterisks indicating noncoplanar ones. (b) Distance between the momentum of the structure factor maximum \mathbf{k}^{\max} and \mathbf{K}^Δ . (c) pf-FRG structure factors for all observed phases. (d) Classical structure factors for all ordered phases obtained from Monte Carlo simulations for the same parameters as in (c). (e) Corresponding common-origin plots of the spin sublattices for each ordered phase, as determined from LT and the semianalytical analysis. (d, e) are directly taken from Ref. [P6]

spin components, respectively. This construction is discussed in detail for phase VI in Sec. 3.1. Phase II realizes a noncoplanar state with twenty-four spin sublattices, and a magnetic unit cell of the same size. Phase VI also has a twenty-four-site magnetic unit cell but only twelve distinct spin orientations, pointing to the vertices of a deformed icosahedron which becomes regular at $(J_2, J_3) = (0, -1)$, as illustrated in the common-origin plot of Fig. 6.12(d).

By contrast, the noncoplanar phases IV and V cannot be obtained from LT. Here, the semianalytical method identifies 72 distinct spin sublattices, parametrized by six variables in phase V and four in phase IV. However, both unconstrained LT and pf-FRG indicate that the region between phases III and VI is instead governed by incommensurate (ICS) momenta, which are inaccessible to finite-size Monte Carlo simulations. It is therefore likely that phases IV and V do

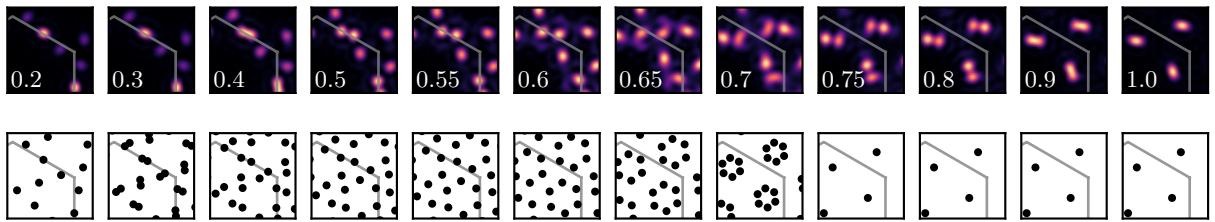


Figure 6.13 – Evolution of the structure factor across the incommensurate regime. Top: pf-FRG structure factors for fixed $J_2/J_1 = 1$ and varying J_3/J_1 (values indicated at the bottom left), showing the evolution from phase VI (left) to phase III. Bottom: Minimal \mathbf{q}^{LT} vectors for the same parameters. The continuous evolution indicates a single phase with incommensurate (ICS) correlation patterns, rather than the distinct phases IV and V identified in Monte Carlo simulations.

not correspond to two separate states, but rather represent a single incommensurate noncoplanar phase, as we discuss further below.

The effects of quantum fluctuations Before discussing putative chiral QSL phases, we first compare the pf-FRG structure factors with their classical counterparts, both shown in Fig. 6.12(c, d). Deep within the ordered phases I, II, III, and VI, the agreement is excellent, as expected. In contrast, in the region between phases III and VI the pf-FRG structure factor evolves smoothly from one to the other, with peaks shifting continuously to incommensurate (ICS) momenta, rather than showing the two distinct phases IV and V identified in the classical analysis. This behavior is illustrated in Fig. 6.13, which also demonstrates that the \mathbf{q}^{LT} vectors show equivalent behavior. A similar situation occurs near the junction of phases I, II, and VI, where pf-FRG again reveals ICS momenta. Together, these results strongly suggest that the classical phase diagram likewise hosts ICS phases in these regions, which cannot be faithfully resolved by finite-size Monte Carlo simulations.

Strikingly, it is precisely in those regions where the classical and quantum structure factors disagree that extended paramagnetic regimes appear. The most prominent example is a broad quantum-disordered region around the triple point where phases I, II, and VI meet. Here, the strong competition between neighboring orders, amplified by quantum fluctuations, appears sufficient to melt the magnetic order entirely. Since this regime extends deeply into the classical noncoplanar phase II, it is a promising candidate for a chiral QSL.

Surprisingly, the extended regime encompassing the classical noncoplanar ICS phases IV and V shows very similar behavior. Although we do not observe a flow breakdown, the structure factor still exhibits relatively sharp peaks in these regions, as shown in Fig. 6.12(c). Additionally, ICS states typically order at lower critical temperatures, which makes such transitions harder for the pf-FRG to detect. As a result, we cannot unambiguously determine whether this region is truly quantum-disordered or an artifact of the method. If genuine, such a paramagnetic state would coincide directly with the classical noncoplanar ICS regime, again making it a strong candidate for a chiral QSL.

6.4.3 Phase diagram of the ferromagnet

We now consider the case of ferromagnetic couplings $J_1 < 0$. The combined quantum phase diagram, classical phase boundaries, and pf-FRG structure factors are presented in Fig. 6.14. For the classical calculation, we restrict to a Luttinger-Tisza analysis, which captures all phases analytically except for a regime characterized by \mathbf{q}^{LT} vectors at incommensurate (ICS) momenta

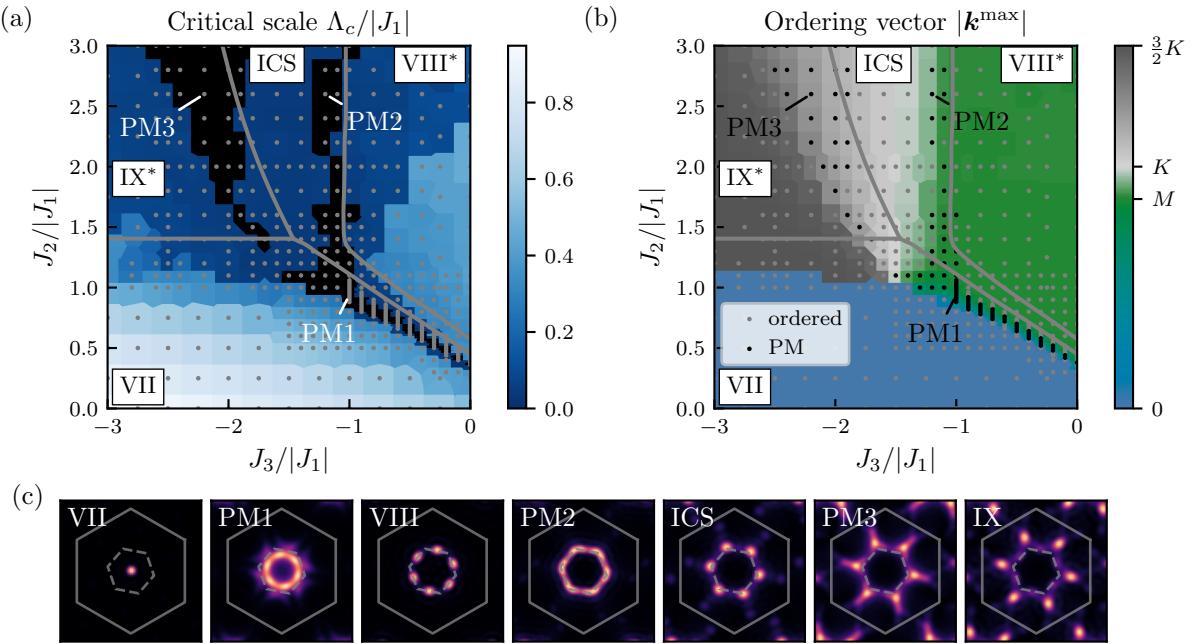
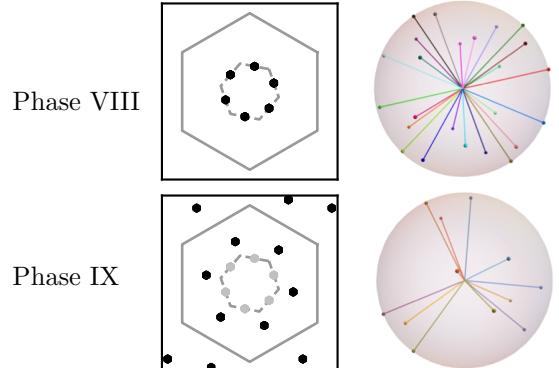


Figure 6.14 – Phase diagram of the ferromagnet ($J_1 < 0$). (a) Critical scale Λ_c , at which the renormalization group flow develops an instability, signaling the onset of long-range order. Black regions indicate the absence of an instability ($\Lambda_c = 0$), corresponding to a PM ground state. Gray dots mark the coupling points where calculations were performed. Gray lines denote the classical phase boundaries, obtained from a Luttinger-Tisza analysis. Roman numerals label the different ordered phases, with asterisks indicating noncoplanar ones. (b) Absolute value of the momentum \mathbf{k}^{\max} where the structure factor is maximal. (c) pf-FRG structure factors for all observed phases. Classical counter-parts are in Fig. 6.15

Figure 6.15 –

Classical noncoplanar ground states of the ferromagnet. Shown are the exact classical structure factors (left) and common origin plots (right) for the two noncoplanar phases found in the phase diagram of the ferromagnet in Fig. 6.14. Both spin configurations can be exactly captured using Luttinger-Tisza with a triple- \mathbf{q} Ansatz. Common origin plots are taken directly from Ref. [P6].



that is also detected in pf-FRG. Since Monte Carlo simulations cannot reliably resolve such ICS regions, we do not attempt a detailed classical characterization here. As before, we first outline the classical phases observed—a ferromagnet, two noncoplanar states, and an extended ICS regime—before turning to the impact of quantum fluctuations.

Classical phases In addition to a large FM regime (phase VII), the Luttinger-Tisza analysis identifies two noncoplanar phases, VIII and IX, whose exact classical structure factors and common-origin plots are shown in Fig. 6.15. Both are described by a triple- \mathbf{q} Ansatz with $\mathbf{q}^{\text{LT}} = (0, \pi)^T, (\pi, 0)^T, (\pi, \pi)^T$ assigned to the three spin components. Phase VIII realizes a

noncoplanar state with a twenty-four-site magnetic unit cell and equally many spin sublattices. While the detailed configuration evolves with J_2 and J_3 , its symmetry and qualitative form remain as illustrated in Fig. 6.15. Phase IX also has a twenty-four-site magnetic unit cell, but only twelve spin sublattices. The spin directions form a polyhedron composed of four equilateral triangles, which we term a deformed truncated tetrahedron. Unlike the Archimedean truncated tetrahedron (with 12 vertices, four regular hexagons, and four triangles), this structure lacks regular hexagons. In this phase, the overall spin configuration remains fixed as J_3 varies, but the size and orientation of the equilateral triangles depend on J_2 .

Effects of quantum fluctuations Deep within the ordered phases, the classical and quantum calculations again show good agreement. However, reminiscent of the nearest-neighbor model with competing ferro- and antiferromagnetic interactions discussed in the previous section, quantum fluctuations significantly shift the phase boundaries by reducing the extent of the FM regime (phase VII). Instead, a paramagnetic region (PM 1) emerges near the boundary of phase VIII. As also discussed in Sec. 6.3.4 this makes the region a promising candidate for multipolar order, and its structure factor indeed closely resembles that of point PM D in the putative spin-nematic regime of the nearest-neighbor phase diagram [Fig. 6.10]. This suggests possible spin-nematic behavior, although we have not explicitly computed the corresponding response.

Of central interest here is the ICS regime between the noncoplanar phases VIII and IX. Traversing this region from phase VII to phase IX, the structure-factor peaks evolve continuously, shifting first from \mathbf{M} to \mathbf{K} , and then from \mathbf{K} to $\frac{3}{2}\mathbf{K}$. In no extended regime do the \mathbf{q}^{LT} or \mathbf{k}^{max} vectors lock exactly onto \mathbf{K} , consistent with an incommensurate phase. At the boundaries of this ICS regime with the noncoplanar phases VIII and IX, two extended paramagnetic regions (PM2 and PM3) emerge. Both show no flow breakdown and only broad features in the structure factor, as illustrated in Fig. 6.14(c). While PM2 lies adjacent to the boundary of noncoplanar phase VIII, PM3 extends largely over the classical noncoplanar IX region. These regimes are therefore strong candidates for chiral QSL ground states.

6.5 Discussion

Our classical analysis demonstrated that introducing long-range cross-plaquette interactions on the maple-leaf lattice stabilizes a variety of noncoplanar magnetic orders. In the $S = 1/2$ limit, we find multiple instances where quantum fluctuations suppress long-range order, opening the possibility of chiral QSL phases emerging through the melting of noncoplanar chiral parent states.

Since the maple-leaf lattice lacks reflection symmetry about any straight line, it would be particularly interesting to explore whether other lattice symmetries could be broken—up to time reversal—to realize the combined PT symmetry required for the specific class of U(1) chiral QSLs [213]. A systematic classification of chiral mean-field Ansätze via the projective symmetry group [245], combined with variational Monte Carlo to evaluate correlations (as recently done for fully symmetric QSLs [237]), could provide a concrete route for comparison with the classical and pf-FRG structure factors.

It would also be valuable to investigate the phase diagram with other modern numerical techniques. The pf-FRG, while powerful, carries significant uncertainties, and previous sections already highlighted how tensor network methods can yield markedly different results in related models. The large unit cell and incommensurate nature of correlations on the maple-leaf lattice, however, make this a particularly challenging problem. As a first step, the pf-FRG framework could be further exploited to probe tendencies toward dimer or, especially relevant for the FM,

spin-nematic order, using the response-function approach introduced in Sec. 4.5.3.

Although not explicitly presented here, we also studied the finite temperature behavior of the classical model [P6] with classical Monte Carlo to investigate the nature of the thermal transition into the ordered ground state. Although the Mermin-Wagner theorem forbids the breaking of the continuous $SU(2)$ symmetry in two dimensions, in noncoplanar phases the spontaneous selection of a certain chirality is a discrete symmetry breaking that is allowed even in two dimensions. We indeed found that the noncoplanar phases show strong indications of a phase transition at relatively large temperatures (above $T > \gtrsim 0.2J_1$). Here, using the recently developed pseudo-majorana FRG (pm-FRG) [14, 140, 141] that generalizes the pf-FRG to finite temperatures would be well suited to probe for similar transitions. It would also have the advantage that, at such large temperatures, the method has even shown to obtain quantitative agreement with exact methods, and an identification of a paramagnetic regime in pm-FRG would thus substantially strengthen the claim of a putative QSL ground state.

Moreover, as a natural next step, the pf-FRG framework itself could be further exploited to explore competing ordering tendencies—such as dimer order or, in the ferromagnetic case, spin-nematic order—via the response-function approach introduced in Sec. 4.5.3.

Chapter 7

Spin-valley magnetism in moiré materials

In this chapter, we turn to the realm of two-dimensional materials. While the previous chapter focused on the maple-leaf lattice, where two-dimensionality is only effective—the underlying minerals are three-dimensional crystals with weak interlayer coupling—here we consider systems that are intrinsically two-dimensional, consisting of only a few atomic layers, or even a single layer. The prime example is graphene: a single sheet of carbon atoms arranged in a honeycomb lattice. Since the groundbreaking discovery in 2004 that high-quality, free-standing graphene sheets can be isolated with relative ease [253], research on two-dimensional materials has expanded explosively. Beyond graphene, a broad family of atomically thin crystals has become experimentally accessible, most prominently transition-metal dichalcogenides (TMDs) [254] and hexagonal boron nitride (h-BN) [255], which is widely used as an atomically flat insulating substrate.

Such materials are of great interest because of their high tunability. For instance, the carrier density—that is, the electronic filling—can be directly controlled by applying a gate voltage. This stands in sharp contrast to three-dimensional systems, where changing the filling typically requires chemical doping and is therefore much more constrained. Even greater flexibility arises when stacking multiple layers of two-dimensional materials: the number of layers, their relative alignment, and the choice of materials provide powerful control knobs that can dramatically alter material properties and give rise to physics entirely different from that of a single isolated layer.

For example, single-layer graphene is a semimetal with linear dispersion near the \mathbf{K} and \mathbf{K}' points discussed in Sec. 2.2.1, where electronic correlations remain relatively weak [71]. In contrast, twisted bilayer graphene (TBG)—comprising two stacked graphene layers rotated by a small relative angle—exhibits markedly different behavior: rotating one layer relative to the other produces a large-scale moiré pattern consisting of locally aligned and misaligned regions (illustrated in Fig. 2.3). This pattern, and the resulting spatial modulation of the interlayer hopping amplitudes and potentials, gives rise to extremely flat electronic bands that strongly enhance otherwise weak electronic correlations. Such flat bands occur only at specific, so-called “magic angles,” most notably around $\theta \approx 1.1^\circ$ [77, 78].

Experiments on magic-angle TBG indeed reveal strong correlation phenomena, including Mott insulating states and unconventional superconductivity, which can be accessed by tuning the electronic filling with a gate voltage [256–259]. Similar correlated behavior has also been observed in other stacked graphene systems with different number of layers and alignments, such as twisted double bilayer graphene (TDBG) [260–263], and—of particular relevance for this chapter—in trilayer graphene aligned with hexagonal boron nitride (TG/h-BN) [264–268]. A further emergent class of platforms are moiré systems of two-dimensional TMDs [269–272], which exhibit analogous correlation-driven effects.

The emergence of flat bands can already be understood within the low-energy tight-binding framework of graphene. As discussed in Sec. 2.2.1, the essential physics derives from the linear dispersion at the Dirac cones located at \mathbf{K} and \mathbf{K}' . Effective theories for moiré materials

therefore typically start from an expansion around these points and subsequently incorporate interlayer couplings and external potentials. Because the states near \mathbf{K} and \mathbf{K}' are degenerate, this expansion introduces an additional bi-valued quantum number—the *valley degree of freedom*—that distinguishes electrons near \mathbf{K} from those near \mathbf{K}' . The resulting Hubbard models thus describe fermions carrying both spin and valley degrees of freedom, hopping on the emergent moiré superlattice. While the precise parameters of these models—and in particular their position along the weak-to-strong coupling axis—remain uncertain, the observation of Mott insulating states motivates studying their strong-coupling limit. As discussed in Sec. 2.2, this naturally yields quantum spin models where the localized moments are described by spin-valley operators, i.e., generators of $SU(4)$ rather than the $SU(2)$ spins of conventional models. Such $SU(4)$ spin-valley models have, for example, been explicitly derived for TBG [48] and TG/h-BN [49].

As discussed in the introductory Chapter 1, an exact $SU(4)$ symmetry generally enhances quantum fluctuations compared to $SU(2)$. However, once symmetry-breaking terms are introduced, this enhancement can be suppressed, and the resulting interactions may strongly reshape the phase diagram—potentially stabilizing distinct spin- and/or valley-ordered states. In the effective spin-valley models for moiré materials, the $SU(4)$ symmetry is indeed typically strongly broken by exchange interactions to $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$ through various coupling terms between the spin and valley sectors.

This raises two central questions: (i) which ordered states are favored by these $SU(4)$ -breaking terms, and (ii) can quantum fluctuations still stabilize paramagnetic phases despite the symmetry reduction. To address the first question, we employ the semiclassical Monte Carlo method for $SU(4)$ models introduced in Sec. 3.2, which also captures the role of thermal fluctuations. For the second, we use the pf-FRG framework generalized to spin-valley systems as described in Sec. 4.3. In this chapter, we apply these approaches to concrete models for two different systems:

In the first section, based on Ref. [P2], we study one of the simplest moiré setups: monolayer graphene subject to a substrate-induced potential. Despite its simplicity, the noninteracting band structure already shows rich physics, with highly localized states coexisting with dispersive one-dimensional chiral channels. A strong-coupling expansion in the presence of Coulomb interactions yields an effective $SU(4)$ spin-valley model with peculiar chiral interactions. We analyze the semiclassical mean-field phase diagram of this model, and for particularly degenerate phases study the role of thermal fluctuations using our semiclassical Monte Carlo implementation. We also compare our results with $SU(4)$ spin-wave calculations, finding good agreement and thus providing a good benchmark for our implementation.

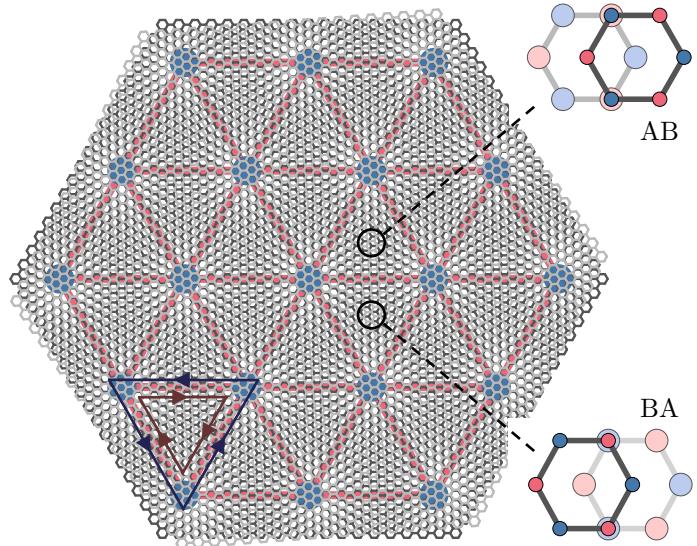
In the second section, based on Ref. [P1], we investigate trilayer graphene aligned with hexagonal boron nitride (TG/h-BN). In this system, both the electronic filling and an external displacement field act as tuning parameters, and experiments report correlated insulating states at large fields and integer fillings. Building on the effective spin-valley model derived in Ref. [49], we explore the phase diagram using a combination of semiclassical Monte Carlo and pf-FRG, with the aim of clarifying the competing ordered states and the potential for quantum-disordered phases.

7.1 Single-layer graphene in a moiré potential

In this section, we consider monolayer graphene placed on an insulating substrate that shares the hexagonal structure of graphene but either has a slightly different lattice constant or is rotated by a small twist angle. This configuration generates a moiré pattern, as illustrated in Fig. 7.1.

Figure 7.1 –

Triangular moiré superlattice and network model. Two hexagonal layers with a small twist angle form a moiré pattern. The interlayer potential varies on the moiré scale due to different stacking arrangements (schematically shown on the right). Blue circles mark AA-stacked regions, where the layers perfectly overlap and localized states reside. Red lines trace the one-dimensional chiral channels (1DCCs), which are valley-polarized such that electrons in opposite valleys move in opposite directions, as indicated by the arrows inside for one triangle on the bottom left.



In contrast to TBG, the insulating nature of the substrate prohibits interlayer tunneling, and the moiré potential acts only through a spatially modulated electrostatic field.

Following Ref. [P2], we begin by briefly outlining the derivation of an effective network model that captures the low-energy physics of this system, including the emergence of highly localized states and one-dimensional chiral channels (1DCCs) that interact in a characteristic manner. In the strong-coupling limit, the effective model reduces to an $SU(4)$ spin-valley Hamiltonian with $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$ symmetry and chiral interactions.

The main part of this section focuses on studying the semiclassical zero-temperature phase diagram of this model, as well as the role of thermal fluctuations in particularly degenerate regions, using a combination of mean-field methods and our semiclassical Monte Carlo implementation discussed in Sec. 3.2. This analysis not only reveals an intriguing mechanism by which fluctuations induce noncoplanar chiral order, but also highlights the capabilities of our Monte Carlo approach. In addition, by comparing with $SU(4)$ spin-wave calculations at low temperatures, we obtain a valuable benchmark for our method.

Unfortunately, the relevant filling for the model under consideration is one electron per site (quarter filling), which cannot be accessed with our current pf-FRG implementation (as discussed in Sec. 4.3). We therefore refrain from analyzing quantum fluctuations in this case.

We note that the central achievement of Ref. [P2] is the identification of the localized states and chiral channels, together with the derivation of the resulting exotic effective interactions. Full credit for this, as well as for the mean-field and $SU(4)$ spin-wave analyses, goes to my coauthors of Ref. [P2], particularly Jeyong Park and Jinhong Park. My contribution was the analysis of thermal fluctuations via semiclassical Monte Carlo simulations, as well as the confirmation of their zero-temperature mean-field phase diagram in certain regions. This also forms the primary focus of the present section. Except for the mean-field phase diagrams in Fig. 7.2, all data shown in the figures were obtained by the author of this thesis and are redrawn from Ref. [P2].

7.1.1 Derivation of the spin-valley model

The starting point for the effective description of the moiré system is the low-energy theory of graphene, obtained by expanding the tight-binding model around the valleys \mathbf{K} and \mathbf{K}' . This

yields the massless Dirac Hamiltonian (see Ref. [71] for details),

$$H_0 = -i\nu (\partial_x s^x \tau^z + \partial_y s^y \tau^0), \quad (7.1)$$

with Pauli matrices τ^a and s^a acting in valley and sublattice space, and $\nu = 3ta/2$ the Fermi velocity.

An insulating substrate induces an electrostatic potential that can be decomposed into a uniform part $V_0 = (V_A + V_B)/2$ and a staggered part $V_s = (V_A - V_B)/2$, where $V_{A/B}$ act on the respective sublattices. Including these contributions leads to the effective single-particle Hamiltonian

$$H = -i\nu (\partial_x s^x \tau^z + \partial_y s^y \tau^0) + V_0(\mathbf{r}) \mathbf{1} + V_s(\mathbf{r}) s^z. \quad (7.2)$$

Here, $V_0(\mathbf{r})$ and $V_s(\mathbf{r})$ vary smoothly in space due to the moiré pattern. The uniform term shifts both sublattices equally, while the staggered term alternates sign depending on the local stacking configuration. For instance V_s has its maximum in the AB and BA stacked regions illustrated in Fig. 7.1

Assuming the moiré potential varies smoothly, it can be represented by a Fourier expansion restricted to the leading (shortest) reciprocal lattice vectors of the triangular superlattice. The resulting band structure exhibits two characteristic types of states. First, there are one-dimensional channel states (1DCCs), which form dispersive bands whose Bloch wave functions are strongly localized along the lines connecting the AA-stacked regions (illustrated as red lines in Fig. 7.1). These channels are valley-polarized: states originating from the \mathbf{K} and \mathbf{K}' points propagate in opposite directions (indicated by arrows in Fig. 7.1). Second, there are highly flat bands associated with states localized in the AA regions (blue dots in Fig. 7.1). The widths of both the 1DCCs and the localized states scale as $\sim 1/L$, with L the moiré lattice spacing. Consequently, for sufficiently large moiré unit cells the localization is strong and hybridization between the different electronic states is weak.

To capture the interaction between localized states and 1DCCs, one can derive an effective network model that includes (i) the kinetic energy of the 1DCCs, (ii) hybridization between two 1DCCs at their crossing points, and (iii) hybridization between localized states and 1DCCs. Due to the strong localization, there will additionally be an energy penalty when more than two electrons occupy these states that can be modeled by an on-site Hubbard interaction U . Since the hybridization is weak, and localization is strong, one can perform a strong coupling expansion to derive the effective interaction between the localized states that, in this case, takes the form of an RKKY interaction [273, 274] mediated by the 1DCSs ¹

The detailed derivation and full Hamiltonian is presented in Ref. [P2]. To study the competition of the interaction, we restrict to nearest neighbor interactions, which results in the spin-valley Hamiltonian

$$H = J_2 \sum_{\langle ij \rangle} (1 + \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j) \left(e^{i\varphi} \tau_i^+ \tau_j^- + \text{h.c.} \right) + J'_2 \sum_{\langle i \rightarrow j \rangle} (1 + \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j) (1 + \tau_i^z \tau_j^z) \quad (7.3)$$

$$+ J_3 \sum_{p=\nabla/\Delta} p \left(\prod_{k=1}^3 P_{p_k}^+ - \prod_{k=1}^3 P_{p_k}^- \right) \boldsymbol{\sigma}_{p_1} \cdot (\boldsymbol{\sigma}_{p_2} \times \boldsymbol{\sigma}_{p_3}) \quad (7.4)$$

¹This is different to the mechanism behind the typical exchange interaction discussed in Sec. 2.2, as the interactions are mediated by the 1DCCs (which are not localized) instead of virtual hopping processes between localized orbitals themselves. This is an interaction typically found in metallic magnets. It typically has the form $\sim \cos(2k_F r)/r^3$ where k_F is the Fermi-momentum, but we neglect the long-range behavior here for simplicity.

Here, $\boldsymbol{\sigma}_i = (\sigma_i^x, \sigma_i^y, \sigma_i^z)^T$ and $\boldsymbol{\tau}_i = (\tau_i^x, \tau_i^y, \tau_i^z)^T$ denote SU(2) operators acting on the spin and valley degrees of freedom, respectively. More precisely, the operators in the Hamiltonian are spin–valley operators of the form $\sigma_i^\mu \tau_i^\kappa$ with $\mu, \kappa \in \{0, 1, 2, 3\} = \{d, x, y, z\}$, as defined in Eq. (3.39). These act on the full local Hilbert space, with $\sigma_i^\mu \equiv \sigma_i^0 \tau_i^0$ and $\tau_i^\kappa \equiv \sigma_i^0 \tau_i^\kappa$. At quarter filling (one electron per moiré unit cell), the spin and valley operators factorize as simple tensor products, $\sigma_i^\mu \tau_i^\kappa = \sigma_i^\mu \otimes \tau_i^\kappa$, which justifies the separation in notation. This factorization does not hold at half-filling, as discussed in detail in Sec. 3.2. The operators τ_i^\pm denote raising and lowering operators in valley space, while $P_i^\pm = (\tau_i^0 \pm \tau_i^z)/2$ projects onto the \pm valley. In the third term of the Hamiltonian, the sum runs over all up- and down-pointing triangles, with p_1, p_2, p_3 labeling the sites of each triangle in anticlockwise order. The prefactor $p = \pm 1$ alternates between up- and down-pointing triangles. The dominant coupling in the effective model is $J_2 > 0$, while J'_2 and J_3 enter as perturbations. Due to symmetry we only have to consider $\varphi \bmod 2\pi/3$ [P2].

The Hamiltonian exhibits an $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$ symmetry, or, more precisely, an $SU(2)_{\text{spin}}^+ \otimes SU(2)_{\text{spin}}^- \otimes U(1)_{\text{valley}}^-$ symmetry, since spin rotations in each valley leave it invariant independently. The corresponding symmetry generators are $P^\pm \boldsymbol{\sigma}$ and τ^z . The symmetry-breaking terms J_2 and J'_2 arise from RKKY-type processes involving electrons propagating to a neighboring site and back. Owing to the chirality of the one-dimensional channels, such processes necessarily involve two valley flips, $\sim \tau_i^- \tau_j^+$ that break $SU(4)$ symmetry. The chiral interaction $\sim J_3$ originates from processes in which electrons traverse closed loops around a triangle, with the staggered sign again reflecting the valley-dependent chirality of the 1DCCs.

7.1.2 Mean-field phase diagram

To obtain the zero-temperature phase diagram of such a model, in principle, the minimization procedure in the semiclassical limit discussed in Sec. 3.2 could be applied. In our semiclassical Monte Carlo implementation, however, we have not yet included complex couplings (relevant for $\varphi \neq 0, \pi$) or three-spin interactions (relevant for $J_3 \neq 0$), although these extensions are conceptually possible. We are therefore restricted to the case of finite J_2 and J'_2 with $\varphi = 0$ and $J_3 = 0$ in our current method.

To still obtain a mean-field phase diagram in the full parameter space, our collaborators instead performed $SU(4)$ mean-field calculations, which are equivalent to our approach: the interactions in the Hamiltonian can be decoupled according to $T_i^a T_j^b \approx \langle T_i^a \rangle T_j^b + T_i^a \langle T_j^b \rangle - \langle T_i^a \rangle \langle T_j^b \rangle$, where T^a denotes components of the spin–valley operators. Choosing a magnetic unit cell of size N_{MF} fixes the periodicity of the expectation values $\langle T_i^a \rangle$. This reduces the problem to $i = 1, \dots, N_{\text{MF}}$ independent single-site Hamiltonians of the form $H_i = \sum_a T_i^a h_i^a(\{\langle T_j^a \rangle\})$, which can be minimized separately in the local four-dimensional Hilbert space. The resulting ground state is a semiclassical product state, $|\Psi\rangle = \bigotimes_i |\psi_i\rangle$, constructed from the solutions on each site. Starting from random initial values of the expectation values, this procedure yields a set of self-consistent equations that can be solved numerically. Our collaborators carried out the mean-field analysis for $N_{\text{MF}} = 1, 2, 3, 4$ and consistently found that only the one-sublattice and three-sublattice solutions minimize the energy. For the region with $\varphi = 0$ and $J_3 = 0$, our Monte Carlo approach perfectly agrees with their method. The mean-field phase diagrams in the full parameter space are shown in Fig. 7.2, which we discuss in the following.

Degenerate ground state manifold of the J_2 -only Hamiltonian H_0 Since J_2 is the dominant coupling, let us begin by considering the J_2 -only Hamiltonian with $J'_2 = J_3 = \varphi = 0$. In this limit,

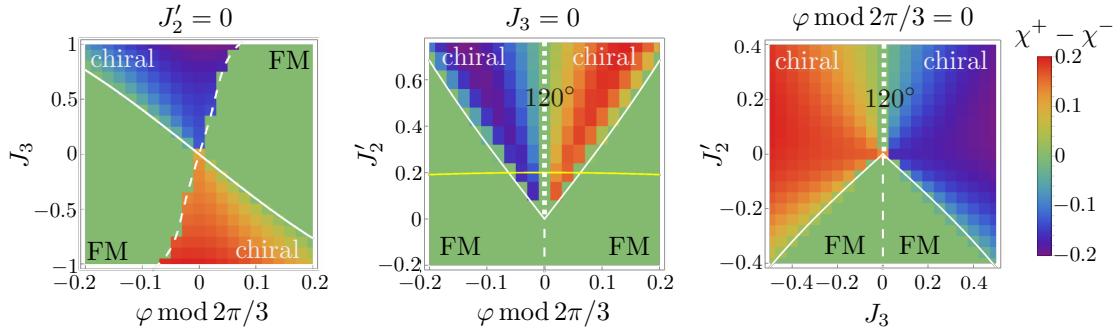


Figure 7.2 – Zero-temperature mean-field phase diagram of the spin–valley model. Ground-state phase diagrams of Hamiltonian (7.3) are shown in different parameter planes. In all cases, the mixed spin–valley degrees of freedom μ^a form 120° order, while the valley-projected spin σP^\pm exhibits ferromagnetic order in one valley. The second valley may host noncoplanar order with finite chirality $\chi^+ - \chi^-$, ferromagnetic order with vanishing chirality (green), or 120° order with vanishing chirality (dotted white lines). Solid lines denote continuous transitions, while dashed white lines mark transitions where the sense of 120° order in μ^a changes between left- and right-circulating. The yellow line indicates how the couplings could be tuned by a gate voltage. Figure and data taken and adapted from Ref. [P2].

the Hamiltonian can be rewritten as

$$H_0 = 2J_2 \sum_{\langle ij \rangle} (1 + \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_j) \left(\tau_i^x \tau_j^x + \tau_i^y \tau_j^y \right) = 2J_2 \sum_{\langle ij \rangle} (\boldsymbol{\mu}_i^1 \cdot \boldsymbol{\mu}_j^2 + \boldsymbol{\mu}_i^2 \cdot \boldsymbol{\mu}_j^1) \quad (7.5)$$

where in the second step we introduced the four-component vectors

$$\boldsymbol{\mu}_i^1 = (\tau_i^x, \tau_i^y \sigma_i^x, \tau_i^y \sigma_i^y, \tau_i^y \sigma_i^z) \quad \boldsymbol{\mu}_i^2 = (\tau_i^y, \tau_i^x \sigma_i^x, \tau_i^x \sigma_i^y, \tau_i^x \sigma_i^z). \quad (7.6)$$

The classical ground state of H_0 is minimized when these vectors form a four-dimensional 120° pattern between neighbors, such that $\boldsymbol{\mu}_i^a \cdot \boldsymbol{\mu}_j^a = \cos(2\pi/3)$ for $a = 1, 2$. This condition, however, does not uniquely fix the ground state but leaves a large degeneracy. In particular, the spin expectation values projected to the two valleys, $\langle \boldsymbol{\sigma} P^\pm \rangle$, also develop different kinds of order. In one valley the spin is always ferromagnetic, while in the other valley, a noncoplanar spin configuration may emerge. This can be diagnosed by the staggered spin chirality

$$\chi^\pm = \frac{1}{2N} \sum_{p=\nabla/\Delta} p \langle \boldsymbol{\sigma}_{p_1} P^\pm \cdot (\boldsymbol{\sigma}_{p_2} P^\pm \times \boldsymbol{\sigma}_{p_3} P^\pm) \rangle \quad (7.7)$$

which takes values in the range $-\frac{1}{8} \leq \chi^\pm \leq \frac{1}{8}$.

For H_0 in the mean-field (semiclassical) limit, all values of the chirality are degenerate. This degenerate manifold can be parameterized by an opening angle $\theta \in [0, \pi]$ between the spins and the z -axis, together with a discrete choice (\pm) specifying which valley hosts ferromagnetic order and which hosts chiral order. The special cases $\theta = 0, \pi$, corresponding to ferromagnetic order, and $\theta = \pi/2$, corresponding to 120° order, both yield $\chi^\pm = 0$ and are included in the degenerate manifold. All other cases $\theta \neq 0, \pi/2, \pi$ correspond to actual noncoplanar chiral order in one valley.

Effects of small perturbations For small perturbations by finite φ , J_2' , or J_3 —and in fact for all parameter regimes considered here—the order remains within the ground-state manifold of H_0 . This means the order in μ^a is always a coplanar 120° pattern, while the valley-projected spin

σP^\pm always exhibits ferromagnetic behavior in one valley, and FM, 120° , or noncoplanar order in the other valley. Consequently, when classifying different phases we only need to denote the spin order in this second valley.

Perturbations to H_0 lift the degeneracy of the manifold and select states of specific chirality. For example, an infinitesimal J_3 favors states with maximal chirality $\chi^\pm = \pm \frac{1}{8}$, a finite φ selects ferromagnetic order (with zero chirality), and a small $J'_2 > 0$ stabilizes 120° order (also with zero chirality). The full phase diagrams, showing the competition between these couplings, are presented in Fig. 7.3, where we fix $J_2 = 1$ throughout.

Particularly relevant for the remainder of this section is the region accessible to our Monte Carlo implementation, which includes only J_2 and J'_2 . This corresponds to the vertical $\varphi = 0$ cut through the phase diagram with $J_3 = 0$ (white line in the middle panel of Fig. 7.3). Interestingly, for $J'_2 > 0$ this line lies exactly on the phase boundary between two states with different chiral orders, yet the ground state itself is a nonchiral 120° configuration. We analyze the effects of thermal fluctuations along this line in parameter space in the following section.

7.1.3 Effects of thermal fluctuations

To investigate the effects of thermal fluctuations, we use our semiclassical Monte Carlo approach. A typical simulation consists of $N_m = 10^6$ thermalization sweeps, followed by $N_m = 4 \cdot 10^6$ measurement sweeps in most regions, but up to $N_m = 10^7$ sweeps in the vicinity of phase transitions. We simulate systems with linear size up to $L = 72$ (5184 sites) and periodic boundary conditions. The details of the algorithm are provided in Sec. 3.2.

Thermal order-by-disorder for H_0 We begin with the J_2 -only Hamiltonian H_0 , which represents the maximally degenerate point where all chiral states are equivalent. Several observables as a function of temperature at this point are shown in Fig. 7.3(a–c). The specific heat displays a sharp peak consistent with a second-order, or possibly very weak first-order, phase transition (as we discuss in more detail below). As $T \rightarrow 0$, the mixed spin–valley degrees of freedom μ_i^n (not shown explicitly) correctly develop 120° order. At the same time, the average valley-projected magnetization $|\langle \sigma P^\pm \rangle|$ approaches its maximal value in both valleys, and the chirality goes to zero, indicating that thermal fluctuations select a ferromagnetic configuration in both valleys in the spin degrees of freedom via an order-by-disorder mechanism. The same selection is found in a $T = 0$ SU(4) spin-wave analysis performed by my coauthors [P2].

At finite temperature, $|\langle \sigma P^\pm \rangle|$ is increasingly suppressed with growing system size L , as shown in Fig. 7.3(c). This behavior is consistent with the expected reduction of the order parameter, $\sim T \ln N$, as dictated by the Mermin–Wagner theorem: since the order parameter breaks a continuous symmetry generated by σP^\pm , it cannot remain finite in two dimensions at nonzero temperature.

More interestingly, however, at finite temperature the chirality (in one valley, spontaneously selected and thus conveniently parametrized by $|\chi^+ - \chi^-|$) becomes finite in a highly singular fashion at the putative phase transition, largely independent of system size [Fig. 7.3(b)]. Remarkably, the same effect also appears in an SU(4) spin-wave analysis when higher-order terms in the $1/M$ expansion are included (with M the number of local bosons used to represent the SU(4) operators). In that framework, my coauthors predicted the low-temperature behavior

$$\langle \chi^+ - \chi^- \rangle \approx \pm 0.22 \frac{T}{J_2} \ln \left(\sqrt{\frac{T_0}{T}} \right). \quad (7.8)$$

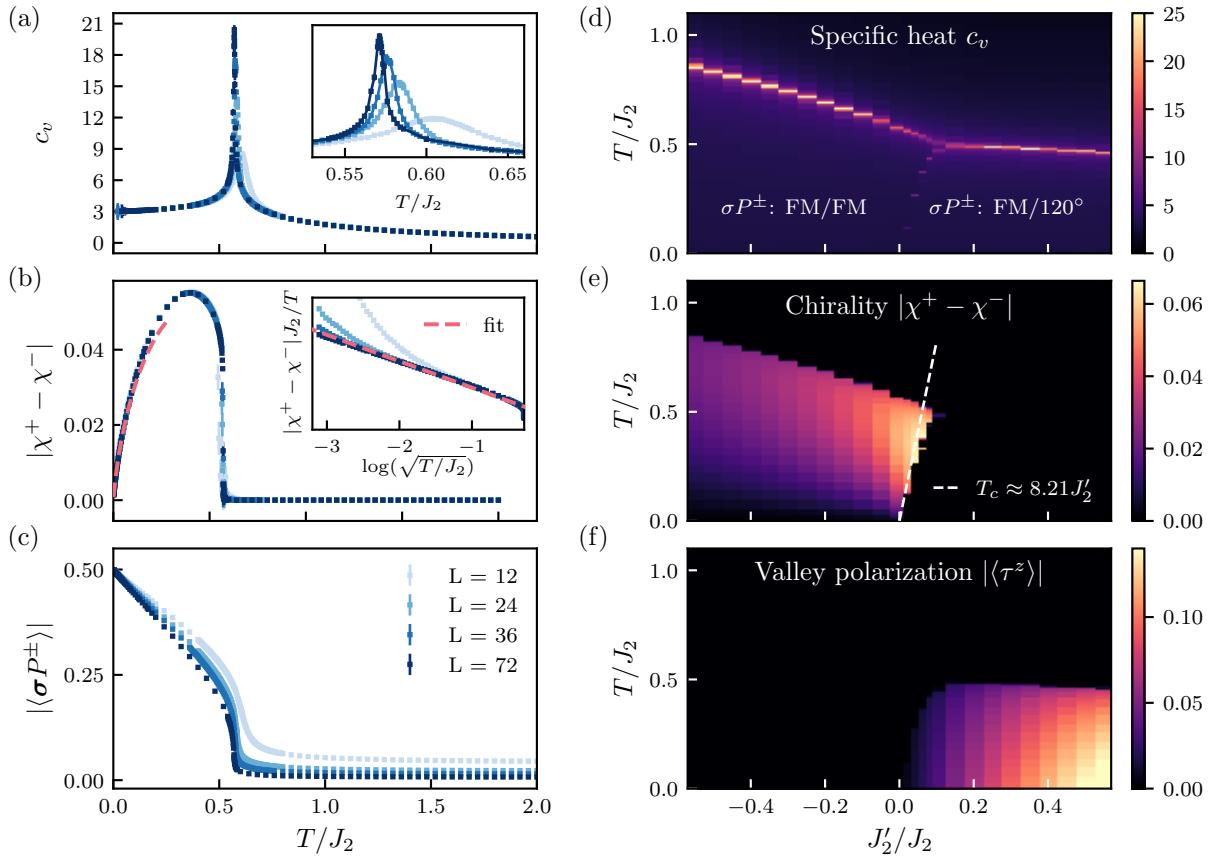


Figure 7.3 – Order-by-disorder and finite-temperature phase diagrams. (a–c) Monte Carlo results for the J_2 model H_0 [Eq. (7.5)] in the semiclassical approximation. Numerical errors are smaller than the symbols. (a) The specific heat shows a pronounced peak that sharpens with system size L , consistent with a thermal phase transition. (b) In the ordered phase, a finite spin chirality develops at finite T , with a singular low- T dependence $\sim T \log(1/\sqrt{T})$. Inset: fit to the SU(4) spin-wave prediction (dashed orange line). (c) An order-by-disorder mechanism selects ferromagnetic spin order in both valleys as $T \rightarrow 0$. At finite T , the long-range order parameter is increasingly suppressed with L , consistent with the Mermin–Wagner theorem. (d–e) Finite-temperature phase diagram versus J'_2 . Besides chirality, the valley polarization $|\langle \tau^z \rangle|$ is shown, distinguishing ferromagnetic order in both σP^\pm ($|\langle \tau^z \rangle| = 0$) from 120° order in one valley ($|\langle \tau^z \rangle| > 0$). The dashed white line in (b) shows the spin-wave estimate for the slope of the critical temperature at low T .

Fitting this expression (with T_0 as the only free parameter) to our Monte Carlo data yields excellent agreement at low temperatures, as illustrated by the dashed red line in Fig. 7.3(b). The system spontaneously selects a definite sign of $\langle \chi^+ - \chi^- \rangle$, thereby breaking the residual \mathbb{Z}_2 symmetry corresponding to a π rotation around τ^x (explicitly, $e^{i\pi\tau_x/2} = i\tau_x$). This discrete symmetry breaking is not forbidden in two dimensions and the transition is likely of Ising type—although we have not carried out a scaling collapse to confirm this explicitly, as it was not the central focus of Ref. [P2].

Finite temperature phase diagram for the $J_2 - J'_2$ Hamiltonian Finally, we also calculated the finite temperature phase diagram including the couplings J'_2 at $J_2 > 0$ and $J_3 = \varphi = 0$, illustrated in Fig. 7.3(d–f). At $T = 0$ this perfectly reproduces the mean-field result: for $J'_2 < 0$ both valleys show the same FM order, indicated by the vanishing chirality in both valleys and no valley

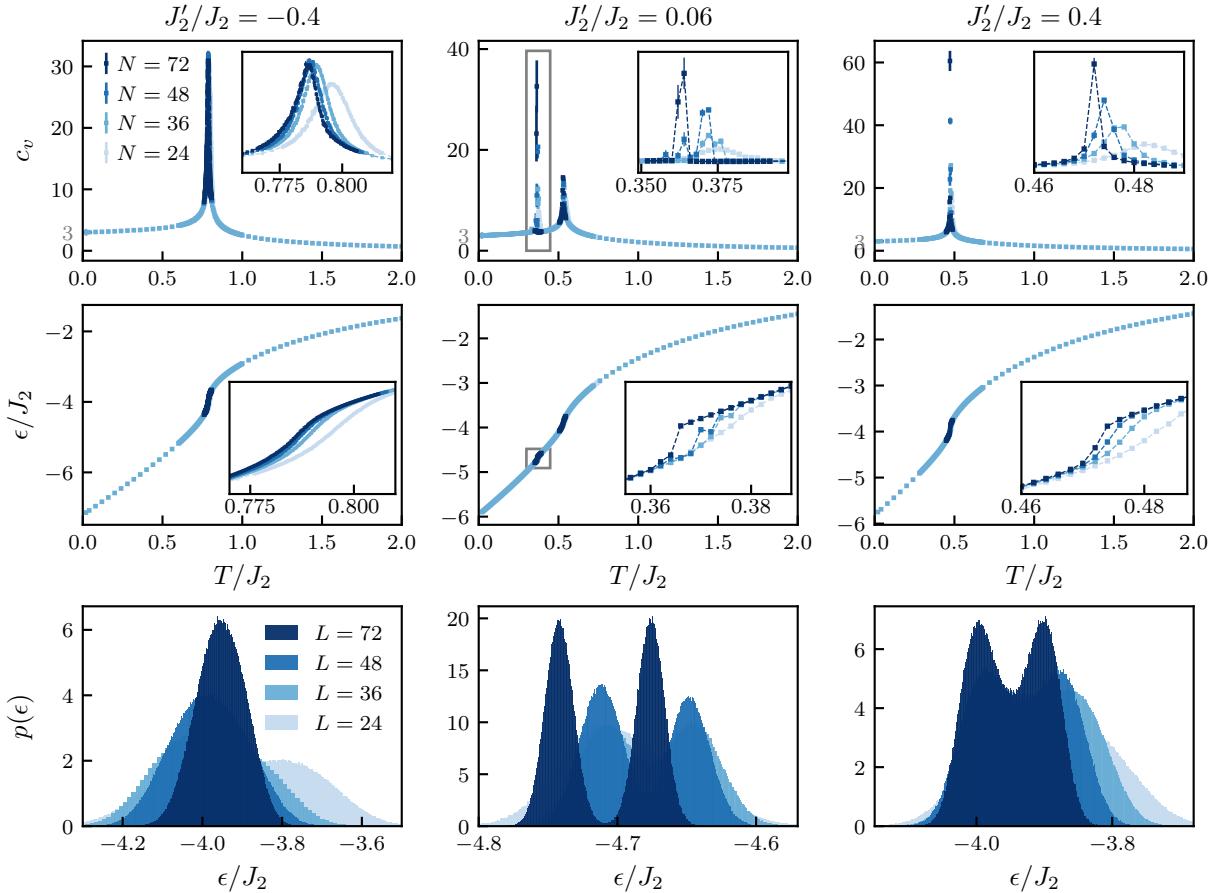
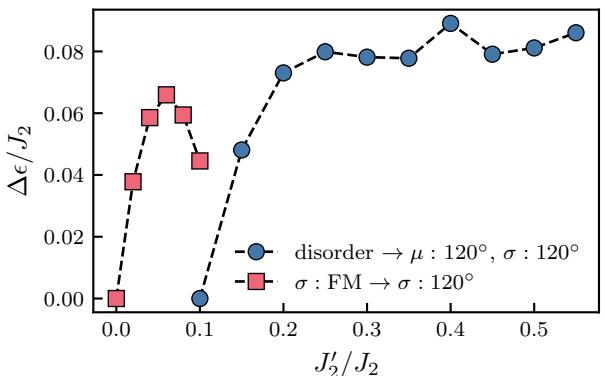


Figure 7.4 – Thermal phase transitions in the $J_2-J'_2$ model. Monte Carlo results for the specific heat c_v , energy per site ϵ , and energy distribution $p(\epsilon)$ at the critical temperature (from histogram reweighting) for different system sizes L . For $J'_2/J_2 = -0.4$ (left), the system transitions from the disordered state into ferromagnetic order both σP^\pm with finite chirality, consistent with a continuous or weakly first-order transition. For $J'_2/J_2 = 0.06$ (middle), an additional transition occurs between ferromagnetic and 120° order in one valley (highlighted by the gray rectangle and shown in the inset). The kink in the energy and the bimodal energy distribution signal a first-order transition. For $J'_2/J_2 = 0.4$ (right), the system directly transitions from the disordered state into 120° order, again with clear signatures of a first-order transition.

Figure 7.5 –

Latent heat of first-order transitions. The latent heat $\delta\epsilon$ is extracted from Monte Carlo data at $L = 72$ by fitting double Gaussians to the energy distribution at the critical point and taking the peak separation. Red squares denote the transition between two ordered states, where σP^\pm changes from FM to 120° order in one valley. Blue circles indicate the direct transition from the disordered state into the same 120° ordered phase.



polarization $|\tau^z| = 0$ (since $0 \langle \boldsymbol{\sigma} P^+ - \boldsymbol{\sigma} P^- \rangle = \langle \boldsymbol{\sigma} \tau^z \rangle = 0$). For $J'_2 > 0$, one of the valleys instead shows coplanar 120° order (also with zero chirality), which breaks the symmetry between $+$ and $-$ valleys indicated by a finite valley polarization $|\langle \tau^z \rangle| > 0$. Interestingly, we find a finite temperature phase with finite chirality as for H_0 .

This phase occurs whenever the ground state has ferromagnetic spin order in both valleys. It also extends into a region with $J'_2 > 0$, where upon lowering the temperature the system first undergoes a transition into the chiral phase and subsequently into the 120° phase. The sequence of transitions reflects the competition between two mechanisms: thermal fluctuations, which favor the chiral state via the order-by-disorder effect discussed above (scaling linearly with T), and the J'_2 interaction, which stabilizes the 120° state. A spin-wave calculation yields an analytic estimate for the critical temperature, $T_c \approx 8.21 J'_2$, below which the J'_2 interaction dominates over thermal fluctuations. This estimate agrees well with the Monte Carlo results at low temperature, as indicated by the white dashed line in Fig. 7.3(e).

Nature of the phase transitions We also analyzed the nature of the phase transitions by calculating the specific heat and energy as functions of temperature, as well as the energy distribution at the critical points, as shown in Fig. 7.3. As already suggested for H_0 , the transition into the finite-temperature chiral phase appears to be continuous, or at most very weakly first order: the energy shows no sharp kink, and the energy distribution exhibits only a single peak (left column). By contrast, the transition into the 120° state displays a pronounced kink in the energy, and the probability distribution at the critical point (obtained by histogram reweighting [95]) develops two well-separated peaks that sharpen with increasing system size. These are strong indicators of phase coexistence and a first-order transition. This applies both to the transition from the chiral state to 120° order (middle column), and the direct transition from the paramagnetic state to 120° order (right column).

We further extracted the latent heat, i.e., the energy required to traverse the first-order transition. For this purpose we fitted double Gaussians to the energy distributions to determine the distance between the two peaks. The results, shown in Fig. 7.5, demonstrate that both transitions exhibit a sizable latent heat, providing additional evidence for their strongly first-order character.

In all phases, the specific heat saturates to $c_v(T \rightarrow 0) = 3$. This corresponds to the presence of six harmonic modes, each contributing $1/2$ according to the equipartition theorem [17, 225]. Such a mode count is expected for an $SU(4)$ model at quarter filling: the local Hilbert space is described by a four-component complex vector, which is parametrized by eight real numbers. After accounting for normalization and an arbitrary global phase, the number of independent degrees of freedom is reduced to $8 - 2 = 6$. By contrast, a classical $SU(2)$ ferromagnet (described by an $O(3)$ vector) has only two harmonic modes.

7.1.4 Discussion

Our analysis reveals that, already at the semiclassical level, truly exotic states can emerge in a moiré system that at first sight appears relatively simple. Among these is the peculiar coplanar 120° order in the mixed spin-valley degrees of freedom $\boldsymbol{\mu}_i^a$, which coexists with ferromagnetic order of the valley-projected spin $P^\pm \boldsymbol{\sigma}$ in one valley, and either ferromagnetic, coplanar 120° , or chiral noncoplanar order in the other valley. The most striking observation is that chiral noncoplanar order is not only stabilized by explicit chiral interactions, but can also arise purely through thermal fluctuations via an order-by-disorder mechanism. This mechanism is captured by higher-order $SU(4)$ spin-wave calculations, whose predictions agree remarkably well with our

Monte Carlo simulations at low temperatures. Furthermore, our zero-temperature phase diagram is in full agreement with the mean-field analysis, providing a strong benchmark for the validity of our implementation.

For the concrete realization of this model, we assumed a hexagonal substrate with a lattice spacing comparable to that of graphene. Our coauthors of Ref. [P2], however, have shown that substrates with lower symmetry, such as h-BN, can also approximately realize similar physics. In particular, they demonstrated that twisted monolayer–bilayer graphene, where the bilayer acts as an effective substrate, is an excellent candidate to realize the model studied here. Moreover, they found that the effective interactions can be tuned by applying a gate voltage. This is exemplified by the yellow line in Fig. 7.2(b), which shows that gating can drive the system through all the ordered phases we identified, including the noncoplanar chiral state. These findings suggest that the phenomena predicted by our model may indeed be experimentally accessible in currently available moiré heterostructures.

An open question is how quantum fluctuations modify the phase diagram, which would require extending our pf-FRG approach to quarter filling. Indeed, Ref. [P2] already demonstrated within first-order spin-wave theory that, for H_0 , quantum fluctuations select the 120° state rather than the ferromagnet—providing a rare example in which thermal and quantum fluctuations stabilize different states. Exploring the interplay between these effects, and in particular whether the finite-temperature chiral phase persists, would be highly interesting. One possible route would be to investigate the phase diagram using the recently developed finite-temperature pseudo-Majorana FRG [140, 141]. However, accurately capturing the noncoplanar nature of the chiral state requires access to a three-spin order parameter, which remains a significant challenge for FRG approaches and most numerical methods applicable to such frustrated magnets.

Another natural question is how such peculiar spin–valley order can be detected experimentally. In the simplest case of valley polarization (finite $\langle \tau^z \rangle$), electrons preferentially occupy one valley over the other, while the accompanying spin order could be probed with standard techniques for magnetic materials, such as neutron scattering. When the valley is not fully polarized and the system instead exhibits a finite $\langle \tau^x \rangle$ or $\langle \tau^y \rangle$, the valley pseudospin mixes the \mathbf{K} and \mathbf{K}' valleys, resulting in intervalley-coherent (IVC) order. Since these momenta differ by a reciprocal lattice vector, IVC order may induce charge-density modulations at the scale of the graphene lattice, which can in principle be detected by scanning tunneling microscopy (STM) [275–277]. Finally, chiral order in the spin sector may be accessible through optical probes [278].

7.2 Spin-valley magnetism in TG/h-BN

We now turn to the study of a moiré system with multiple graphene layers. In particular, we focus on ABC-stacked trilayer graphene (TG) aligned with h-BN. In ABC stacking, the three graphene layers are arranged such that different sublattices overlap in each layer [see Fig. 7.6(a)]. When the top graphene layer is aligned with h-BN, whose lattice constant is slightly larger than that of graphene, a triangular moiré superlattice is formed [Fig. 7.6(b)]. In TG/h-BN, the electronic properties can be tuned not only by controlling the filling via a gate voltage, but also by applying a perpendicular displacement field D , which shifts the relative potential Δ_V between the top and bottom graphene layers. For sufficiently large D and at integer fillings $\nu = -1, -2$ (corresponding to one or two holes per moiré unit cell), correlated insulating states have been observed [91, 264]. Depending on the sign of D and the resulting potential offset Δ_V , the system can host either topologically trivial bands ($\Delta_V < 0$) or bands with a finite Chern number ($\Delta_V > 0$) [49, 91, 264, 279]. In the following, we restrict ourselves to the topologically trivial side, where effective

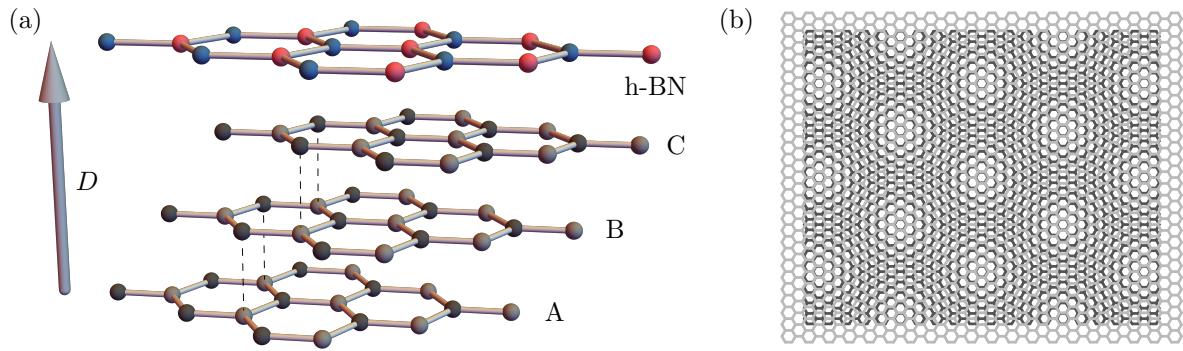


Figure 7.6 – ABC trilayer graphene aligned with hexagonal boron nitride. (a) Illustration of the ABC stacking sequence, in which different A/B sublattices overlap at each layer. The top layer is aligned with h-BN, and a perpendicular displacement field D can tune the potential between the top and bottom graphene layers. (b) Moiré pattern formed between the top graphene layer and the aligned h-BN, induced by the slightly larger lattice constant of h-BN compared to graphene.

lattice models can be derived, and to a filling of $\nu = -2$, where our pf-FRG implementation can be applied.

In this regime, the nature of the insulating states in TG/h-BN, are still under debate. In Ref. [49] they predicted a ferromagnet for large negative D , and possibly a quantum spin-valley liquid closer to the Mott transition. More recent work using single-site DMFT proposes an antiferromagnetic order that breaks C_3 symmetry. This is consistent with spectroscopy measurements that rule out a ferromagnetic ground state, but are consistent antiferromagnetic or intervalley-coherent (IVC) order.

To try to shed some light on this issue, we study the effective model of Ref. [49]. Starting from band structure calculations around the \mathbf{K} and \mathbf{K}' points, they derived a triangular-superlattice Hubbard model for holes carrying both spin and valley quantum numbers. Tuning the displacement field to negative values drives a metal-to-Mott-insulator transition. In the Mott regime, a strong-coupling expansion yields an effective spin–valley Hamiltonian containing SU(4)-symmetric interactions, supplemented by additional valley XXZ- and Dzyaloshinskii–Moriya–type terms and an on-site Hund’s coupling. Together, these strongly break the SU(4) symmetry down to $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$.

After giving a brief description of the derivation of this model, we begin by analyzing its phase diagram with parameters concretely estimated for TG/h-BN in Ref. [49], using our combined semiclassical and pf-FRG implementation to compute the phase diagram as a function of displacement field. However, since closely related models arise in other moiré systems (see for example the previous section or twisted bilayer graphene [48]), and since Ref. [49] themselves emphasize that parameter estimates remain uncertain, we also explore the broader coupling space, with particular focus on understanding the role of different SU(4)-breaking interactions.

The results in this section are based in Ref. [P1], from which most figures in this section are redrawn. The underlying data were all obtained by the author of this thesis.

7.2.1 Derivation of the spin-valley model

To derive an effective lattice model for TG/h-BN, Ref. [49] follow a similar strategy as described for the moiré system in the previous section. They start from the continuum description of a single graphene layer obtained by expanding around the two valleys at \mathbf{K} and \mathbf{K}' (the massless Dirac Hamiltonian in Eq. (7.1)). Interlayer hoppings between the graphene sheets are then

included, together with a potential offset Δ_V between the top and bottom layer, which can be tuned experimentally by a perpendicular displacement field D [see Fig. 7.6]. Alignment of the top graphene layer with the h-BN substrate generates the triangular moiré potential as discussed previously. Diagonalization of the resulting Hamiltonian in momentum space yields an effective band structure of TG/h-BN. The band width—which becomes particularly narrow for large negative Δ_V —can be tuned by the displacement field. For $\Delta_V > 0$, the lowest moiré bands acquire a finite Chern number. While this topological regime is of great interest, it precludes the construction of exponentially localized Wannier orbitals and hence of a simple real-space lattice model [280]. In contrast, for $\Delta_V < 0$ the bands are topologically trivial, allowing the construction of Wannier orbitals and an effective tight-binding model on the triangular moiré superlattice [279, 280].

Starting from such a tight-binding description, by projecting the Coulomb interaction onto the valence band, Ref. [49] derive an extended spin–valley Hubbard model containing both on-site and nearest-neighbor Coulomb repulsion, as well as on-site and inter-site Hund’s couplings that couple different valleys. Based on this model, they predict that tuning the displacement field to negative values drives a Mott transition: for a potential offset $\Delta_V < -20$ meV, the system evolves from a metallic state into a correlated insulator. Experimental evidence for such a transition was reported in Ref. [264], where at integer fillings $\nu = -1, -2$ per moiré unit cell and large negative displacement field the system indeed exhibits correlated insulating states.

As our focus is on the insulating states, we consider the effective spin-valley model derived in Ref. [49] via a strong coupling expansion of the Hubbard model (as outlined for a conventional Hubbard model in Sec. 2.2.3). The resulting Hamiltonian is

$$\begin{aligned} H = & \frac{J_1}{8} \sum_{\langle ij \rangle} (1 + \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j) (1 + \boldsymbol{\tau}_i \boldsymbol{\tau}_j) + \frac{J_2}{8} \sum_{\langle\langle ij \rangle\rangle} (1 + \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j) (1 + \boldsymbol{\tau}_i \boldsymbol{\tau}_j) \\ & + \frac{1}{8} \sum_{\langle ij \rangle} J_{ij}^{\text{XXZ}} (1 + \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j) (\tau_i^x \tau_j^x + \tau_i^y \tau_j^y) + \frac{1}{8} \sum_{\langle ij \rangle} J_{ij}^{\text{DM}} (1 + \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j) (\tau_i^x \tau_j^y - \tau_i^y \tau_j^x) \\ & - \frac{J_H}{4} \sum_i (n_{+i} n_{-i} + \boldsymbol{\sigma}_{+i} \boldsymbol{\sigma}_{-i}) , \end{aligned} \quad (7.9)$$

where $\langle \cdot \rangle$ and $\langle\langle \cdot \rangle\rangle$ denote summation over nearest- and next-nearest-neighbor bonds of the triangular moiré superlattice, respectively. The operators $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$, as well as their on-site products, represent the combined spin–valley operators defined in Eq. (3.39). The interactions proportional to J_1 and J_2 are SU(4)-symmetric nearest- and next-nearest-neighbor Heisenberg couplings. The nearest-neighbor exchange J_1 receives two competing contributions: a ferromagnetic part from a Hund’s-type direct exchange, and an antiferromagnetic part from superexchange processes (see Sec. 2.2.3 for an intuitive discussion). Depending on the microscopic parameters and the applied displacement field, J_1 may therefore be either positive or negative. By contrast, the next-nearest-neighbor exchange J_2 arises solely from superexchange and is thus expected to be antiferromagnetic ($J_2 > 0$).

The interaction terms proportional to J_{ij}^{XXZ} and J_{ij}^{DM} break the SU(4) symmetry down to $\text{SU}(2)_{\text{spin}} \otimes \text{U}(1)_{\text{valley}}$ and also originate primarily from superexchange processes. In TG/h-BN, these couplings are not independent of J_1 and acquire a characteristic bond dependence,

$$J_{ij}^{\text{XXZ}} = (J_1 + K) [\cos(2\varphi_{ij}) - 1] + J'_H , \quad J_{ij}^{\text{DM}} = (J_1 + K) \sin(2\varphi_{ij}) , \quad (7.10)$$

where K and J'_H denote contributions of different Hund’s couplings. The phase φ_{ij} originates from a complex nearest-neighbor hopping amplitude $t_{ij}^1 = |t_1| e^{i\varphi_{ij}}$ within a given valley and

changes sign for bonds related by a C_6 rotation. This induces a valley-contrasting flux of magnitude $|\phi| \equiv 3|\varphi_{ij}| \equiv 3\varphi$, which is staggered between left- and right-pointing triangles. Consequently, the associated exchange interactions break not only the internal SU(4) symmetry but also the C_6 symmetry of the triangular lattice down to C_3 .

In addition to the nearest- and next-nearest-neighbor interactions, the Hamiltonian contains an on-site Hund's coupling J_H . Here, $n_{\pm i} = f_{i\pm l}^\dagger f_{i\pm l}$ and $\sigma_{\pm i}^\mu = f_{is\pm}^\dagger \theta_{ss'}^\mu f_{is'\pm}$ denote the density and spin operators in the \pm valley sectors, respectively. This interaction can be rewritten, up to a constant shift, in terms of the spin-valley operators defined in Eq. (3.39) as

$$n_{+i}n_{-i} + \boldsymbol{\sigma}_{+i}\boldsymbol{\sigma}_{-i} = \frac{1}{4}(1 + \boldsymbol{\sigma}_i\boldsymbol{\sigma}_i)(1 - \tau_i^z\tau_i^z), \quad (7.11)$$

which makes explicit that the Hund's term also breaks the SU(4) symmetry to $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$.

The parameters J_1 , J_2 , and the flux φ are tunable by the displacement field. By contrast, the couplings K , J_H and J'_H depend only weakly on D , and we use the constant values $K = 0.4 \text{ meV}$, $J_H = 0.136 \text{ meV}$ and $J'_H = 0.05 \text{ meV}$ from Ref. [49] throughout. The resulting exchange couplings J_1 , J_2 , J_{ij}^{XXZ} , and J_{ij}^{DM} as functions of the induced potential difference Δ_V are shown in Fig. 7.7(a). All couplings acquire sizable values of similar magnitude somewhere in this phase diagram. While J_2 remains antiferromagnetic and J_{ij}^{XXZ} ferromagnetic across the full range, both J_1 and J_{ij}^{DM} change sign. This illustrates that the competition between the various couplings renders the spin-valley ground state in realistic models highly nontrivial. In the following, we aim to clarify the role of the different couplings through our analysis.

7.2.2 Ordered states at half-filling

Before we discuss the precise phase diagram, let us shortly consider which type of ordered states can even be expected in such a model. In the following, we restrict our analysis to a filling of $\nu = -2$ holes per moiré unit cell, which can be directly treated within our pf-FRG implementation. This corresponds to half-filling of the four-fold degenerate valence band and realizes the six-dimensional representation of $\mathfrak{su}(4)$, as discussed in detail in Sec. 3.2.

In principle, long-range order could develop in any of the 15 generators σ^a , τ^b , and $\sigma^a\tau^b$. Since the Hamiltonian is invariant under global rotations generated by $\boldsymbol{\sigma}$, τ^z , and $\boldsymbol{\sigma}\tau^z$, states related by these transformations are degenerate. It is therefore sufficient to group the generators into the following symmetry-distinct sectors:

$$\boldsymbol{\sigma} = (\sigma^x, \sigma^y, \sigma^z), \quad (7.12)$$

$$\boldsymbol{\tau}^\perp = (\tau^x, \tau^y), \quad (7.13)$$

$$\tau^z = (\tau^z), \quad (7.14)$$

$$\boldsymbol{\sigma}\boldsymbol{\tau}^\perp = (\sigma^x\tau^x, \sigma^y\tau^x, \sigma^z\tau^x, \sigma^x\tau^y, \sigma^y\tau^y, \sigma^z\tau^y), \quad (7.15)$$

$$\boldsymbol{\sigma}\tau^z = (\sigma^x\tau^z, \sigma^y\tau^z, \sigma^z\tau^z). \quad (7.16)$$

With the normalization of the spin-valley operators in Eq. (3.39), each generator at half-filling has eigenvalues -2 , 0 , and 2 . This can be seen, for example, by considering $\sigma^z = \sum_{s,s',l} f_{isl}^\dagger \theta_{ss'}^z f_{is'l}$, which reduces to $n_{\uparrow} - n_{\downarrow}$, where $n_s = \sum_l f_{isl}^\dagger f_{isl}$ counts the number of electrons with spin s . Two spin-up (down) electrons yield $\langle \sigma^z \rangle = \pm 2$, while one up and one down electron give $\langle \sigma^z \rangle = 0$. Importantly, if a state maximizes the magnetization in one generator, e.g. $|\langle \sigma^z \rangle| = 2$, then the expectation values of all other generators necessarily vanish. This implies that a conventional spin ferromagnet cannot coexist with a fully valley-polarized state. This situation is qualitatively

different from the quarter-filled case (one electron per site) discussed in the previous section, where different sectors can reach maximal polarization simultaneously, allowing, for example, full spin and valley polarization at once.

Our model retains a continuous $SU(2)_{\text{spin}}$ symmetry, the spontaneous breaking of which is forbidden in two dimensions by the Mermin–Wagner theorem [281]. We therefore expect ordering to occur primarily in the valley or spin–valley sectors. A finite expectation value of τ^z or $\sigma\tau^z$ corresponds to a valley-polarized state, breaking a discrete Ising-like \mathbb{Z}_2 symmetry. In contrast, ordering in τ^\perp or $\sigma\tau^\perp$ signals the breaking of the continuous $U(1)_{\text{valley}}$ symmetry, which may occur via a Berezinskii–Kosterlitz–Thouless (BKT) transition [282, 283]. Such order corresponds to inter-valley coherence (IVC), since it involves superpositions of states from different valleys. A state with only $\langle\tau^\perp\rangle \neq 0$ represents a pure IVC phase without spin polarization, while finite $\langle\sigma\tau^\perp\rangle$ indicates a mixed spin–valley IVC order.

Beyond simple ferromagnetic order, where the expectation values of the generators are uniform across all sites, the system may also realize states with spatially modulated order parameters, such as 120° order or even incommensurate (ICS) patterns. Such modulations can occur within a single sector, or more generally in the full $SU(4)$ spin–valley space. Throughout, we label ordered states by the sector in which the order develops (e.g., $\sigma\tau^\perp$) together with the type of spatial modulation (e.g., 120°). The explicit form of such states will be discussed in more detail below. When more than one sector develops finite expectation values simultaneously—specifically the combination $\sigma\tau^\perp$, τ^\perp , and σ —we refer to the resulting state as a *mixed order*.

We note that the eigenstates of different generators are not fully linearly independent. For instance, we find that finite expectation values of $\langle\sigma\tau^z\rangle$ and $\langle\tau^z\rangle$ automatically imply a finite σ . While I don't understand the precise relations between different spin-valley operators, the existence of such constraints can be anticipated: the local Hilbert space at half filling is six-dimensional, so local states are described by six-component complex vectors. After accounting for normalization and an arbitrary phase, this leaves only ten independent degrees of freedom. Consequently, not all fifteen generators of $SU(4)$ can vary independently. In this regard, $SU(2)$ in its $S = 1/2$ representation is a special case, since the number of local degrees of freedom and the number of generators both equal three, which also implies that the semiclassical limit coincides with the usual $S \rightarrow \infty$ limit.

7.2.3 Key results

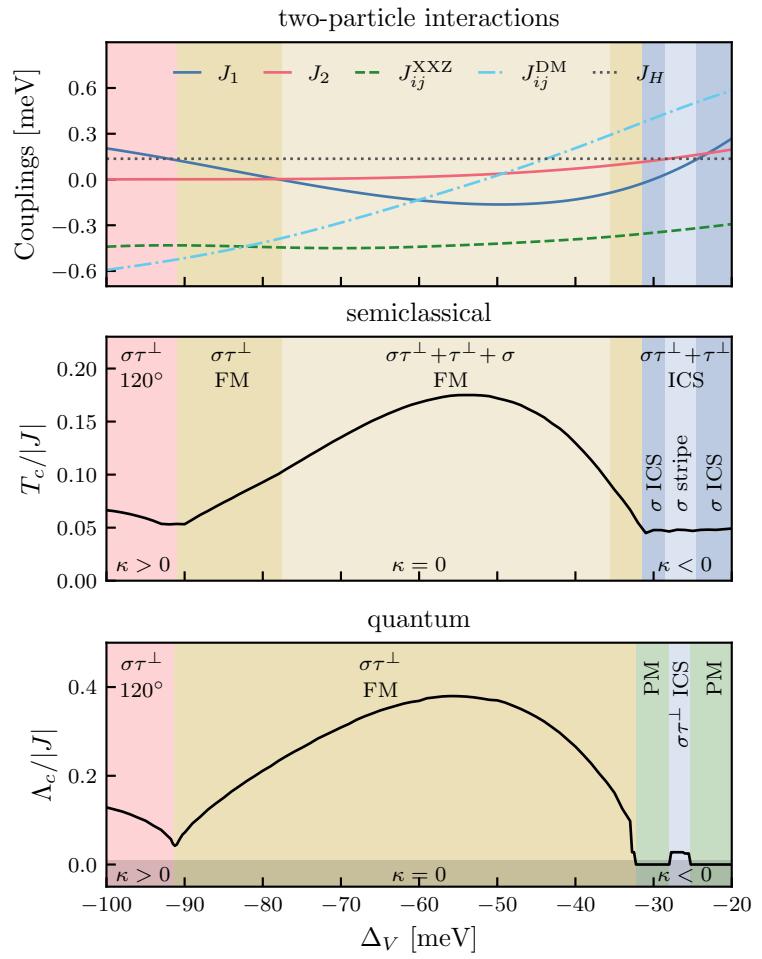
We begin by summarizing the key results for the approximate phase diagram of TG/h-BN obtained from our semiclassical Monte Carlo simulations and the pf-FRG. A more detailed account of the underlying calculations is presented in the following sections.

The semiclassical and quantum phase diagrams for the estimated parameters of TG/h-BN, as a function of the potential difference Δ_V , are shown in Fig. 7.7. We consistently find *only* intervalley-coherent (IVC) order, while valley-polarized states are absent ($\langle\tau^z\rangle = \langle\sigma\tau^z\rangle = 0$). This is natural, since the dominant $SU(4)$ symmetry-breaking interactions, J_{ij}^{XXZ} and J_{ij}^{DM} , act exclusively on the in-plane valley components τ^\perp . For large negative Δ_V , where the valley DM interaction J_{ij}^{DM} is strongest, the $\sigma\tau^\perp$ sector develops ideal 120° order. At intermediate Δ_V , ferromagnetic J_1 and J_{ij}^{XXZ} dominate, leading first to FM order in $\sigma\tau^\perp$, and subsequently to mixed FM order. In these regimes, the pf-FRG results are broadly consistent with the semiclassical findings, with the exception that pf-FRG typically captures only the dominant ordering channel and thus does not resolve the full mixed order.

Near the Mott transition at $\Delta_V \approx -20$ meV, both J_1 and J_2 become positive and comparable in magnitude. In the semiclassical limit this drives the system into noncollinear, mixed incom-

Figure 7.7 –

Phase diagram of the TG/h-BN model as a function of the potential difference Δ_V . (a) Couplings in Hamiltonian (7.9) as estimated for TG/h-BN [49]. The model is expected to undergo a Mott transition at $\Delta_V = -20$ meV. (b) Transition temperature T_c (obtained by fitting specific heat peaks) and ground-state order from semiclassical Monte Carlo calculations. (c) Critical scale and quantum ground states from pf-FRG. The labels denote the spin-valley operator (divided into symmetry-equivalent sectors defined in Eq. (7.16)) that develops a finite local expectation value, together with the corresponding spatial order. All ordered states break the valley U(1) symmetry and are therefore a type of intervalley coherent order, forming 120° , ferromagnetic (FM), or incommensurate (ICS) patterns. Around $\Delta_V = -27$ meV, the spin sector additionally exhibits stripe order in the semiclassical limit. In the quantum phase diagram, semiclassical ICS order always melts into a quantum paramagnetic ground state (PM). The coplanar state regions acquire a finite vector chirality κ , whose sign is tied to the sign of J_{ij}^{DM} .



mensurate order, with an additional small stripe-ordered region in the spin sector (possibly a finite-size artifact). This scenario is consistent with single-site DMFT calculations [279], which predict an AFM state breaking C_3 symmetry near the transition. When quantum fluctuations are included, however, the semiclassical ICS regime is destabilized and melts into a paramagnetic state with no spin or valley order, opening the possibility of a quantum spin-valley liquid close to the Mott transition, as tentatively proposed in Ref. [49]. The strong-coupling expansion is most reliable at large $|\Delta_V|$. Closer to the Mott transition, higher-order corrections may become important and could in principle stabilize other phases. Nevertheless, quantum spin liquid behavior is also expected near Mott transitions in other strongly correlated systems [4, 284], suggesting that such higher-order processes might in fact help stabilize the disordered regime. Our results therefore strengthen the case for searching for spin-valley liquid states in correlated moiré heterostructures, even in the presence of substantial SU(4)-breaking interactions.

Figure 7.7 demonstrates that the nature of the ground-state order is highly sensitive to the SU(4) symmetry-breaking couplings. A dominant ferromagnetic $J_{ij}^{\text{XXZ}} < 0$ favors collinear ferromagnetic IVC order, whereas large J_{ij}^{DM} stabilizes noncollinear IVC states with a finite *vector chirality* κ [defined in Eqs. (7.19) and (7.20)], such as 120° or ICS spiral order. The sign of the

chirality is strictly determined by the sign of J_{ij}^{DM} . Which type of noncollinear order emerges depends on the ratio J_2/J_1 : small values favor 120° order, while larger ratios stabilize ICS states. We confirm this behavior by systematically varying the SU(4)-breaking couplings at fixed J_1 and J_2 , as shown in Fig. 7.11. In the quantum limit the results agree well with the semiclassical analysis, with the important exception of the ICS regime, where quantum fluctuations consistently melt the order into a quantum paramagnetic phase.

7.2.4 Semiclassical analysis

We now present a detailed discussion of the ground states and the role of thermal fluctuations obtained from the semiclassical Monte Carlo simulations introduced in Sec. 3.2. Our typical simulations are performed on lattices of up to $N = 36^2$ sites with periodic boundary conditions. For each temperature, we use up to $N_m = 4 \cdot 10^6$ Monte Carlo sweeps, increasing to $N_m = 20 \cdot 10^6$ sweeps in the vicinity of phase transitions. To accurately determine the semiclassical ground state, we employ simulated annealing combined with stochastic gradient descent as also detailed in Sec. 3.2.

Ferromagnetic states A large region of the TG/h-BN phase diagram around $\Delta_V \approx -60$ meV exhibits ferromagnetic (FM) order. Interestingly, the character of this FM order changes between the boundaries and the center of the region. At the boundaries (approximately $-91 \lesssim \Delta_V \lesssim -77$ meV and $-36 \lesssim \Delta_V \lesssim -31$ meV), the ground state is an eigenstate of either τ^\perp or $\sigma\tau^\perp$. While these states are degenerate in the bare Hamiltonian, thermal fluctuations lift the degeneracy via an order-by-disorder mechanism, selecting $\sigma\tau^\perp$ order. The corresponding magnetization

$$|\langle \sigma\tau^\perp \rangle| = \frac{1}{N} \left| \sum_i \langle \sigma\tau^\perp \rangle \right|, \quad (7.17)$$

reaches its maximal value of 2 as $T \rightarrow 0$, while all other sectors vanish. This can be shown in Fig. 7.8(a) which shows the magnetization in all three relevant sectors. This means the ground state is simply an eigenstate of $\sigma\tau^\perp$ on all sites.

In the center of the FM region, however, a *mixed* FM state appears, where several sectors acquire finite magnetization simultaneously. Up to symmetry transformations, the ground state can be written as

$$|\Psi\rangle_{\text{FM}}^{\text{mixed}} \sim |\sigma^x\tau^x\rangle + |\tau^x\rangle + \delta |\sigma^x\rangle, \quad (7.18)$$

where $|\sigma^\mu\tau^\nu\rangle$ denotes the eigenstate of $\sigma^\mu\tau^\nu$ with eigenvalue +2. The parameter δ , which depends on Δ_V , lies in the interval $\delta \in [0.455, 0.538]$ meV. This value was determined by matching the analytic energy of $|\Psi\rangle_{\text{FM}}^{\text{mixed}}$ to the numerical minimization, with perfect agreement [see Fig. 7.8(c)]. The same figure also clearly reveals the regimes where mixed FM order is energetically favorable to pure $\sigma\tau^\perp$ FM order.

We note that $|\sigma^x\tau^x\rangle$, $|\tau^x\rangle$, and $|\sigma^x\rangle$ are not linearly independent, and their linear combinations can be written in different forms. Notably, even the simpler combination $|\sigma^x\tau^x\rangle + |\tau^x\rangle$ yields a finite spin expectation value $\langle \sigma^x \rangle$, which lowers the energy of the on-site Hund's coupling $\sim J_H$. This provides a natural explanation for the emergence of the mixed order.

States with a finite vector chirality Outside the FM regions, all one-sublattice magnetizations vanish strictly [see Fig. 7.8(a)], so different order parameters are needed to characterize the ground states. In these regimes J_{ij}^{DM} is large, favoring coplanar spin-valley configurations in

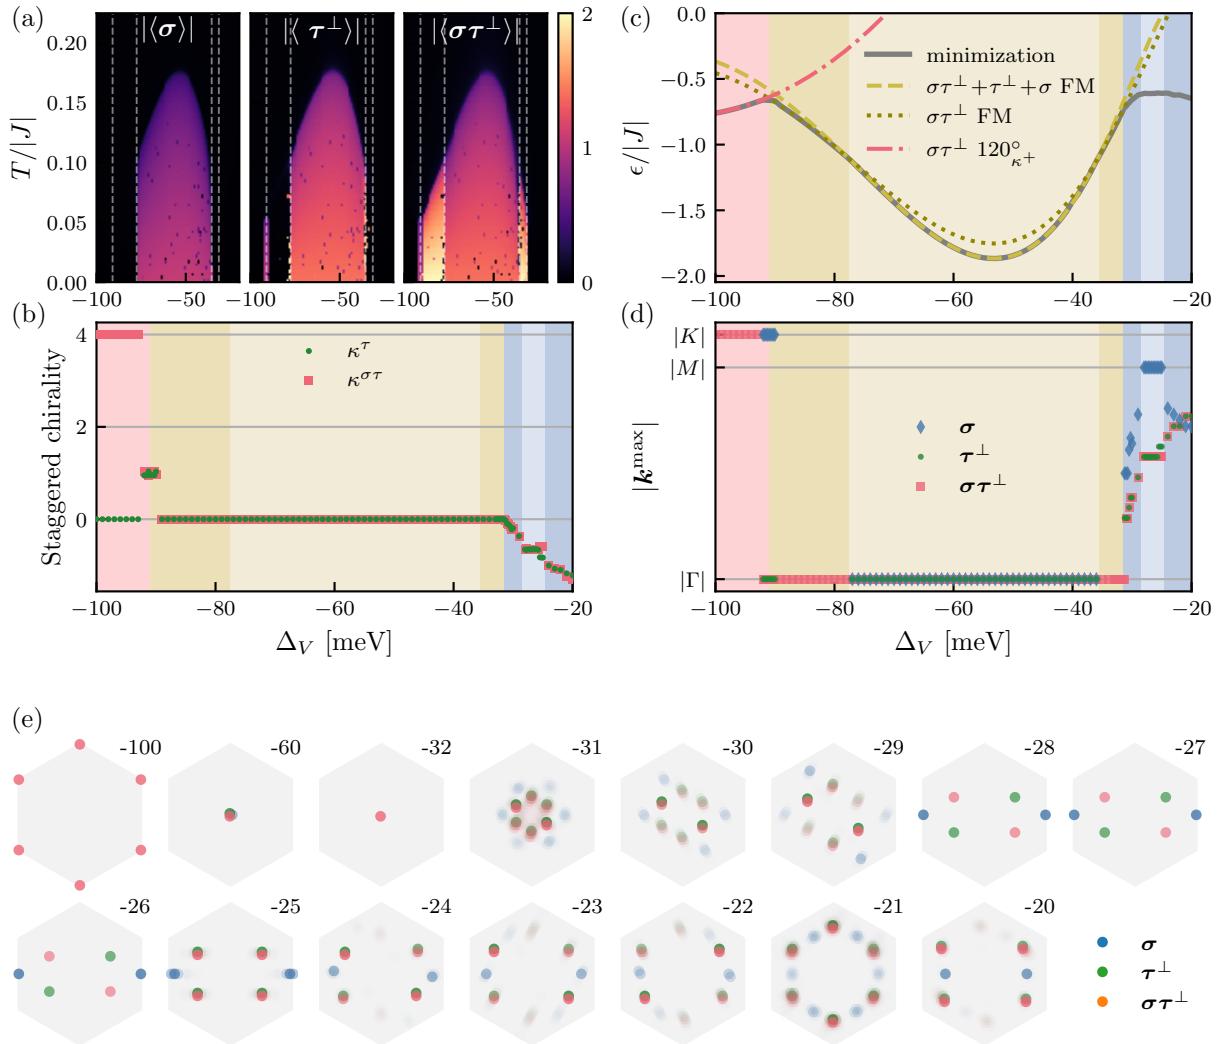


Figure 7.8 – Semiclassical phases and observables of TG/h-BN. (a) Average magnetization in the three relevant spin-valley sectors as a function of temperature and Δ_V . (b) Staggered vector chirality, as defined in Eq. (7.20), characterizing coplanar order. (c) Ground-state energy per site ϵ obtained from numerical minimization (gray lines) compared to analytical values for several candidate states. (d) Magnitude of the momentum \mathbf{k}^{\max} at which the structure factor is maximal for the nonvanishing sectors. (e) Full structure factor of the different sectors for several values of Δ_V indicated in the top right. Overlapping points are slightly shifted for visibility.

which local expectation values rotate from site to site, thereby producing a finite chirality. The canonical example of such an arrangement is the 120° state. To quantify these coplanar states, we compute the z -component of the *staggered vector chirality*, defined for a general three-component vector \mathbf{v} as

$$\kappa(\mathbf{v}) = \frac{1}{3\sqrt{3}N} \sum_{p=\nabla/\Delta} (-1)^p (\mathbf{v}_{p_1} \times \mathbf{v}_{p_2} + \mathbf{v}_{p_2} \times \mathbf{v}_{p_3} + \mathbf{v}_{p_3} \times \mathbf{v}_{p_1})^z, \quad (7.19)$$

where the sum runs over all up- and down-pointing triangles, with (p_1, p_2, p_3) labeling their sites in counterclockwise order, and the factor $(-1)^p$ alternates between up and down triangles. For normalized vectors $|\mathbf{v}| = 2$, this definition yields the maximal value $\kappa(\mathbf{v}) = \pm 4$ for ideal 120° order in the xy plane, while collinear or disordered states give $\kappa(\mathbf{v}) = 0$.

Based on this measure, we define the sector-resolved chiralities as

$$\kappa^\tau \equiv \kappa(\langle \boldsymbol{\tau}^\perp \rangle), \quad \kappa^{\sigma\tau} \equiv \kappa(\langle \sigma^x \boldsymbol{\tau}^\perp \rangle) + \kappa(\langle \sigma^y \boldsymbol{\tau}^\perp \rangle) + \kappa(\langle \sigma^z \boldsymbol{\tau}^\perp \rangle), \quad (7.20)$$

whose behavior as a function of Δ_V is shown in Fig. 7.8(b).

For $\Delta_V < -91$ meV, the ground state exhibits 120° order in the $\sigma\boldsymbol{\tau}^\perp$ sector with chirality $\kappa^{\sigma\tau} = +4$. An explicit expression of such a state can be obtained by starting from $|\sigma^x \tau^x\rangle$ and applying rotations in the valley xy -plane as

$$|\Psi_a\rangle_{120^\circ}^{\sigma\tau^\perp} = e^{-i\tau^z\theta^a} |\sigma^x \tau^x\rangle, \quad (7.21)$$

with angles $\theta^a = (0, 2\pi/3, 4\pi/3)$ assigned to the three sublattices of the 120° pattern. While eigenstates of $\boldsymbol{\tau}^\perp$ also belong to the classical ground-state manifold, thermal fluctuations again select the $\sigma\boldsymbol{\tau}^\perp$ states via order-by-disorder. The overall sign of the chirality is determined by the sign of J_{ij}^{DM} , which lifts the degeneracy between states of opposite chirality. The identification of the above state as the true ground state for $\Delta_V < -91$ meV is confirmed by comparing its energy to the numerically obtained ground-state energy, as shown in Fig. 7.8(c).

Closer to the Mott transition, for $\Delta_V > -31$ meV, the system also develops finite chirality, now in both $\boldsymbol{\tau}^\perp$ and $\sigma\boldsymbol{\tau}^\perp$. Unlike in the 120° phase, the chirality does not remain fixed but decreases continuously with increasing Δ_V , signaling the onset of a mixed incommensurate (ICS) order. This type of order is best characterized in momentum space. To this end, we compute the structure factors of the relevant sectors, defined as the Fourier transforms of the corresponding real-space correlations: $\langle \boldsymbol{\sigma}_i \boldsymbol{\sigma}_j \rangle$, $\langle \boldsymbol{\tau}_i^\perp \boldsymbol{\tau}_j^\perp \rangle$, and $\langle \boldsymbol{\sigma}_i \boldsymbol{\tau}_i^\perp \boldsymbol{\sigma}_j \boldsymbol{\tau}_j^\perp \rangle$. Representative results are shown in Fig. 7.8(e), while the dominant wave vector magnitudes $|\mathbf{k}^{\text{max}}|$ extracted from these structure factors are summarized in Fig. 7.8(d).

As expected, in the ferromagnetic phase the structure factor peaks at the Γ point, and in the 120° phase at the \mathbf{K} and \mathbf{K}' points. In the ICS regime, the ordering vectors shift continuously to incommensurate positions: between Γ and \mathbf{K} for $\boldsymbol{\tau}^\perp$ and $\sigma\boldsymbol{\tau}^\perp$, and between Γ and \mathbf{M} for $\boldsymbol{\sigma}$. Only in a narrow parameter window do peaks for $\boldsymbol{\sigma}$ appear exactly at the \mathbf{M} point, corresponding to stripe order.

Thermal fluctuations As already mentioned, thermal fluctuations preferentially select $\sigma\boldsymbol{\tau}^\perp$ over pure $\boldsymbol{\tau}^\perp$ order, despite their degeneracy at the classical level. To additionally assess the thermal stability of the different ordered phases, we computed the specific heat, energy traces, and energy distributions at the critical scale in the 120° , FM, and ICS regimes of the phase diagram. The analysis follows the procedure outlined in Sec. 7.1.3 and is discussed in more detail in Appendix D (see Fig. D.1 and Fig. D.2). In all regimes we find clear peaks in the specific heat that sharpen with increasing system size, suggesting the presence of thermal phase transitions rather than crossovers. In the 120° and FM regimes, no evidence of phase coexistence is seen in the energy distributions, consistent with continuous transitions, which in this case would likely be of BKT type. By contrast, at the transition into the ICS regime the energy shows a clear kink, and the energy distribution has a double peak structure, both strong indications of a first-order transition.

In every phase, the low-temperature specific heat saturates to $c_v(T \rightarrow 0) = 5$ (see Fig. D.3), corresponding to ten harmonic zero modes [17, 225], in agreement with the six-dimensional complex local Hilbert space, which is parametrized by ten independent real degrees of freedom as discussed above.

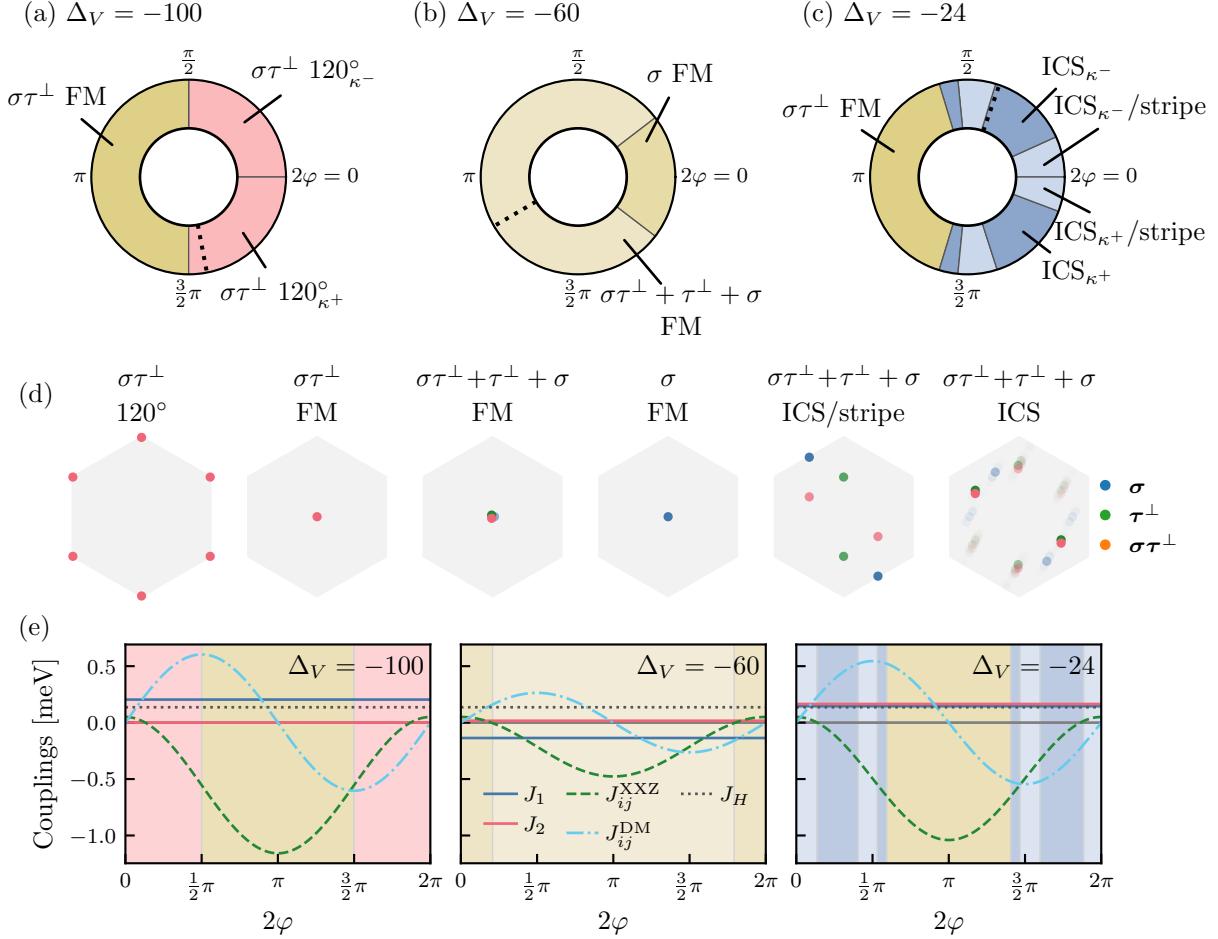


Figure 7.9 – Semiclassical phase diagram as a function of φ . (a–c) Phases and phase boundaries for J_1 and J_2 fixed to their TG/h-BN estimates at $\Delta_V = -100, -60, -24$, but with varying phase φ . Varying φ effectively tunes the couplings J_{ij}^{XXZ} and J_{ij}^{DM} according to Eq. (7.10) as shown in (e). The dotted black lines indicate the TG/h-BN estimates for φ at these values of Δ_V . κ^\pm denotes the sign of the chirality (when nonzero). All ICS phases exhibit order in all three sectors, while in the ICS/stripe phase the σ sector instead shows stripe order. (d) Ground-state (static) structure factors. Overlapping points are slightly shifted for visibility.

The role of the $\text{SU}(4)$ symmetry-breaking couplings Our analysis shows that $\text{SU}(4)$ -breaking couplings play a central role in determining the ground-state order. Moreover, as emphasized in Ref. [49], precise parameter estimates for TG/h-BN remain uncertain, which motivates exploring a broader parameter space beyond the specific estimates studied earlier. Such an extended analysis may also be of relevance for other moiré materials with similar physics.

To isolate the role of these interactions, we computed phase diagrams at fixed J_1 and J_2 while varying the phase φ_{ij} , which tunes the effective couplings J_{ij}^{XXZ} and J_{ij}^{DM} via Eq. (7.10). For J_1 and J_2 , we chose values estimated for TG/h-BN at $\Delta_V = -100, -60, -24$ meV [see Fig. 7.7(a)], which correspond to the three principal 120° , FM, and ICS/PM regimes, respectively. The resulting phase diagrams as a function of $2\varphi \in [0, 2\pi]$ are shown in Fig. 7.9, along with the precise values for the five couplings.

All three phase diagrams exhibit a similar trend. For $2\varphi \approx \pi$ (left side of the circles), J_{ij}^{XXZ} dominates and is negative, stabilizing collinear IVC order. In the region $2\varphi \in (-\pi/2, \pi/2)$

(right side of the circles) J_{ij}^{DM} becomes larger than J_{ij}^{XXZ} and chiral states are favored. The nature of the chiral states then depends on the values of J_1 and J_2 , where small J_2/J_1 (as for $\Delta_V = -100$ meV) prefers 120° and $J_2/J_1 1$ (as for $\Delta_V = -24$ meV) prefers ICS order.

At $\Delta_V = -100$ meV, the 120° state remains stable for $2\varphi \in (-\pi/2, \pi/2)$, while outside this range FM order takes over. At $\Delta_V = -60$ meV, FM order persists for all 2φ , but its character changes: while the original state is a mixed ferromagnet, around $2\varphi = 0$ a pure spin ferromagnet (σ) appears. The latter would be forbidden at finite temperatures by the Mermin–Wagner theorem, and we attribute the appearance of this phase to finite-size effects. At $\Delta_V = -24$ meV, the mixed ICS state is stable across a wide range of 2φ , though small regions of stripe order in σ again also appear. Here, FM order reemerges only near $2\varphi \approx \pi$.

In all cases, the chirality of noncollinear states flips sign exactly at $\varphi = 0$, where J_{ij}^{DM} changes sign. Importantly, we never observe valley-polarized states (finite τ^z or $\sigma\tau^z$), confirming that all phases (except the spin FM) remain of intervalley-coherent (IVC) type.

7.2.5 Quantum fluctuations

We now turn to the impact of quantum fluctuations on the spin–valley ordered phases. For this purpose, we employ the pf-FRG extension to SU(4) models described in Sec. 4.3. Compared to the SU(2) case, the enlarged set of generators in SU(4) significantly increases the number of independent vertices. In the present $SU(2)_{\text{spin}} \otimes U(1)_{\text{valley}}$ symmetric model, this requires keeping track of twelve symmetry-inequivalent vertex components in spin-valley space (instead of two as for SU(2) Heisenberg models). As a consequence, the numerical effort is considerably larger, and we are somewhat more restricted in lattice size and frequency resolution than for the maple-leaf Heisenberg model discussed in the previous section. In practice, we use a frequency grid of $40 \times 30 \times 30$ and consider finite clusters with linear size up to $L = 12$. Further technical details of the pf-FRG implementation for this model are provided in Appendix D.

Observables The central output of the spin–valley pf-FRG is the flow of the static spin–valley correlations $\chi_{ij}^{\mu\nu\kappa\eta}(\omega = 0) \sim \langle \sigma_i^\mu \tau^\kappa \sigma_j^\nu \tau_j^\eta \rangle$ as defined in Eq. (A.4). Since the symmetries of the Hamiltonian are preserved along the flow, it is sufficient to monitor a single representative susceptibility for each sector, which we define as

$$\begin{aligned} \chi_{ij}^\sigma &= \chi_{ij}^{\mu\mu dd} \sim \langle \sigma_i \sigma_j \rangle, \\ \chi_{ij}^{\tau^\perp} &= \chi_{ij}^{ddxx} = \chi_{ij}^{ddy} \sim \langle \tau_i^\perp \tau_j^\perp \rangle, & \chi_{ij}^{\sigma\tau^\perp} &= \chi_{ij}^{\mu\mu xx} = \chi_{ij}^{\mu\mu yy} \sim \langle \sigma_i \tau_i^\perp \sigma_j \tau_j^\perp \rangle, \\ \chi_{ij}^{\tau^z} &= \chi_{ij}^{ddzz} \sim \langle \tau_i^z \tau_j^z \rangle, & \chi_{ij}^{\sigma\tau^z} &= \chi_{ij}^{\mu\mu zz} \sim \langle \sigma_i \tau_i^z \sigma_j \tau_j^z \rangle. \end{aligned} \quad (7.22)$$

To distinguish ordered from paramagnetic states, we monitor the flow of the corresponding structure factors. We identify ordering when the flow develops negative curvature (i.e., when its second derivative becomes negative) at the largest accessible system size $L = 12$, following Ref. [146]. This criterion differs from the flow–breakdown condition introduced in Sec. 4.5.1, which we had not yet developed at the time of writing Ref. [P1]. Although slightly less strict, it yields qualitatively consistent results.

The absence of a flow breakdown signals a paramagnetic regime without any conventional order in the spin–valley space, consistent with a putative quantum spin–valley liquid, or other quantum-paramagnetic ground states. In contrast, when a breakdown occurs, the type of order is inferred from the momentum dependence of the structure factor and from the sectors that exhibit significant spectral weight. While this allows us in principle to identify the same ordered states as

in the semiclassical analysis, in practice subdominant components are usually suppressed at the breakdown scale, making mixed orders (with simultaneous contributions from multiple sectors) more difficult to resolve.

The out-of-plane susceptibilities $\chi_{ij}^{\tau^z}$ and $\chi_{ij}^{\sigma\tau^z}$, which would indicate valley polarization, remain negligibly small and are omitted in the following. The off-diagonal components $\chi^{ddxy} = -\chi^{ddyx} \sim \langle \tau^x \tau^y \rangle$ and $\chi^{\mu\mu xy} = -\chi^{\mu\mu yx} \sim \langle \sigma\tau^x \sigma\tau^y \rangle$ are indeed finite but very small. Nevertheless, they enter the definition of the staggered chiralities in Eq. (7.20) and thereby allow us to distinguish states with opposite chirality. Representative examples of structure-factor flows and chiralities in the different phases are shown in Appendix D.

Quantum phase diagrams The resulting quantum phase diagram for the TG/h-BN parameters is shown in Fig. 7.10, displaying the evolution of the critical scale, the absolute value of the dominant ordering vectors $|\mathbf{k}^{\max}|$, and the momentum-resolved structure factors in the three relevant sectors σ , τ^\perp , and $\sigma\tau^\perp$. For large negative Δ_V , both the 120° and FM phases are reproduced, with the main difference to the semiclassical case being that only the $\sigma\tau^\perp$ sector remains dominant and no mixed order is resolved (as discussed above). In contrast, in the ICS regions close to the Mott transition we observe correlations with incommensurate peaks and negative chirality, but the flow of the structure factor shows no breakdown, signaling the stabilization of a quantum paramagnetic phase. This regime is interspersed by a small ordered region, corresponding to the ICS/stripe order also found semiclassically.

To assess the robustness of these phases against variations in the SU(4)-breaking couplings, we compute phase diagrams as a function of φ_{ij} for different fixed values of J_1 and J_2 , analogous to the semiclassical case. The results, shown in Fig. 7.11, again reveal very good agreement with the semiclassical limit in all regions with 120° or FM order. The sign of the chirality also agrees perfectly across all chiral noncollinear phases. The main difference arises for parameters corresponding to $\Delta_V = -24$ meV in the ICS regions.

As illustrated in Fig. 7.12, which shows the evolution of the critical scale and structure factor peaks for varying φ in this regime, the semiclassical ICS phases consistently melt into quantum paramagnetic phases (green areas). This paramagnetic regime persists over a wide range of $2\varphi_{ij}$. This shows that competing antiferromagnetic SU(4)-symmetric interactions J_1 and J_2 favor quantum paramagnetic states which remain robust even in the presence of sizable SU(4)-breaking couplings, particularly of the J_{ij}^{DM} type.

7.2.6 Discussion

The main results of our combined semiclassical and quantum calculations can be summarized as follows. Considering the model of Ref. [49] for TG/h-BN in the strong-coupling limit, we find that at large negative displacement fields the ground state exhibits either ferromagnetic or 120° order. Importantly, these are not conventional spin orders confined to the electronic spin degree of freedom, but rather mixed spin-valley orders of inter-valley coherent (IVC) character, involving superpositions of the two valleys (signaled by finite τ^\perp or $\sigma\tau^\perp$). As discussed in the previous section, such IVC order can induce charge-density modulations on the scale of the graphene lattice constant, which can in principle be detected using scanning tunneling microscopy (STM) [275–277].

Close to the Mott transition, our pf-FRG calculations reveal a quantum paramagnetic phase stabilized by the competition of SU(4)-symmetric nearest- and next-nearest-neighbor interactions, J_1 and J_2 . That such interactions can give rise to paramagnetic, putative quantum spin-valley liquid regimes was already demonstrated in a previous pf-FRG study [104]. However, as

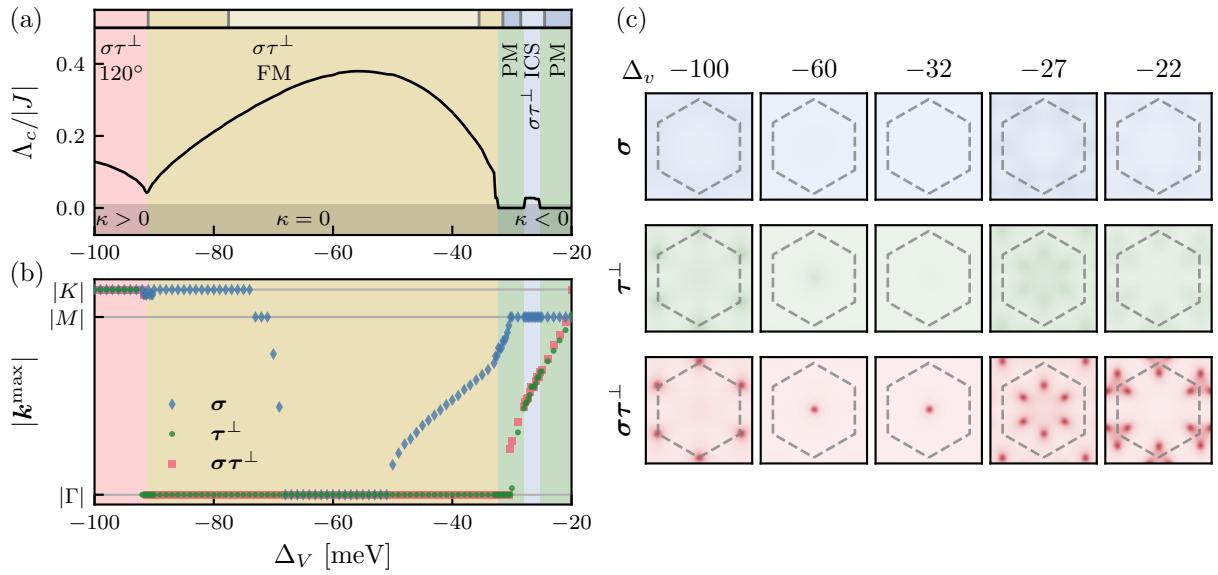


Figure 7.10 – Quantum phase diagram of TG/h-BN as a function of Δ_V , obtained from pf-FRG. (a) Critical scale Λ_c indicating the onset of long-range order. In the PM regime, the absence of a flow breakdown ($\Lambda_c = 0$) signals the lack of order in all spin-valley sectors. Top top panel shows the semiclassical phase boundaries. (b) Magnitude of the momentum \mathbf{k}^{\max} at which the structure factor is maximal (shown for all sectors, including subdominant ones). (c) Momentum-resolved structure factor for the three relevant sectors and several values of Δ_V at low cutoffs. Only the $\sigma\tau^\perp$ sector shows significant spectral weight. The rightmost panel displays the structure factor in the PM region ($\Delta_V = -22$ meV).

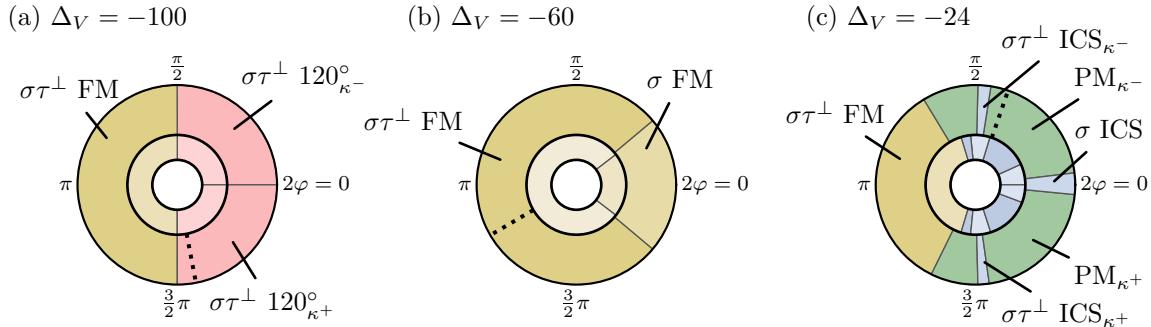


Figure 7.11 – Quantum phase diagram as a function of φ . Quantum analog of Fig. 7.9 obtained from pf-FRG (where the precise couplings are illustrated). Semiclassical phases are indicated by the inner circles. Phases and phase boundaries are shown for J_1 and J_2 fixed to their TG/h-BN estimates at $\Delta_V = -100, -60, -24$, while varying the phase φ . The semiclassical and quantum phase diagrams for (a) and (b) are essentially identical, whereas in (c) all semiclassical ICS regions melt into a quantum PM regime. The corresponding structure factors and critical scales for (c) are presented in Fig. 7.12.

shown in the two material examples discussed in this chapter, realistic systems typically feature sizable SU(4)-symmetry breaking exchange terms in the strong-coupling limit. We demonstrated that even in the presence of such terms, relatively wide paramagnetic regimes remain that show neither spin nor valley order. This warrants the continued search for exotic quantum spin-valley entangled states in moiré heterostructures, where strong correlations and competing interactions provide a natural platform for their realization.

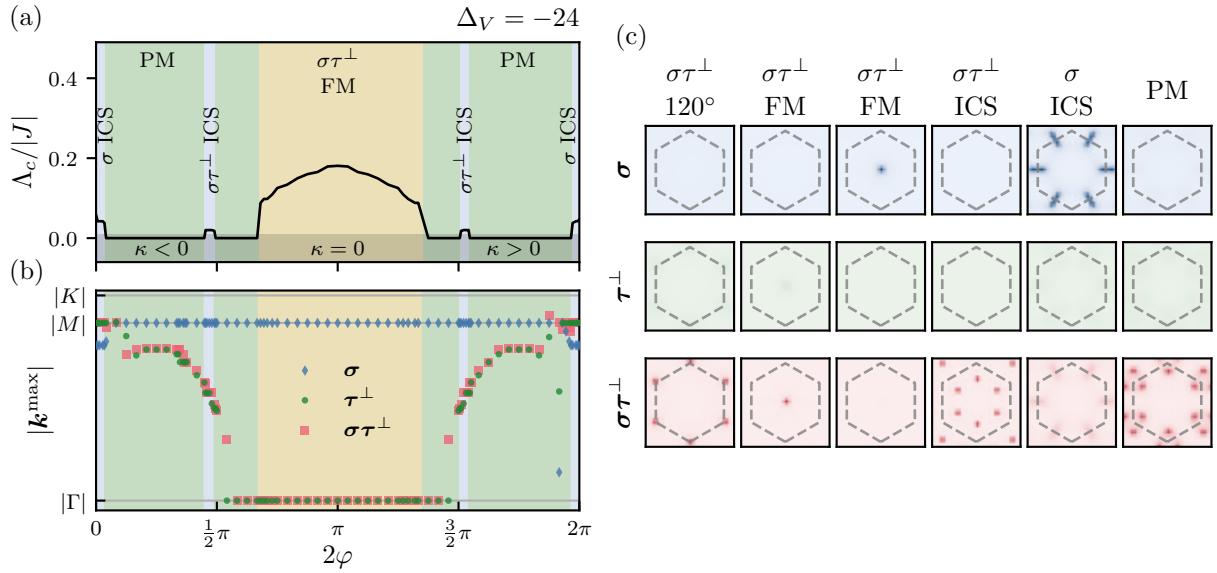


Figure 7.12 – Emergence of large quantum paramagnetic regimes and quantum structure factors. (a) Critical scale for J_1 and J_2 fixed to the TG/h-BN estimates at $\Delta_V = -24$ meV while varying the phase φ (as in Fig. 7.11(c)). Large PM regimes that show no flow breakdown (green background) emerge. The sign of the vector chirality is indicated at the bottom. (b) Evolution of the momentum k^{\max} at which the structure factor is maximal for each sector. (c) Quantum structure factors in the different phases shown in Fig. 7.11(a–c).

The pf-FRG alone cannot unambiguously determine the nature of the paramagnetic regimes. While putative quantum spin-valley liquid physics would be an exciting possibility, symmetry-breaking valence-bond solid (VBS) orders are equally likely. For the fully SU(4)-symmetric model, for example, DMRG calculations suggest that the ground state is a VBS that breaks translational invariance [53]. Performing similar DMRG or tensor network calculations may help clarify the nature of paramagnetic regime. Our pf-FRG results point such future research toward the parameter regimes where exotic quantum ground states may potentially be realized.

An important open question is whether the strong-coupling description is the correct framework for capturing the correlated states in TG/h-BN. The origin of the correlated insulating phases is still under active debate, with proposals ranging from Stoner instabilities [267] (which would mean that itinerant electrons are actually responsible for the magnetic behavior instead of localized electrons of a Mott insulator), to the relevance of Mott correlations [279], or genuine strong-coupling physics [285]. Additional tuning parameters, or different material compositions, may tip the balance between these scenarios. Our study contributes to this discussion by clarifying which phases may be realized in the strong-coupling scenario.

Chapter 8

Concluding remarks

In the first part of this thesis, we demonstrated how the pf-FRG can be generalized to a wide range of quantum spin models, most notably $SU(2)$ models with anisotropic, nondiagonal interactions relevant to materials with strong spin-orbit coupling, and various $SU(4)$ spin-valley models relevant for moiré materials and spin-orbit-entangled Mott insulators. In the second part, this framework enabled us to study the role of quantum fluctuations in three material classes: pyrochlore rare-earth oxides, maple-leaf magnets, and moiré materials. We uncovered rich phase diagrams featuring both exotic ordered and paramagnetic states: signatures of spin-nematic order and putative higher-rank spin liquids in the pyrochlore model; multiple dimer phases, along with putative spin-nematic and chiral spin liquids on the maple-leaf lattice; and entangled spin-valley orders together with a putative spin-valley quantum spin liquid in $SU(4)$ models for moiré materials.

This overview highlights both the main strength and the main weakness of the pf-FRG. Its strength lies in its very broad applicability. Its weakness is reflected in the frequent use of the word ‘putative’: although the method reliably detects quantum paramagnetic regions, it rarely allows conclusive statements about their precise nature. Let us briefly comment on both points.

Starting with the strength, current implementations [117, C3] can already treat essentially arbitrary spin models with two-spin interactions. Since model parameters enter only through the initial conditions and lattice symmetries are detected automatically, implementing new lattices or models requires only a proper definition of the unit cell and interactions. A notable limitation, however, is the lack of support for systems with more than one symmetry-inequivalent site. This is crucial for non-Archimedean lattices, materials with multiple magnetic ions, or the study of different dimer orders that require explicit symmetry breaking. Conceptually, such an extension would require site-dependent self-energies. While this increases numerical cost, it is straightforward in principle and would be a natural next step toward making our pf-FRG implementation [C3] even more general.

Turning to the main weakness, pf-FRG currently characterizes paramagnetic regions only by qualitative comparison with other techniques and known results, or by analyzing symmetry-breaking tendencies via response functions. Both lead to merely *putative* conclusions. Overcoming this limitation is particularly challenging. A more decisive characterization of quantum paramagnetic and spin-liquid phases would require access to *dynamical* correlation functions on the real-frequency axis. These could probe the excitation spectrum, reveal fractionalization, and distinguish gapped from gapless phases. They would also enable direct comparison with experiments and could, for example, help resolve the puzzle of mixed ferro- and antiferromagnetic correlations in the pyrochlore rare-earth compound $Yb_2Ti_2O_7$, discussed in Chapter 5.

As outlined in Chapter 4, obtaining dynamical correlations requires either analytic continuation from imaginary to real frequencies—a notoriously ill-conditioned problem that has so far failed for pf-FRG data—or a direct real-frequency implementation via the Keldysh formalism [286]. The latter has been demonstrated for low-dimensional itinerant models [287, 288]

and recently shown to be numerically feasible within pf-FRG [158], though important spectral features (e.g. sharp magnon lines or gaps) were not reproduced. Possible improvements include enforcing the single-occupancy constraint—for instance via the Popov–Fedotov trick [137, 138] or pseudo-Majorana representations [140, 141]—and incorporating magnetic fields to access magnetically ordered regimes, as has been demonstrated for conventional pf-FRG in Ref. [157]. While such extensions are likely to be numerically very demanding, if they succeed in enabling reliable calculations of dynamical correlation functions, they could greatly enhance the impact of pf-FRG—both for theoretical studies of quantum paramagnets and for direct comparison with experiment.

Even within the current pf-FRG framework, however, there remains much more to explore. For example, in analogy to our analysis of the pinch-line spin liquid in Chapter 5, one could investigate the influence of quantum fluctuations across the full ‘Atlas of Classical Pyrochlore Spin Liquids’ proposed in Ref. [122], which contains many additional higher-rank classical spin liquids described by different emergent gauge theories. Related proposals for higher-rank spin liquids on the *breathing* pyrochlore lattice [192] are likewise directly accessible with our method. Since their key signatures include characteristic multifold pinch-point singularities in the structure factor, pf-FRG is particularly well-suited to study them.

Focusing again on the pyrochlore lattice, we showed that quantum fluctuations can significantly shift phase boundaries, which is especially relevant for materials whose estimated parameters place them close to such boundaries. While we discussed the case of $\text{Yb}_2\text{Ti}_2\text{O}_7$, a natural next candidate is $\text{Er}_2\text{Sn}_2\text{O}_7$, which likewise appears near a phase boundary [152, 193]. To obtain more conclusive insights, however, calculations of dynamical correlation functions would once again be highly desirable.

Further development is also possible in the realm of $\text{SU}(4)$ spin models. A major limitation of our current implementation is its restriction to half-filling (equivalent to two electrons per site). This case is relevant for the insulating phases in TG/h-BN studied in Chapter 7. Many other systems of interest, however, require quarter filling (equivalent to one electron per site). Among moiré materials, the most prominent example is twisted bilayer graphene [48]. Beyond moiré physics, the strongly spin–orbit–entangled Mott insulator ZrCl_3 has also been argued to realize an approximately $\text{SU}(4)$ -symmetric model at one electron per site [46], and related models have been proposed as promising candidates for spin–orbital entangled quantum spin liquids [47]. A worthwhile direction for future work would be to explore whether techniques that enforce different fillings—such as the Popov–Fedotov trick [137, 138] or pseudo-Majorana representations [140, 141]—can be generalized to $\text{SU}(4)$ models.

Appendix A

Appendix for Chapter 4

A.1 Susceptibilities from pf-FRG

To calculate the spin-spin correlations

$$\chi_{ij}^{\Lambda ab}(\omega) = \int_0^\infty d\tau e^{i\omega\tau} \left\langle T_\tau S_i^a(\tau) S_j^b(0) \right\rangle. \quad (\text{A.1})$$

in terms of the vertex, one has to replace the spin operator on the right-hand-sight with it's pseudo-fermion representation in Eq. (1.22), use tree-expansion of the two-particle correlation function in Eq. (4.17), use the connection between connected and disconnected correlation functions

$$G(x'_1, x'_2, x_1, x_2) = G^{(c)}(x'_1, x'_2, x_1, x_2) + G(x'_1, x_1)G(x'_2, x_2) + G(x'_2, x_1)G(x'_1, x_2), \quad (\text{A.2})$$

insert the vertex parametrization of Eq. (4.48) and finally perform the sum over spin indices. This yields

$$\begin{aligned} \chi_{ij}^{\Lambda ab}(\omega) = & -\frac{1}{2} \frac{1}{2\pi} \int d\omega' G^\Lambda(\omega') G^\Lambda(\omega' + \omega) \delta_{ij} \delta_{ab} \\ & - \frac{1}{4} \left(\frac{1}{2\pi} \right)^2 \iint d\omega' d\omega'' G^\Lambda(\omega' + \omega/2) G^\Lambda(\omega' - \omega/2) G^\Lambda(\omega'' + \omega/2) G^\Lambda(\omega'' - \omega/2) \\ & \left[4 \Gamma_{ij}^{\Lambda ab}(\omega' + \omega'', \omega, \omega' - \omega'') - \delta_{ij} \sum_{\kappa\eta} \Gamma_{ii}^{\Lambda\kappa\eta}(\omega' + \omega'', \omega'' - \omega', -\omega) \text{Tr}(\boldsymbol{\theta}^\kappa \boldsymbol{\theta}^\mu \boldsymbol{\theta}^\eta \boldsymbol{\theta}^\nu) \right]. \end{aligned} \quad (\text{A.3})$$

A detailed derivation can be found in Ref. [116]. We have slightly simplified the frequency arguments by employing a shift of the integration variables by $\omega/2$. For spin-valley spin-valley

correlations the same derivation yields

$$\begin{aligned}
 \chi_{ij}^{\mu\kappa\nu\eta}(\omega) &:= \int dt e^{i\omega t} \left\langle T_\tau (\sigma_i^\mu \otimes \tau_i^\kappa)(t) (\sigma_j^\nu \otimes \tau_j^\eta)(0) \right\rangle \\
 &= -4 \frac{1}{2\pi} \int d\omega' G(\omega') G(\omega' + \omega) \delta_{ij} \delta^{\mu\nu} \delta^{\kappa\eta} \\
 &\quad - \left(\frac{1}{2\pi} \right)^2 \int d\omega' d\omega'' G^\Lambda(\omega' + \omega/2) G^\Lambda(\omega' - \omega/2) G^\Lambda(\omega'' + \omega/2) G^\Lambda(\omega'' - \omega/2) \\
 &\quad \times \left[16 \Gamma_{ij}^{\mu\kappa\nu\eta}(\omega' + \omega'', \omega, \omega' - \omega'') \right. \\
 &\quad \left. - \delta_{ij} \sum_{\bar{\mu}\bar{\nu}\bar{\kappa}\bar{\eta}} \Gamma_{ii}^{\bar{\mu}\bar{\kappa}\bar{\nu}\bar{\eta}}(\omega' + \omega'', \omega'' - \omega', -\omega) \text{Tr}(\Theta^{\bar{\mu}}\Theta^\mu\Theta^{\bar{\nu}}\Theta^\nu) \text{Tr}(\Theta^{\bar{\kappa}}\Theta^\kappa\Theta^{\bar{\eta}}\Theta^\eta) \right], \tag{A.4}
 \end{aligned}$$

where the larger prefactors arise from the additional sums over the valley indices.

Appendix B

Appendix for Chapter 5

This appendix collects technical details and supporting material for Chapter 5. We first provide the remaining definitions needed to fully specify the nearest-neighbor Hamiltonian in various parametrizations, including its decomposition into irrep order parameters. Finally, we present supplemental pf-FRG data that further illustrate how the quantum phase diagrams were obtained.

B.1 Hamiltonian in the global and local frame

Here we provide the remaining definitions of the Hamiltonian in both the local (5.9) and global (5.11) frames, following the conventions of Ref. [152]. For completeness, we also summarize the relations between the different coupling parametrizations.

The basis sites of the tetrahedral unit cell [Fig. 5.1(a)] are positioned relative to the tetrahedron center as

$$\mathbf{r}_0 = \frac{a}{8} (1, 1, 1), \quad \mathbf{r}_1 = \frac{a}{8} (1, -1, -1), \quad \mathbf{r}_2 = \frac{a}{8} (-1, 1, -1), \quad \mathbf{r}_3 = \frac{a}{8} (-1, -1, 1), \quad (\text{B.1})$$

with a the lattice constant. In the local frame, the $\tilde{\mathbf{z}}$ axis of spin $\tilde{\mathbf{S}}_i$ points from the tetrahedron center to site \mathbf{r}_i , i.e. along the local $\langle 111 \rangle$ direction. Local $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ axes follow the convention of Ref. [174], where all $\tilde{\mathbf{y}}$ axes lie in the same plane. The local and global spins are related via

$$\tilde{\mathbf{S}}_\mu = \mathbf{R}_\mu \mathbf{S}_\mu, \quad (\text{B.2})$$

with rotation matrices \mathbf{R}_μ ($\mu = 0, 1, 2, 3$) given by

$$\begin{aligned} \mathbf{R}_0 &= \frac{1}{\sqrt{6}} \begin{pmatrix} -2 & 1 & 1 \\ 0 & -\sqrt{3} & \sqrt{3} \\ \sqrt{2} & \sqrt{2} & \sqrt{2} \end{pmatrix}, & \mathbf{R}_1 &= \frac{1}{\sqrt{6}} \begin{pmatrix} -2 & -1 & -1 \\ 0 & \sqrt{3} & -\sqrt{3} \\ \sqrt{2} & -\sqrt{2} & -\sqrt{2} \end{pmatrix}, \\ \mathbf{R}_2 &= \frac{1}{\sqrt{6}} \begin{pmatrix} 2 & 1 & -1 \\ 0 & -\sqrt{3} & -\sqrt{3} \\ -\sqrt{2} & \sqrt{2} & -\sqrt{2} \end{pmatrix}, & \mathbf{R}_3 &= \frac{1}{\sqrt{6}} \begin{pmatrix} 2 & -1 & 1 \\ 0 & \sqrt{3} & \sqrt{3} \\ -\sqrt{2} & -\sqrt{2} & \sqrt{2} \end{pmatrix}. \end{aligned} \quad (\text{B.3})$$

Rows correspond to the local $\tilde{\mathbf{x}}, \tilde{\mathbf{y}}, \tilde{\mathbf{z}}$ axes. The bond-dependent phases factors γ_{ij} in the local Hamiltonian are

$$\gamma = \begin{pmatrix} 0 & 1 & -e^{-i\pi/3} & -e^{i\pi/3} \\ 1 & 0 & -e^{i\pi/3} & -e^{-i\pi/3} \\ -e^{-i\pi/3} & -e^{i\pi/3} & 0 & 1 \\ -e^{i\pi/3} & -e^{-i\pi/3} & 1 & 0 \end{pmatrix}. \quad (\text{B.4})$$

In the global frame, coupling matrices $\mathbf{J}_{\mu\nu}$ follow from the above basis transformation. For example, \mathbf{J}_{01} is given in Eq.(5.11). Other $\mathbf{J}_{\mu\nu}$ are generated by lattice symmetries, namely C_3 rotations around local $\tilde{\mathbf{z}}$ axes and C_2 rotations around the global z axis,

$$\mathbf{C}_3 = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad \mathbf{C}_2 = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (\text{B.5})$$

Combined with a possible lattice inversion along the corresponding bond (implemented as a matrix transpose), the remaining coupling matrices can be obtained from \mathbf{J}_{01} through the symmetry transformations that map the respective bonds onto one another, namely

$$\begin{aligned} \mathbf{J}_{02} &= \mathbf{C}_3 \mathbf{J}_{01} \mathbf{C}_3^T, & \mathbf{J}_{03} &= \mathbf{C}_3^T \mathbf{J}_{01} \mathbf{C}_3, & \mathbf{J}_{12} &= \mathbf{C}_3^T \mathbf{C}_2^T \mathbf{J}_{01}^T \mathbf{C}_2 \mathbf{C}_3, \\ \mathbf{J}_{13} &= \mathbf{C}_3 \mathbf{C}_2^T \mathbf{J}_{01}^T \mathbf{C}_2 \mathbf{C}_3^T, & \mathbf{J}_{23} &= \mathbf{C}_2 \mathbf{J}_{01}^T \mathbf{C}_3^T. \end{aligned} \quad (\text{B.6})$$

which reproduces the coupling matrices in Ref. [152].

The relation between couplings in local and global frames is

$$\begin{pmatrix} J_{zz} \\ J_{\pm} \\ J_{\pm\pm} \\ J_{z\pm} \end{pmatrix} = \frac{1}{6} \begin{pmatrix} -4 & 2 & -4 & -8 \\ 2 & -1 & -1 & -2 \\ 1 & 1 & -2 & 2 \\ \sqrt{2} & \sqrt{2} & \sqrt{2} & -\sqrt{2} \end{pmatrix} \begin{pmatrix} J_1 \\ J_2 \\ J_3 \\ J_4 \end{pmatrix}. \quad (\text{B.7})$$

Finally, a π rotation around any local $\tilde{\mathbf{z}}$ axis flips $\tilde{S}^{\pm} \rightarrow -\tilde{S}^{\pm}$, leaving \tilde{S}^z unchanged, hence $J_{z\pm} \rightarrow -J_{z\pm}$ [189]. This defines a dual parametrization $(\bar{J}_1, \bar{J}_2, \bar{J}_3, \bar{J}_4)$, related to the original (J_1, \dots, J_4) via

$$\begin{pmatrix} \bar{J}_1 \\ \bar{J}_2 \\ \bar{J}_3 \\ \bar{J}_4 \end{pmatrix} = \frac{1}{9} \begin{pmatrix} 5 & -4 & -4 & 4 \\ -8 & 1 & -8 & 8 \\ -4 & -4 & 5 & 4 \\ 2 & 2 & 2 & 7 \end{pmatrix} \begin{pmatrix} J_1 \\ J_2 \\ J_3 \\ J_4 \end{pmatrix}. \quad (\text{B.8})$$

B.2 Irrep decomposition and order-parameter fields

The order-parameter fields used in the irrep decomposition of the classical Hamiltonian in Eq. (5.14) are defined in terms of global spin operators, following Ref. [152], as

$$\begin{aligned}
 m_{A_2} &= \frac{1}{2\sqrt{3}} (S_0^x + S_0^y + S_0^z + S_1^x - S_1^y - S_1^z - S_2^x + S_2^y - S_2^z - S_3^x - S_3^y + S_3^z) , \\
 \mathbf{m}_E &= \begin{pmatrix} \frac{1}{2\sqrt{6}} (-2S_0^x + S_0^y + S_0^z - 2S_1^x - S_1^y - S_1^z + 2S_2^x + S_2^y - S_2^z + 2S_3^x - S_3^y + S_3^z) \\ \frac{1}{2\sqrt{2}} (-S_0^y + S_0^z + S_1^y - S_1^z - S_2^y - S_2^z + S_3^y + S_3^z) \end{pmatrix} , \\
 \mathbf{m}_{T_{1A}} &= \begin{pmatrix} \frac{1}{2} (S_0^x + S_1^x + S_2^x + S_3^x) \\ \frac{1}{2} (S_0^y + S_1^y + S_2^y + S_3^y) \\ \frac{1}{2} (S_0^z + S_1^z + S_2^z + S_3^z) \end{pmatrix} , \\
 \mathbf{m}_{T_{1B}} &= \begin{pmatrix} \frac{-1}{2\sqrt{2}} (S_0^y + S_0^z - S_1^y - S_1^z - S_2^y + S_2^z + S_3^y - S_3^z) \\ \frac{-1}{2\sqrt{2}} (S_0^x + S_0^z - S_1^x + S_1^z - S_2^x - S_2^z + S_3^x - S_3^z) \\ \frac{-1}{2\sqrt{2}} (S_0^x + S_0^y - S_1^x + S_1^y + S_2^x - S_2^y - S_3^x - S_3^y) \end{pmatrix} , \\
 \mathbf{m}_{T_2} &= \begin{pmatrix} \frac{1}{2\sqrt{2}} (-S_0^y + S_0^z + S_1^y - S_1^z + S_2^y + S_2^z - S_3^y - S_3^z) \\ \frac{1}{2\sqrt{2}} (S_0^x - S_0^z - S_1^x - S_1^z - S_2^x + S_2^z + S_3^x + S_3^z) \\ \frac{1}{2\sqrt{2}} (-S_0^x + S_0^y + S_1^x + S_1^y - S_2^x - S_2^y + S_3^x - S_3^y) \end{pmatrix} , \\
 \mathbf{m}_{T_{1-}} &= \cos \theta \mathbf{m}_{T_{1A}} - \sin \theta \mathbf{m}_{T_{1B}} , \\
 \mathbf{m}_{T_{1+}} &= \sin \theta \mathbf{m}_{T_{1A}} + \cos \theta \mathbf{m}_{T_{1B}} ,
 \end{aligned} \tag{B.9}$$

where the angle

$$\theta = \frac{1}{2} \arctan \left(\frac{\sqrt{8}J_3}{2J_1 + 2J_2 + J_3 - 2J_4} \right) \tag{B.10}$$

is chosen such that the coupling between $\mathbf{m}_{T_{1A}}$ and $\mathbf{m}_{T_{1B}}$ vanishes, yielding $\mathbf{m}_{T_{1-}}$ and $\mathbf{m}_{T_{1+}}$. Physically, θ is the canting angle of the spins in the T_1 ground state, forming a splayed ferromagnet around the $\langle 111 \rangle$ axis (or symmetry equivalents). The prefactors a_λ of the order-parameter fields in the Hamiltonian are

$$\begin{aligned}
 a_{A_2} &= -2J_1 + J_2 - 2J_3 + 4J_4 , \\
 a_E &= -2J_1 + J_2 + J_3 + 2J_4 , \\
 a_{T_{1-}} &= (2J_1 + J_2) \cos^2 \theta - (J_2 + J_3 - 2J_4) \sin^2 \theta + \sqrt{2}J_3 \sin 2\theta , \\
 a_{T_{1+}} &= (2J_1 + J_2) \sin^2 \theta - (J_2 + J_3 - 2J_4) \cos^2 \theta - \sqrt{2}J_3 \sin 2\theta , \\
 a_{T_2} &= -J_2 + J_3 - 2J_4 ,
 \end{aligned} \tag{B.11}$$

from which the classical ground-state in $\mathbf{q} = 0$ ordered phases can be directly deduced.

B.3 Supplemental data for the pf-FRG calculation

Here, we provide additional figures for a better insight in how the quantum phase diagrams in Chapter 5 were created. Fig. B.1 shows representative RG flows of the structure factor across ordered and paramagnetic regimes, illustrating how the flow-breakdown criterion is applied to distinguish those two regimes. Fig. B.2 presents cuts through the phase diagram, indicating the critical scale Λ_c and the relative magnitudes of the irrep susceptibilities, which are used to classify the different paramagnetic regimes.

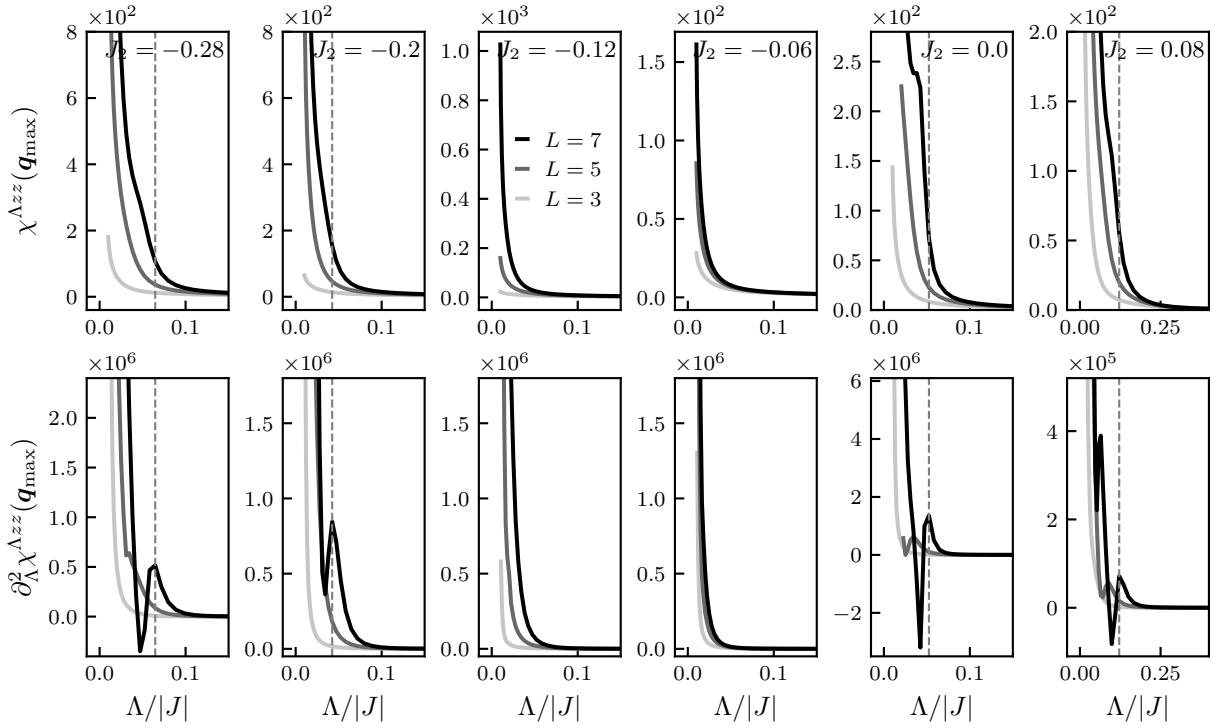


Figure B.1 – RG flow of the structure factor at the momentum \mathbf{k}_{\max} where it is maximal, for fixed $J_1/|J_3| = 0.12$, $J_4 = 0$, and $J_3 < 0$, with varying $J_2/|J_3|$. Dashed gray lines indicate the critical scale Λ_c at which a flow breakdown is identified using the criterion of Sec. 4.5.1. The two leftmost and rightmost panels correspond to parameters in the ordered T_{1-} and E phases, respectively, while the central panels show the paramagnetic regime, where the flow at derivatives remains smooth down to $\Lambda/|J| = 0.01$.

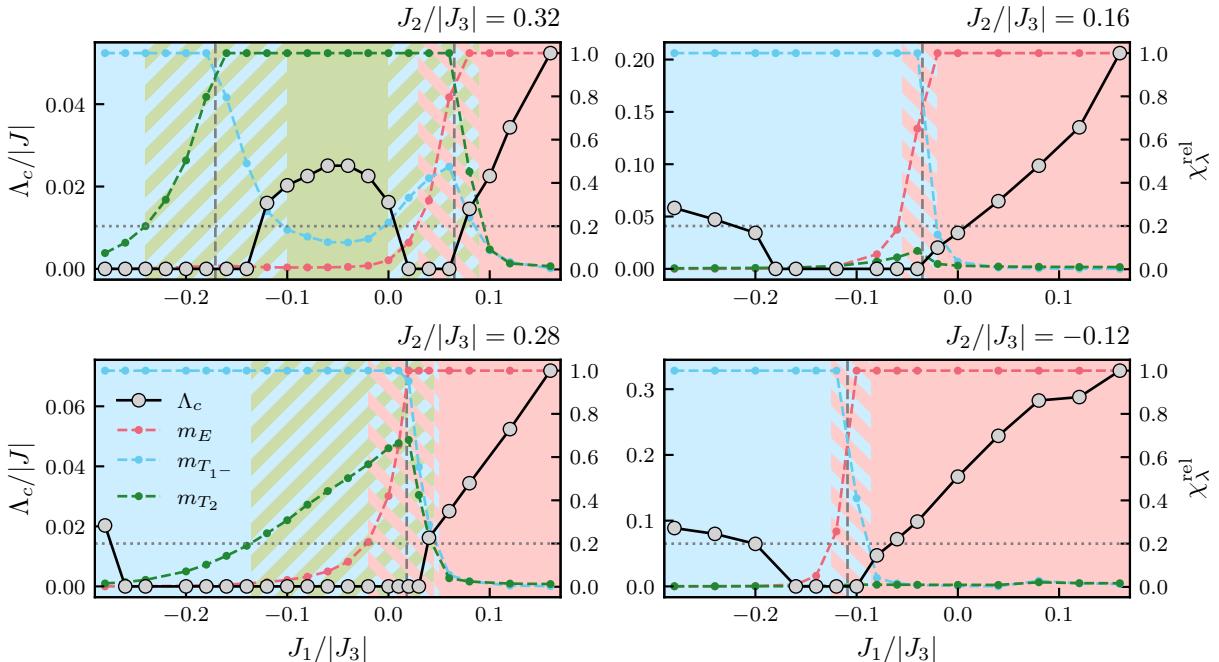


Figure B.2 – Cuts through the phase diagram in Fig. 5.2(b). The black dots indicate the critical scale Λ_c , while the colored dots represent the relative values of the irrep susceptibilities in the low-cutoff limit. Regions where multiple $\chi_{\lambda}^{\text{rel}}$ exceed 20% (dotted horizontal line) are highlighted with hatched backgrounds. Dashed lines mark the couplings where the dominant irrep susceptibility changes.

Appendix C

Appendix for Chapter 6

In this appendix, we first specify the conventions used for the maple-leaf lattice and then provide supplemental data from the pf-FRG calculations presented in Chapter 6. These results help to illustrate more clearly how the phase diagrams and quantum phase boundaries shown in the main text were constructed.

C.1 Definition of the maple-leaf lattice

In all calculations, we adopt the convention in which the maple-leaf lattice is defined by the two lattice vectors

$$\mathbf{a}_1 = \left(\frac{3\sqrt{3}}{2}, -\frac{1}{2} \right), \quad \mathbf{a}_2 = \left(\sqrt{3}, 2 \right), \quad (\text{C.1})$$

together with six basis sites within the hexagonal unit cell,

$$\boldsymbol{\delta}_1 = (0, 0), \quad \boldsymbol{\delta}_2 = \left(\frac{\sqrt{3}}{2}, -\frac{1}{2} \right), \quad \boldsymbol{\delta}_3 = (\sqrt{3}, 0), \quad \boldsymbol{\delta}_4 = (\sqrt{3}, 1), \quad \boldsymbol{\delta}_5 = \left(\frac{\sqrt{3}}{2}, \frac{3}{2} \right), \quad \boldsymbol{\delta}_6 = (0, 1), \quad (\text{C.2})$$

with lattice spacing set to $a = 1$. In momentum space, this corresponds to the reciprocal lattice vectors

$$\mathbf{G}_1 = \frac{4\pi}{7} \left(\frac{2}{\sqrt{3}}, -1 \right), \quad \mathbf{G}_2 = \frac{4\pi}{7} \left(\frac{1}{2\sqrt{3}}, \frac{3}{2} \right), \quad (\text{C.3})$$

from which the first Brillouin zone is obtained (shown as dashed lines in all structure-factor and \mathbf{q}^{LT} plots in the main text). The extended Brillouin zone, corresponding to the first Brillouin zone of the triangular lattice that reduces to the maple-leaf lattice upon depletion, follows from the reciprocal lattice vectors

$$\mathbf{G}_1^{\Delta} = 2\pi \left(\frac{1}{\sqrt{3}}, -1 \right), \quad \mathbf{G}_2^{\Delta} = 2\pi \left(\frac{1}{\sqrt{3}}, 1 \right). \quad (\text{C.4})$$

Equivalently, it can be obtained by scaling the original Brillouin zone by $\sqrt{7}$ and applying a rotation by the angle

$$\phi = \arccos\left(\frac{5}{2\sqrt{7}}\right) \quad (\text{C.5})$$

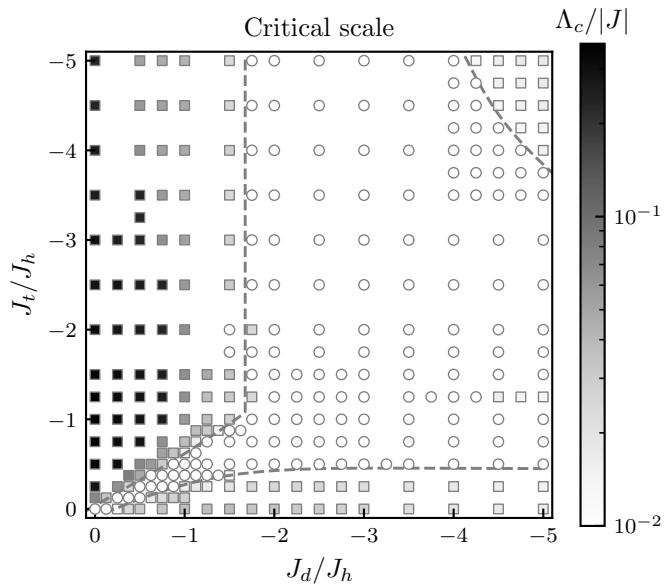
A peculiarity of the maple-leaf lattice is that its nearest-neighbor bonds are rotated with respect to the lattice vectors by precisely this odd angle.

C.2 Evolution of the critical scale for the ferro–antiferromagnetic Heisenberg model

Since it was not included in the main text, we present in Fig. C.1 the evolution of the critical scale for the ferro–antiferromagnetic Heisenberg model, which was used to determine the approximate

Figure C.1 –

Critical scale from pf-FRG in the nearest neighbor ferro–antiferromagnetic Heisenberg model The color scale represents the critical scale Λ_c , as determined by the flow breakdown criterion used to identify the PM phase in the pf-FRG phase diagram shown in Fig. 6.10. Square markers indicate points in the PM phase where no flow breakdown was detected. Dashed lines serve as visual guides, approximating the phase boundaries to the PM phase.



boundaries of the PM regime shown in Fig. 6.10. Examples of the flow are provided in Sec. 4.5.1, where the procedure for extracting the critical scales is also discussed. The figure is redrawn from Ref. [U1].

C.3 Supplemental data: cross-plaquette interactions

This section provides additional data supplementing Sec. 6.4. All figures are redrawn from Ref. [P6], and all underlying data were generated by the author of this thesis.

Cuts through the quantum phase diagrams To aid the interpretation of the full quantum phase diagram of the AFM ($J_1 > 0$) shown in Fig. 6.12, we present in Fig. C.2 the evolution of the critical scale Λ_c and the ordering vector \mathbf{q}^{\max} along vertical (fixed J_3/J_1) and horizontal (fixed J_2/J_1) cuts through parameter space. These cuts highlight the regions with incommensurate order (ICS), where the ordering vector lies neither at a symmetry point of the first nor of the extended Brillouin zone of the maple-leaf lattice. They also show the continuous evolution of the ordering vector between phases III and VI, rather than the two distinct phases IV and V found in the classical analysis (see, e.g., the upper right panel for fixed $J_2/J_1 = 0.8$).

Dips in the critical scale appear at the phase boundaries between phases I and VI, and between phases II and III (or the nearby ICS regime), consistent with phase transitions. In contrast, no clear feature is visible at the I–II boundary, suggesting a crossover rather than a transition. This is consistent with the classical ground-state energy, which shows only a very weak kink along the same cut [P6], indicative of a weak first-order transition that may not be well resolved by the pf-FRG critical scale alone. The pf-FRG ordering vectors, however, do exhibit a sharp jump at this boundary, though slightly shifted compared to the classical result.

Analogous cuts through the FM phase diagram ($J_1 < 0$) [Fig. 6.14] are shown in Fig. C.3. Here, too, the ICS regimes are characterized by a continuous evolution of the structure factor. In particular, the lower right panel demonstrates that the ICS region between phases PM2 and PM3 develops structure-factor peaks exactly at the commensurate K point only at a single parameter value, without forming an extended commensurate region.

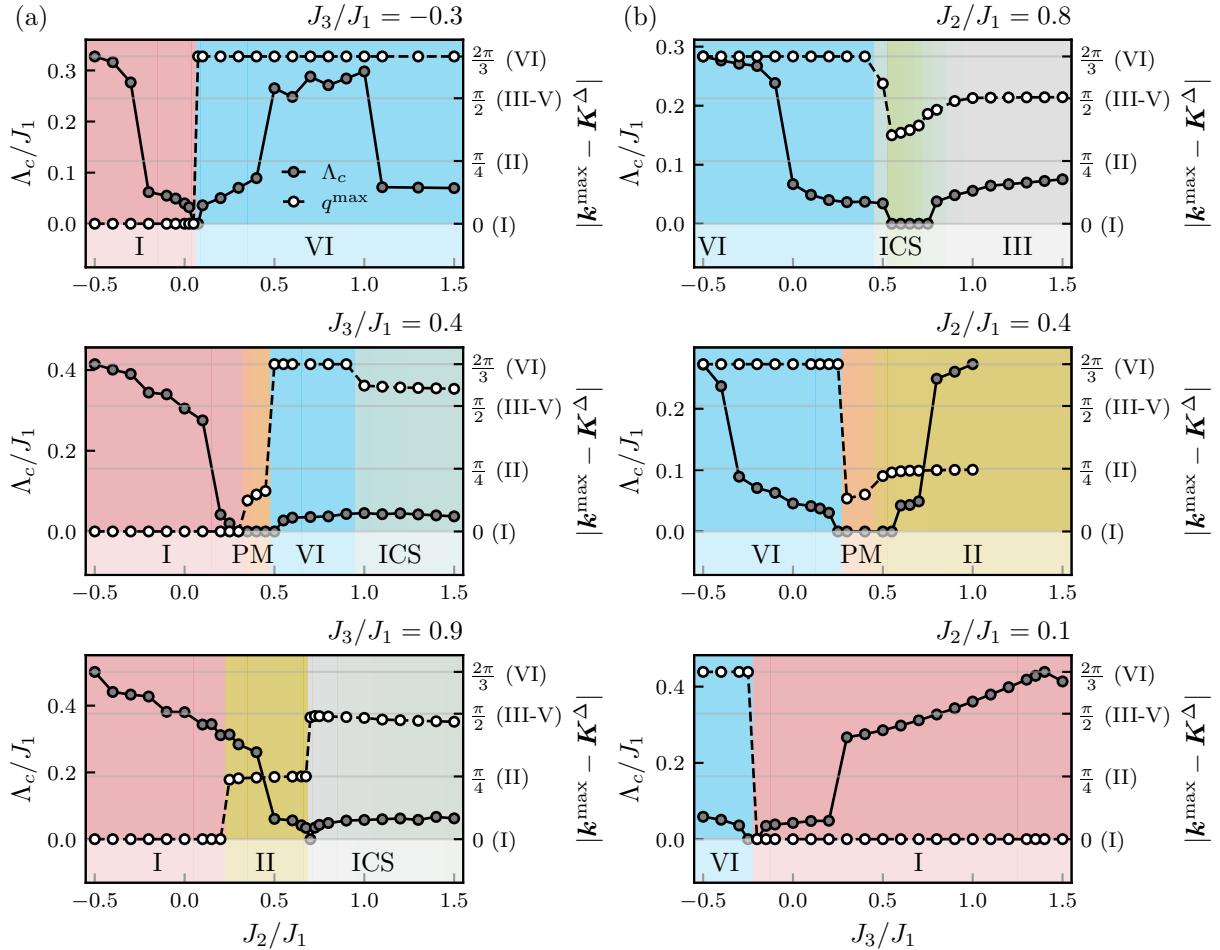


Figure C.2 – Cuts through the quantum phase diagram of the AFM shown in [Fig. 6.12]. Gray circles mark the critical scale Λ_c , and white circles trace the evolution of the ordering vector \mathbf{q}^{\max} for (a) vertical cuts ($J_3/J_1 = \text{const.}$) and (b) horizontal cuts ($J_2/J_1 = \text{const.}$). Beyond the ordered phases found in the classical analysis, extended regions with incommensurate (ICS) ordering vectors appear.

Comparison of pf-FRG and unconstrained Luttinger-Tisza To further substantiate the structure factors obtained from our pf-FRG calculations—particularly in the incommensurate regimes where they disagree with the classical analysis—we also compare them with the \mathbf{q}^{LT} vectors minimizing the Luttinger-Tisza energy [Fig. C.4 for the AFM, Fig. C.5 for the FM]. Although these vectors do not represent the true classical ground state in the ICS regimes, where the strong spin-length constraint is violated, they still provide a useful semiclassical approximation, since the Luttinger-Tisza energy serves as a lower bound to the exact ground-state energy (see Sec. 3.1).

We find very good agreement with the pf-FRG structure-factor peaks—with one caveat: \mathbf{q}^{LT} vectors are periodic under the reciprocal lattice vectors \mathbf{G}_i , whereas the structure factor is periodic under \mathbf{G}_i^{Δ} (both defined above). This mismatch reduces the number of observable peaks compared to \mathbf{q}^{LT} vectors (see Sec. 4.5.2). As a consequence, phases II and VI cannot be distinguished based solely on \mathbf{q}^{LT} vectors, since they share the same magnetic unit cell. The underlying spin-sublattice structure is instead encoded in the Luttinger-Tisza eigenvectors, which are not shown here. Taken together, these observations strongly suggest that the Monte Carlo phases IV and V do not represent distinct phases, but rather a single ICS phase.

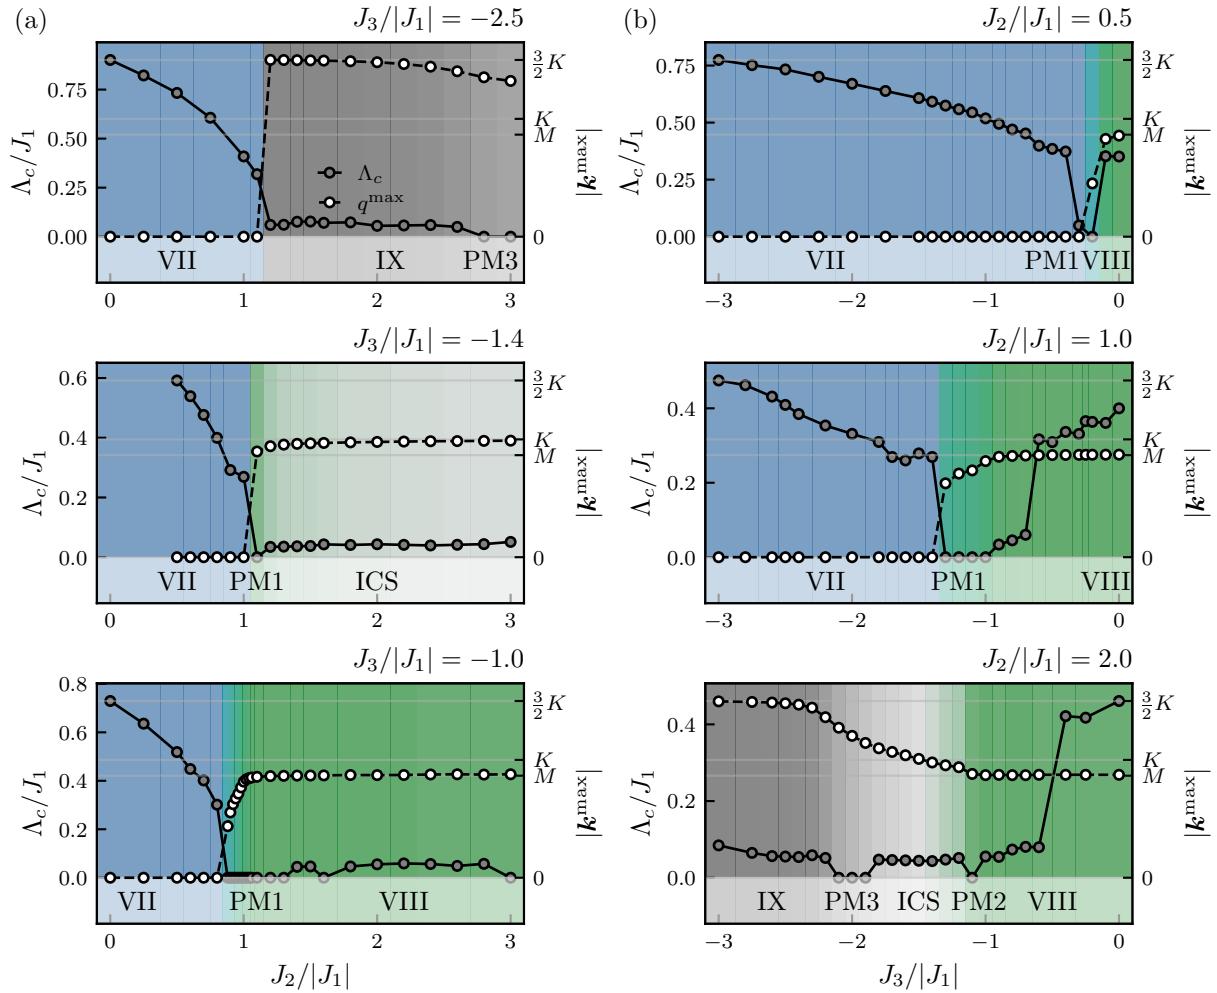


Figure C.3 – Cuts through the quantum phase diagram of the FM shown in [Fig. 6.14]. Gray circles mark the critical scale Λ_c , and white circles trace the evolution of the ordering vector q^{\max} for (a) vertical cuts ($J_3/J_1 = \text{const.}$) and (b) horizontal cuts ($J_2/J_1 = \text{const.}$)

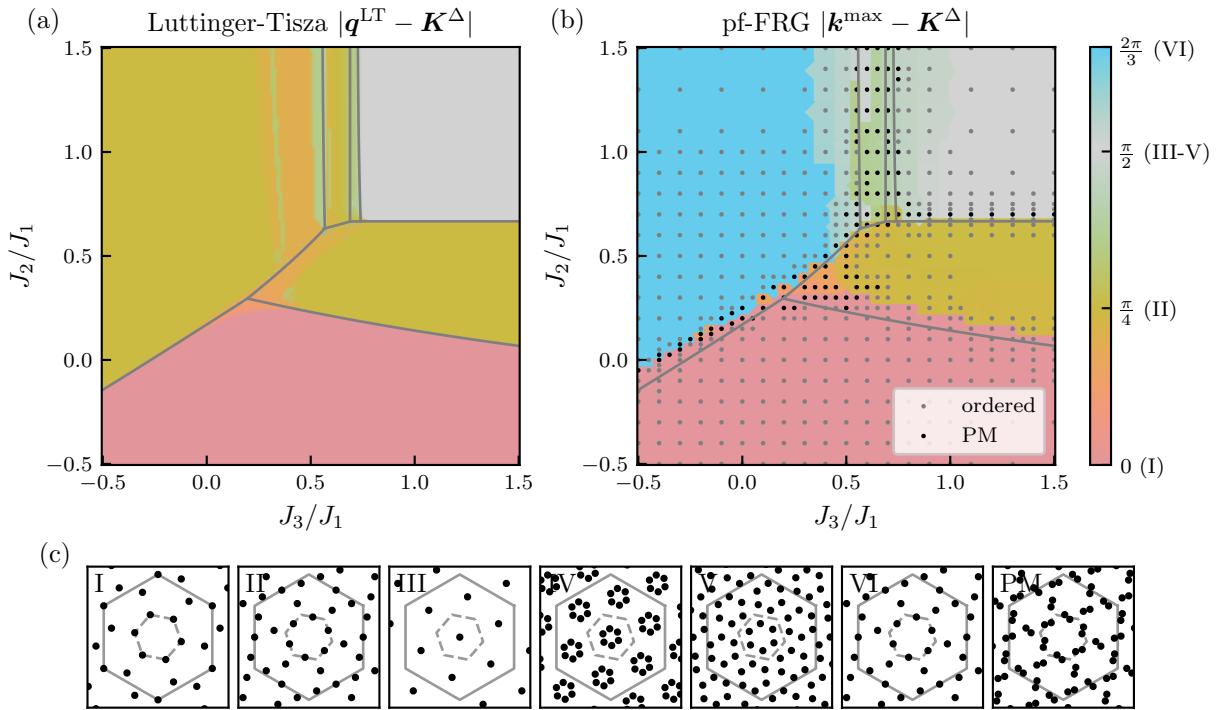


Figure C.4 – Comparison of Luttinger–Tisza and pf-FRG for the AFM. (a) LT \mathbf{q}^{LT} vectors minimizing the Fourier-transformed interaction matrix. (b) pf-FRG phase diagram [also shown in Fig. 6.12(b)]. (c) Minimal \mathbf{q}^{LT} vectors at the same parameters as in the quantum case [Fig. 6.12(c)]. In phases IV and V the strong spin-length constraint is violated, so these \mathbf{q}^{LT} vectors do not necessarily represent the true classical ground state. In all other phases the constraint is satisfied, yielding the exact classical ground state.

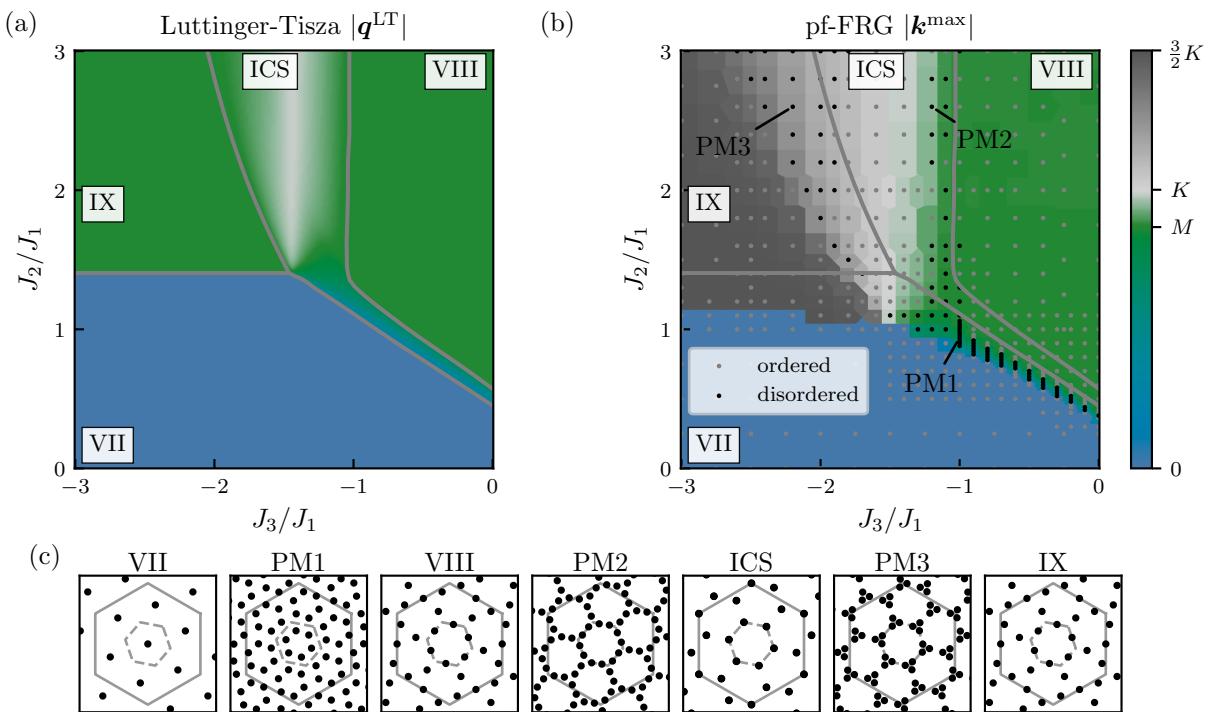


Figure C.5 – Comparison of Luttinger-Tisza and pf-FRG for the FM. Same as Fig. C.4, but for $J_1 < 0$. The corresponding quantum phase diagram is shown in Fig. 6.14. The spin-length constraint is satisfied in all phases except the ICS regimes.

Appendix D

Appendix for Chapter 7

We provide supplemental data for Chapter 7. All figures in this appendix are redrawn from Ref. [P1] and the data is obtained by the author of this thesis.

D.1 Vertex parametrization for an $SU(2) \otimes U(1)$ symmetric spin-valley model

To efficiently parametrize the vertex in the pf-FRG calculation, we must identify the components $\Gamma^{\mu\nu\kappa\eta}$ that are equivalent under the spin-valley symmetry of the model. This follows directly from the discussion in Sec. 4.2.7 on how spin symmetries constrain the vertex functions, which can be applied independently to the spin and valley sectors of the spin-valley model.

According to Eq. (4.93), the $SU(2)$ spin symmetry restricts the spin indices of the vertex to density ($d = dd$) and spin ($s = xx = yy = zz$) terms. Likewise, as shown in Eq. (4.83), the $U(1)$ valley symmetry allows six distinct components: dd , $xx = yy$, zz , $xy = -yx$, dz , and zd . Accordingly, in the most general vertex parametrization, which (following Eq. (4.116)) can be expressed as (only stating spin indices)

$$\Gamma(x'_1, x'_2; x_1, x_2) = \sum_{\mu, \kappa, \eta=0}^3 \Gamma^{\mu\nu\kappa\eta} \theta_{s'_1 s_1}^\mu \theta_{s'_2 s_2}^\mu \theta_{l'_1 l_1}^\kappa \theta_{l'_2 l_2}^\eta, \quad (D.1)$$

we can identify twelve symmetry-equivalent vertex components ($d, x, y, z = 0, 1, 2, 3$) as

$$\begin{array}{ll} \Gamma^{sxx} = \Gamma^{syy} & \Gamma^{dxx} = \Gamma^{dyy} \\ \Gamma^{szz} & \Gamma^{dzz} \\ \Gamma^{sxy} = -\Gamma^{syx} & \Gamma^{dxy} = -\Gamma^{dyx} \\ \Gamma^{sdz} & \Gamma^{ddz} \\ \Gamma^{szd} & \Gamma^{dzd} \\ \Gamma^{sdd} & \Gamma^{ddd} \end{array} \quad (D.2)$$

Although the density components are zero at the start of the flow, they acquire finite values as Λ decreases. The sums over the spin and valley indices are performed numerically, yielding flow equations for each component individually.

D.2 Thermodynamics in TG/h-BN

In this section, we present additional data for the phase transitions observed semiclassically in TG/h-BN across the three principal regimes identified in Fig. 7.8: the 120° , FM, and ICS regimes.

Fig. D.1 shows the specific heat and energy as functions of temperature, as well as the energy distribution at the critical temperature. Both the 120° and FM phases exhibit signatures consistent with a continuous transition (which would have to be of BKT type), whereas the ICS phase displays clear evidence of a first-order transition, as indicated by the double-peak structure in the energy histogram. The distance between these peaks corresponds to the latent heat of the transition, which is shown as a function of Δ_V in Fig. D.2.

Fig. D.3 presents the specific heat down to very low temperatures for all three regimes. In each case, the specific heat approaches $c_v(T \rightarrow 0) = 5$, indicating the presence of five harmonic zero modes. This is consistent with the ten parameters that locally describe a state in the SU(4) model with the six-dimensional representation considered here, since each state is represented by a six-component complex vector (12 parameters), where normalization and an arbitrary phase remove two degrees of freedom.

D.3 Supplemental data for the semiclassical and quantum phase diagrams

Here, we present additional observables underlying the classical and quantum phase diagrams of the TG/h-BN model.

Fig. D.4 shows examples of the structure-factor flows for different phases. A flow breakdown is identified when the flow develops negative curvature (i.e., when the second derivative becomes negative). Only in the paramagnetic (PM) phase does the flow remain smooth and convex down to the lowest considered cutoff, $\Lambda/|J| = 0.01$.

Fig. D.5 displays the flow of the vector chiralities, as defined in Eq. (7.20), for all observed phases. The results are fully consistent with the chiralities obtained from semiclassical Monte Carlo simulations: collinear (FM) phases correctly maintain zero chirality throughout the flow, and the sign of the chiral phases is always opposite to J_{ij}^{DM} .

Fig. D.6 presents supplemental data corresponding to the classical (Fig. 7.9) and quantum (Fig. 7.11) phase diagrams as functions of the phase φ . Shown are (i) the critical temperature, obtained by fitting peaks in the specific heat; (ii) the ground-state energy from numerical minimization, together with analytically determined energies for the different phases; and (iii) the critical scale extracted from the pf-FRG.

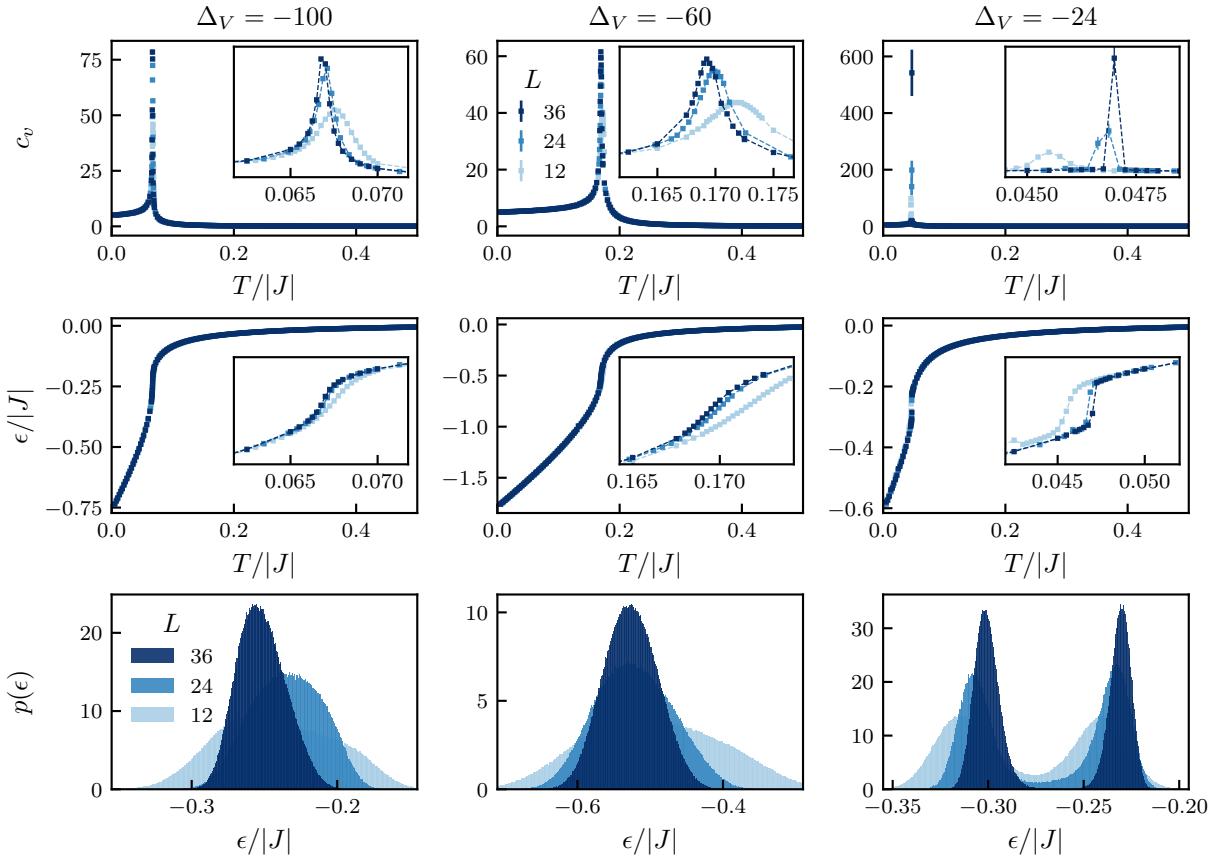


Figure D.1 – Thermodynamics of the spin-valley model obtained from semiclassical Monte Carlo simulations for TG/h-BN-inspired coupling parameters, shown for three different potential differences $\Delta_V = -100, -60, -24$ meV. These stabilize a 120° ordered state (left column), a spin-valley ferromagnet (middle column), and incommensurate (ICS) order (right column). The top row shows the specific heat, the middle row the energy per site, and the bottom row the energy histogram at the thermal phase transition. The double-peak structure in the latter indicates a first-order transition.

Figure D.2 –
Latent heat of the first-order transition into ICS/stripe order. The energy distribution at the transition temperature exhibits a double-peak structure, as shown in the top panels for $\Delta_V = -32, -31$, and -29 meV. The latent heat is defined as the difference between the peak positions, which we determine by fitting double Gaussians to the energy distributions. All results are obtained for a lattice size of $L = 24$.

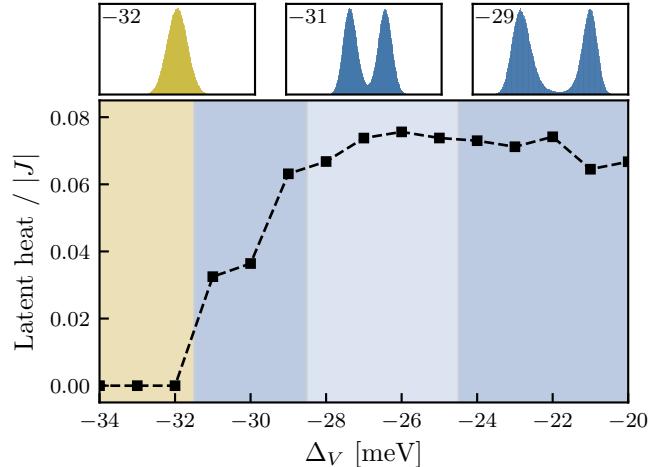


Figure D.3 –

Specific heat saturation. Shown are the exact classical structure factors (left) and common origin plots (right) for the two noncoplanar phases found in the phase diagram of the ferromagnet in Fig. 6.14. Both spin configurations can be exactly captured using Luttinger-Tisza with a triple- \mathbf{q} Ansatz. Common origin plots are taken directly from Ref. [P6].

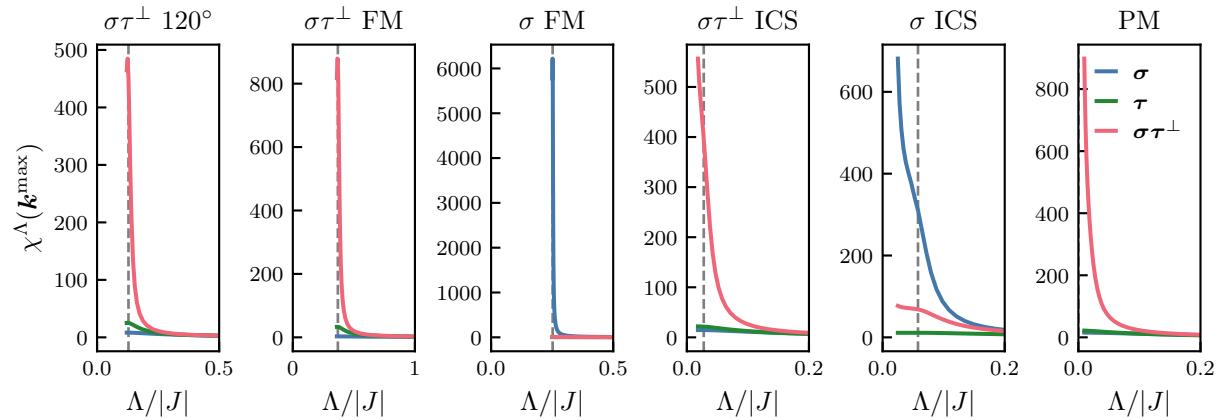
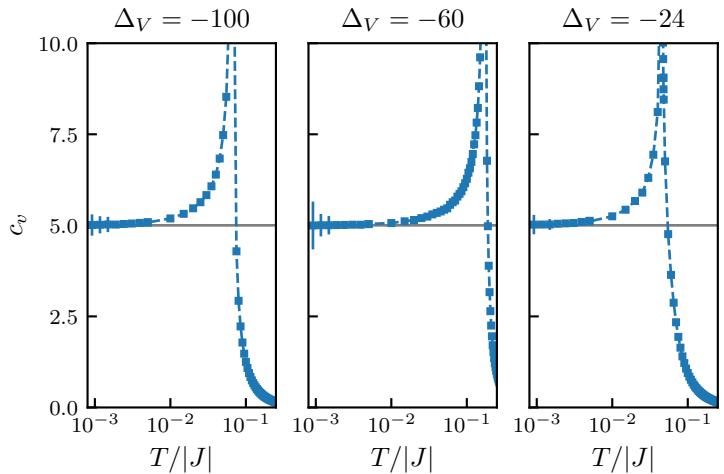


Figure D.4 – Flows of the spin-valley structure factors in TG/h-BN. Shown are the pf-FRG structure factor flows for the different phases of the model, evaluated at the momentum of maximal intensity \mathbf{k}^{\max} . The dotted gray line indicates the critical scale Λ_c , which marks the onset of magnetic order and is determined from the point of maximum negative curvature in the flow. In the paramagnetic (PM) regime, the flow remains fully convex down to the lowest numerically accessible cutoff $\Lambda/|J| = 0.01$. This behavior signals the absence of magnetic ordering. Typically, we find that one sector clearly dominates, while the flows of the other sectors are strongly suppressed.

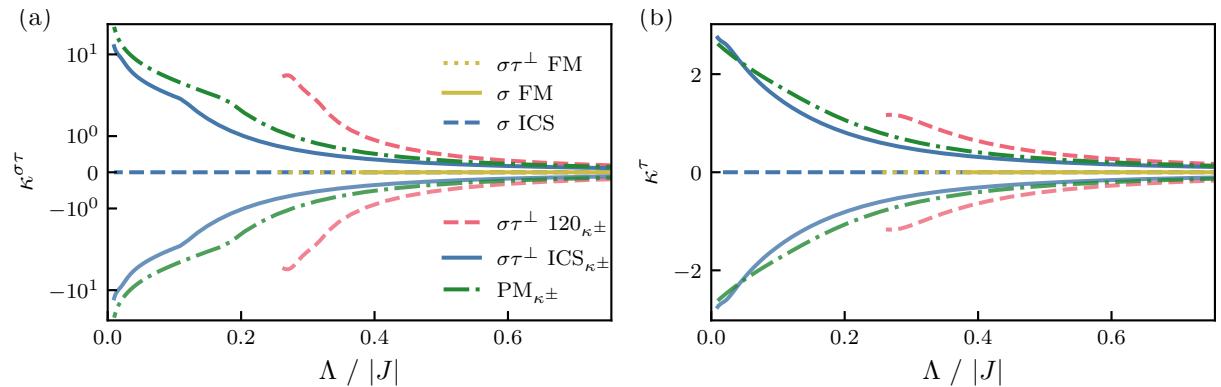


Figure D.5 – Renormalization group flow of the staggered chirality deep within the different phases of the TG/h-BN model. Ferromagnetic (FM) states, or states that order exclusively in the spin sector, always exhibit zero chirality. For all other phases, the chirality is exactly opposite in sign to the coupling J_{ij}^{DM} .

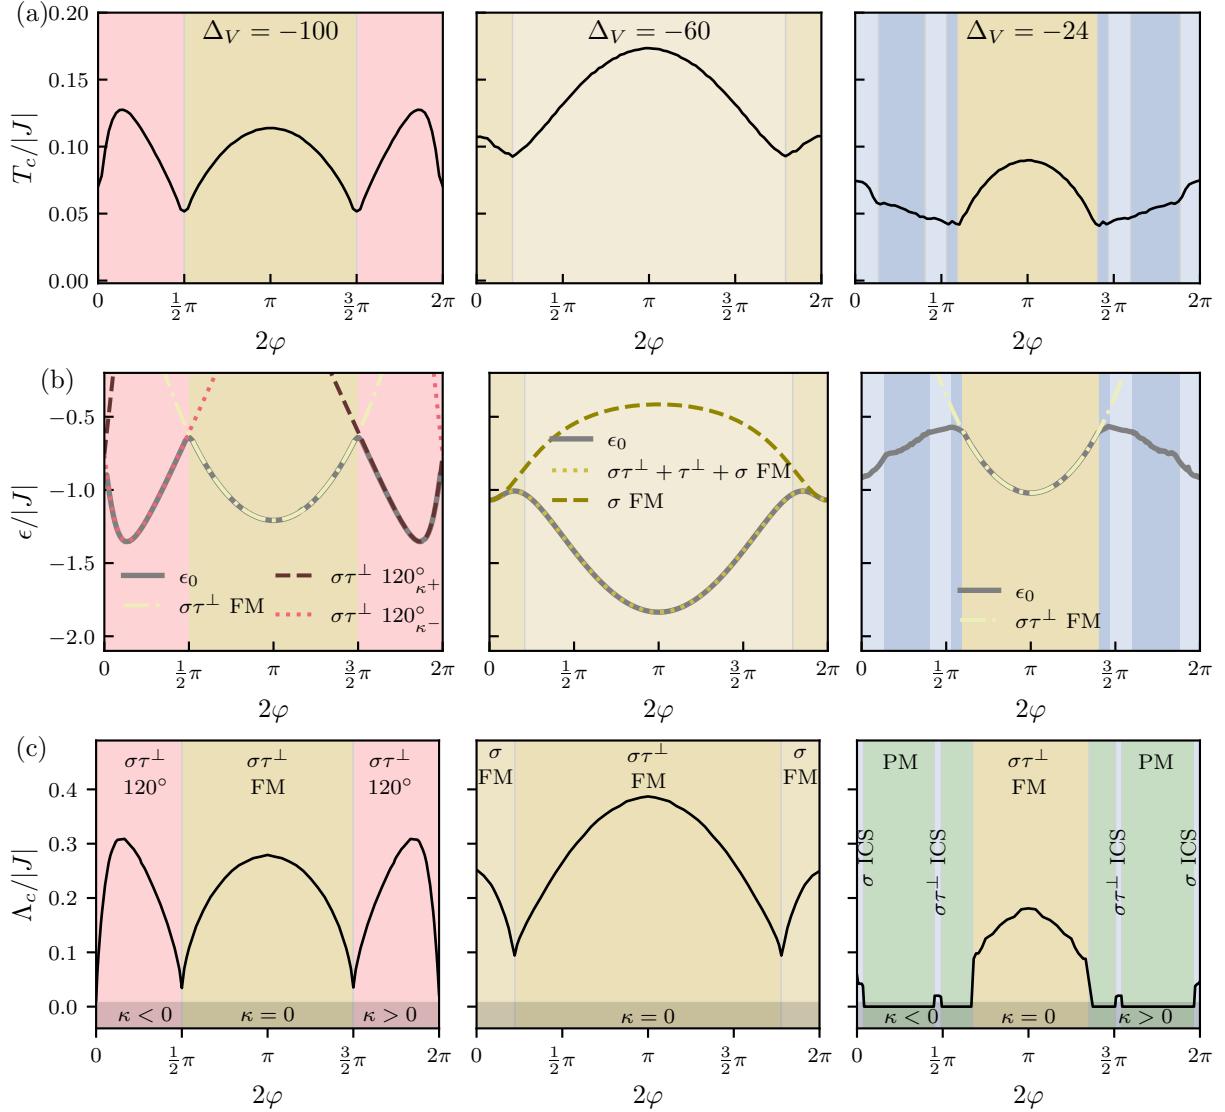


Figure D.6 – Classical and quantum phase diagram as a function of φ . The couplings J_1 and J_2 are fixed to the TG/h-BN estimates for different values of Δ_V , as shown in the top panel of each column, while the phase φ is varied. The precise dependence of the couplings on φ is shown in Fig. 7.9(e). (a) Critical temperature from semiclassical Monte Carlo simulations, obtained by fitting the peaks in the specific heat. (b) Ground-state energy per site ϵ_0 from numerical minimization (gray line), compared to analytically determined energies of the different observed phases. (c) Critical scale from pf-FRG in the same parameter regimes.

Bibliography

- [P1] **L. Gresista**, D. Kiese, S. Trebst, and M. M. Scherer, Spin-valley magnetism on the triangular moiré lattice with SU(4) breaking interactions, *Physical Review B* **108**, 045102 (2023).
- [P2] J. Park, **L. Gresista**, S. Trebst, A. Rosch, and J. Park, Network of chiral one-dimensional channels and localized states emerging in a moiré system, *2D Materials* **10**, 035033 (2023).
- [P3] **L. Gresista**, D. Kiese, and S. Trebst, Moments and multiplets in moiré materials, *The European Physical Journal B* **95**, 119 (2022).
- [P4] **L. Gresista**, D. Lozano-Gómez, M. Vojta, S. Trebst, and Y. Iqbal, Quantum effects on pyrochlore higher-rank U(1) spin liquids: Pinch-line singularities, spin nematics, and connections to oxide materials, *Physical Review Research* **7**, 033109 (2025).
- [P5] **L. Gresista**, C. Hickey, S. Trebst, and Y. Iqbal, Candidate quantum disordered intermediate phase in the Heisenberg antiferromagnet on the maple-leaf lattice, *Physical Review B* **108**, L241116 (2023).
- [P6] M. Gembé, **L. Gresista**, H.-J. Schmidt, C. Hickey, Y. Iqbal, and S. Trebst, Noncoplanar orders and quantum disordered states in maple-leaf antiferromagnets, *Physical Review B* **110**, 085151 (2024).
- [U1] **L. Gresista**, D. Kiese, S. Trebst, and Y. Iqbal, Unconventional orders in the maple-leaf ferro-antiferromagnetic Heisenberg model, *in preparation*.
- [C1] **L. Gresista** and D. Kiese, lgresista/SemiClassicalMC: SemiClassicalMC v.0.1.0, Zenodo (2023).
- [C2] **L. Gresista**, ClusterMeanFieldTheory.jl repository, <https://github.com/lgresista/ClusterMeanFieldTheory.jl>.
- [C3] D. Kiese, T. Müller, and **L. Gresista**, dominikkiese/PFFRGSolver.jl: v0.5.1, Zenodo (2023).
- [1] S. Blundell, *Magnetism in Condensed Matter*. Oxford University Press, 2011.
- [2] A. Fert, N. Reyren, and V. Cros, Magnetic skyrmions: Advances in physics and potential applications, *Nature Reviews Materials* **2**, 17031 (2017).
- [3] K. Guslienko, Magnetic Hopfions: A Review, *Magnetism* **4**, 383–399 (2024).
- [4] L. Savary and L. Balents, Quantum spin liquids: A review, *Reports on Progress in Physics* **80**, 016502 (2016).
- [5] P. Anderson, Resonating valence bonds: A new kind of insulator? *Materials Research Bulletin* **8**, 153–160 (1973).
- [6] X. G. Wen, TOPOLOGICAL ORDERS IN RIGID STATES, *International Journal of Modern Physics B* **04**, 239–271 (1990).

- [7] X.-G. Wen, Quantum orders and symmetric spin liquids, *Physical Review B* **65**, 165113 (2002).
- [8] A. Yu. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2–30 (2003).
- [9] D. Aasen et al. Roadmap to fault tolerant quantum computation using topological qubit arrays. arXiv: 2502.12252.
- [10] P. W. Anderson, The Resonating Valence Bond State in La_2CuO_4 and Superconductivity, *Science* **235**, 1196–1198 (1987).
- [11] A. Mann, High-temperature superconductivity at 25: Still in suspense, *Nature* **475**, 280–282 (2011).
- [12] P. A. Lee, N. Nagaosa, and X.-G. Wen, Doping a Mott insulator: Physics of high-temperature superconductivity, *Reviews of Modern Physics* **78**, 17–85 (2006).
- [13] J. Reuther and P. Wölfle, $J_1 - J_2$ frustrated two-dimensional Heisenberg model: Random phase approximation and functional renormalization group, *Physical Review B* **81**, 144410 (2010).
- [14] T. Müller, D. Kiese, N. Niggemann, B. Sbierski, J. Reuther, S. Trebst, R. Thomale, and Y. Iqbal, Pseudo-fermion functional renormalization group for spin models, *Reports on Progress in Physics* **87**, 036501 (2024).
- [15] A. Kitaev, Anyons in an exactly solved model and beyond, *Annals of Physics* **321**, 2–111 (2006).
- [16] J. Villain, R. Bidaux, J.-P. Carton, and R. Conte, Order as an effect of disorder, *Journal de Physique* **41**, 1263–1272 (1980).
- [17] J. T. Chalker, P. C. W. Holdsworth, and E. F. Shender, Hidden order in a frustrated system: Properties of the Heisenberg Kagomé antiferromagnet, *Physical Review Letters* **68**, 855–858 (1992).
- [18] R. Moessner and J. T. Chalker, Properties of a Classical Spin Liquid: The Heisenberg Pyrochlore Antiferromagnet, *Physical Review Letters* **80**, 2929–2932 (1998).
- [19] L. Pauling, The Structure and Entropy of Ice and of Other Crystals with Some Randomness of Atomic Arrangement, *Journal of the American Chemical Society* **57**, 2680–2684 (1935).
- [20] S. T. Bramwell and M. J. P. Gingras, Spin Ice State in Frustrated Magnetic Pyrochlore Materials, *Science* **294**, 1495–1501 (2001).
- [21] M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K. W. Godfrey, Geometrical Frustration in the Ferromagnetic Pyrochlore $\text{Ho}_2\text{Ti}_2\text{O}_7$, *Physical Review Letters* **79**, 2554–2557 (1997).
- [22] 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, Magnetic Coulomb phase in the spin ice $\text{Ho}_2\text{Ti}_2\text{O}_7$, *Science* **326**, 415–417 (2009).
- [23] D. J. P. Morris et al., Dirac Strings and Magnetic Monopoles in the Spin Ice $\text{Dy}_2\text{Ti}_2\text{O}_7$, *Science* **326**, 411–414 (2009).
- [24] J. G. Rau and M. J. Gingras, Frustrated Quantum Rare-Earth Pyrochlores, *Annual Review of Condensed Matter Physics* **10**, 357–386 (2019).

[25] K. T. K. Chung. Mapping the Phase Diagram of a Frustrated Magnet: Degeneracies, Flat Bands, and Canting Cycles on the Pyrochlore Lattice. arXiv: 2411.03429.

[26] D. Lozano-Gómez, O. Benton, M. J. P. Gingras, and H. Yan. An Atlas of Classical Pyrochlore Spin Liquids. arXiv: 2411.03547.

[27] H. Yan, O. Benton, R. Moessner, and A. H. Nevidomskyy, Classification of classical spin liquids: Typology and resulting landscape, *Physical Review B* **110**, L020402 (2024).

[28] H. Yan, O. Benton, A. H. Nevidomskyy, and R. Moessner, Classification of classical spin liquids: Detailed formalism and suite of examples, *Physical Review B* **109**, 174421 (2024).

[29] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Valence bond ground states in isotropic quantum antiferromagnets, *Communications in Mathematical Physics* **115**, 477–528 (1988).

[30] P. Ghosh, T. Müller, and R. Thomale, Another exact ground state of a two-dimensional quantum antiferromagnet, *Physical Review B* **105**, L180412 (2022).

[31] P. Ghosh, Triplon analysis of magnetic disorder and order in maple-leaf Heisenberg magnet, *Journal of Physics: Condensed Matter* **36**, 455803 (2024).

[32] Y. Zhou, K. Kanoda, and T.-K. Ng, Quantum spin liquid states, *Reviews of Modern Physics* **89**, 025003 (2017).

[33] M. Hermele, M. P. A. Fisher, and L. Balents, Pyrochlore photons: The $U(1)$ spin liquid in a $S = 1/2$ three-dimensional frustrated magnet, *Physical Review B* **69**, 064404 (2004).

[34] L. Capriotti, A. E. Trumper, and S. Sorella, Long-Range Néel Order in the Triangular Heisenberg Model, *Physical Review Letters* **82**, 3899–3902 (1999).

[35] S. R. White and A. L. Chernyshev, Ne’el Order in Square and Triangular Lattice Heisenberg Models, *Physical Review Letters* **99**, 127004 (2007).

[36] N. Astrakhantsev, T. Westerhout, A. Tiwari, K. Choo, A. Chen, M. H. Fischer, G. Carleo, and T. Neupert, Broken-Symmetry Ground States of the Heisenberg Model on the Pyrochlore Lattice, *Physical Review X* **11**, 041021 (2021).

[37] I. Hagymási, R. Schäfer, R. Moessner, and D. J. Luitz, Possible Inversion Symmetry Breaking in the $S = 1/2$ Pyrochlore Heisenberg Magnet, *Physical Review Letters* **126**, 117204 (2021).

[38] G. Khaliullin, Orbital Order and Fluctuations in Mott Insulators, *Progress of Theoretical Physics Supplement* **160**, 155–202 (2005).

[39] G. Jackeli and G. Khaliullin, Mott Insulators in the Strong Spin-Orbit Coupling Limit: From Heisenberg to a Quantum Compass and Kitaev Models, *Physical Review Letters* **102**, 017205 (2009).

[40] S. Trebst and C. Hickey, Kitaev materials, *Physics Reports* **950**, 1–37 (2022).

[41] J. Wen, S.-L. Yu, S. Li, W. Yu, and J.-X. Li, Experimental identification of quantum spin liquids, *npj Quantum Materials* **4**, 12 (2019).

[42] T.-H. Han, J. S. Helton, S. Chu, D. G. Nocera, J. A. Rodriguez-Rivera, C. Broholm, and Y. S. Lee, Fractionalized excitations in the spin-liquid state of a kagome-lattice antiferromagnet, *Nature* **492**, 406–410 (2012).

[43] M. Fu, T. Imai, T.-H. Han, and Y. S. Lee, Evidence for a gapped spin-liquid ground state in a kagome Heisenberg antiferromagnet, *Science* **350**, 655–658 (2015).

- [44] K. W. Plumb, J. P. Clancy, L. J. Sandilands, V. V. Shankar, Y. F. Hu, K. S. Burch, H.-Y. Kee, and Y.-J. Kim, α – RuCl₃: A spin-orbit assisted Mott insulator on a honeycomb lattice, *Physical Review B* **90**, 041112 (2014).
- [45] A. Banerjee et al., Excitations in the field-induced quantum spin liquid state of -RuCl₃, *npj Quantum Materials* **3**, 8 (2018).
- [46] M. G. Yamada, M. Oshikawa, and G. Jackeli, Emergent SU(4) Symmetry in α –ZrCl₃ and Crystalline Spin-Orbital Liquids, *Physical Review Letters* **121**, 097201 (2018).
- [47] M. G. Yamada, M. Oshikawa, and G. Jackeli, SU(4)-symmetric quantum spin-orbital liquids on various lattices, *Physical Review B* **104**, 224436 (2021).
- [48] J. W. F. Venderbos and R. M. Fernandes, Correlations and electronic order in a two-orbital honeycomb lattice model for twisted bilayer graphene, *Physical Review B* **98**, 245103 (2018).
- [49] Y.-H. Zhang and T. Senthil, Bridging Hubbard model physics and quantum Hall physics in trilayer graphene/h-BN moiré superlattice, *Physical Review B* **99**, 205150 (2019).
- [50] K. I. Kugel' and D. I. Khomskii, The Jahn-Teller effect and magnetism: Transition metal compounds, *Soviet Physics Uspekhi* **25**, 231–256 (1982).
- [51] I. Affleck and J. B. Marston, Large-n limit of the Heisenberg-Hubbard model: Implications for high- T_c superconductors, *Physical Review B* **37**, 3774–3777 (1988).
- [52] C. A. Gallegos, S. Jiang, S. R. White, and A. L. Chernyshev, Phase Diagram of the Easy-Axis Triangular-Lattice J1 - J2 Model, *Physical Review Letters* **134**, 196702 (2025).
- [53] A. Keselman, L. Savary, and L. Balents, Dimer description of the SU(4) antiferromagnet on the triangular lattice, *SciPost Physics* **8**, 076 (2020).
- [54] A. W. Sandvik, Computational Studies of Quantum Spin Systems, *AIP Conference Proceedings* **1297**, 135–338 (2010).
- [55] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, *Physical Review Letters* **94**, 170201 (2005).
- [56] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Reviews of Modern Physics* **93**, 045003 (2021).
- [57] M. B. Hastings, An area law for one-dimensional quantum systems, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P08024 (2007).
- [58] P. Calabrese and J. Cardy, Entanglement entropy and quantum field theory, *Journal of Statistical Mechanics: Theory and Experiment* **2004**, P06002 (2004).
- [59] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Annals of Physics* **326**, 96–192 (2011).
- [60] R. Orús, Tensor networks for complex quantum systems, *Nature Reviews Physics* **1**, 538–550 (2019).
- [61] P. Corboz and F. Mila, Tensor network study of the Shastry-Sutherland model in zero magnetic field, *Physical Review B* **87**, 115144 (2013).
- [62] L. Wang, Z.-C. Gu, F. Verstraete, and X.-G. Wen, Tensor-product state approach to spin- $\frac{1}{2}$ square J_1-J_2 antiferromagnetic Heisenberg model: Evidence for deconfined quantum criticality, *Physical Review B* **94**, 075143 (2016).

- [63] R. Haghshenas, S.-S. Gong, and D. N. Sheng, Single-layer tensor network study of the Heisenberg model with chiral interactions on a kagome lattice, *Physical Review B* **99**, 174423 (2019).
- [64] A. A. Abrikosov, Electron scattering on magnetic impurities in metals and anomalous resistivity effects, *Physics Physique Fizika* **2**, 5–20 (1965).
- [65] D. P. Arovas and A. Auerbach, Functional integral theories of low-dimensional quantum Heisenberg models, *Physical Review B* **38**, 316–332 (1988).
- [66] C. Wetterich, Exact evolution equation for the effective potential, *Physics Letters B* **301**, 90–94 (1993).
- [67] M. L. Baez and J. Reuther, Numerical treatment of spin systems with unrestricted spin length S : A functional renormalization group study, *Physical Review B* **96**, 045144 (2017).
- [68] F. L. Buessen, D. Roscher, S. Diehl, and S. Trebst, Functional renormalization group approach to $SU(N)$ Heisenberg models: Real-space renormalization group at arbitrary N , *Physical Review B* **97**, 064415 (2018).
- [69] R. P. Feynman, R. B. Leighton, and M. L. Sands, *The Feynman Lectures on Physics - Vol. 2. Basic Books*, 2011.
- [70] A. Altland and B. Simons, *Condensed Matter Field Theory*. Cambridge University Press, 2023.
- [71] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim, The electronic properties of graphene, *Reviews of Modern Physics* **81**, 109–162 (2009).
- [72] R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, 2004.
- [73] N. Marzari, A. A. Mostofi, J. R. Yates, I. Souza, and D. Vanderbilt, Maximally localized Wannier functions: Theory and applications, *Reviews of Modern Physics* **84**, 1419–1475 (2012).
- [74] F. Aryasetiawan, M. Imada, A. Georges, G. Kotliar, S. Biermann, and A. I. Lichtenstein, Frequency-dependent local interactions and low-energy effective models from electronic structure calculations, *Physical Review B* **70**, 195104 (2004).
- [75] C. L. Cleveland and R. Medina A., Obtaining a Heisenberg Hamiltonian from the Hubbard model, *American Journal of Physics* **44**, 44–46 (1976).
- [76] Y. Kuramoto, *Perturbation Theory and Effective Hamiltonian*, Y. Kuramoto, Ed. Springer Japan, 2020.
- [77] R. Bistritzer and A. H. MacDonald, Moiré bands in twisted double-layer graphene, *Proceedings of the National Academy of Sciences* **108**, 12233–12237 (2011).
- [78] E. Y. Andrei and A. H. MacDonald, Graphene bilayers with a twist, *Nature Materials* **19**, 1265–1275 (2020).
- [79] J. M. Luttinger and L. Tisza, Theory of Dipole Interaction in Crystals, *Physical Review* **70**, 954–964 (1946).
- [80] J. M. Luttinger, A Note on the Ground State in Antiferromagnetics, *Physical Review* **81**, 1015–1018 (1951).
- [81] D. H. Lyons and T. A. Kaplan, Method for Determining Ground-State Spin Configurations, *Physical Review* **120**, 1580–1585 (1960).

- [82] T. A. Kaplan and N. Menyuk, Spin ordering in three-dimensional crystals with strong competing exchange interactions, *Philosophical Magazine* **87**, 3711–3785 (2007).
- [83] H.-J. Schmidt and J. Richter, Classical ground states of spin lattices, *Journal of Physics A: Mathematical and Theoretical* **55**, 465005 (2022).
- [84] T. L. C. Müller, Quantum magnetism in three dimensions: Exploring phase diagrams and real materials using Functional Renormalization, *Universität Würzburg*, 2023.
- [85] K. Millard and H. S. Leff, Infinite-Spin Limit of the Quantum Heisenberg Model, *Journal of Mathematical Physics* **12**, 1000–1005 (1971).
- [86] E. H. Lieb, The classical limit of quantum spin systems, *Communications in Mathematical Physics* **31**, 327–340 (1973).
- [87] I. Kimchi and A. Vishwanath, Kitaev-Heisenberg models for iridates on the triangular, hyperkagome, kagome, fcc, and pyrochlore lattices, *Physical Review B* **89**, 014414 (2014).
- [88] N. Niggemann, M. Hering, and J. Reuther, Classical spiral spin liquids as a possible route to quantum spin liquids, *Journal of Physics: Condensed Matter* **32**, 024001 (2019).
- [89] C. Xu and L. Balents, Topological Superconductivity in Twisted Multilayer Graphene, *Physical Review Letters* **121**, 087001 (2018).
- [90] H. C. Po, L. Zou, A. Vishwanath, and T. Senthil, Origin of Mott Insulating Behavior and Superconductivity in Twisted Bilayer Graphene, *Physical Review X* **8**, 031089 (2018).
- [91] L. Chen, H. Hu, and Q. Si, Fragile Insulator and Electronic Nematicity in a Graphene Moire System. *arXiv: 2007.06086*.
- [92] L. Classen, C. Honerkamp, and M. M. Scherer, Competing phases of interacting electrons on triangular lattices in moiré heterostructures, *Physical Review B* **99**, 195120 (2019).
- [93] E. M. Stoudenmire, S. Trebst, and L. Balents, Quadrupolar correlations and spin freezing in $S = 1$ triangular lattice antiferromagnets, *Physical Review B* **79**, 214436 (2009).
- [94] C. Hickey and A. Paramekanti, Thermal Phase Transitions of Strongly Correlated Bosons with Spin-Orbit Coupling, *Physical Review Letters* **113**, 265302 (2014).
- [95] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge University Press, 2014.
- [96] F. F. Carsten Bauer, BinningAnalysis.jl repository, <https://github.com/carstenbauer/BinningAnalysis.jl>.
- [97] M. E. Muller, A note on a method for generating points uniformly on n-dimensional spheres, *Commun. ACM* **2**, 19–20 (1959).
- [98] J. D. Alzate-Cardona, D. Sabogal-Suárez, R. F. L. Evans, and E. Restrepo-Parra, Optimal phase space sampling for Monte Carlo simulations of Heisenberg spin systems, *Journal of Physics: Condensed Matter* **31**, 095802 (2019).
- [99] S. Geman and D. Geman, Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-6**, 721–741 (1984).
- [100] S. R. Sklan and C. L. Henley, Nonplanar ground states of frustrated antiferromagnets on an octahedral lattice, *Physical Review B* **88**, 024407 (2013).
- [101] M. F. Lapa and C. L. Henley, Ground States of the Classical Antiferromagnet on the Pyrochlore Lattice. *arXiv: 1210.6810*.

[102] S. D. Axen, M. Baran, and R. Bergmann, Manifolds.jl (v0.10.23), Zenodo (2021).

[103] R. Bergmann, Manopt.jl (v0.5.23), Zenodo (2025).

[104] D. Kiese, F. L. Buessen, C. Hickey, S. Trebst, and M. M. Scherer, Emergence and stability of spin-valley entangled quantum liquids in moiré heterostructures, *Physical Review Research* **2**, 013370 (2020).

[105] Y.-Z. Ren, N.-H. Tong, and X.-C. Xie, Cluster mean-field theory study of J_1-J_2 Heisenberg model on a square lattice, *Journal of Physics: Condensed Matter* **26**, 115601 (2014).

[106] D. Yamamoto, G. Marmorini, and I. Danshita, Quantum Phase Diagram of the Triangular-Lattice X X Z Model in a Magnetic Field, *Physical Review Letters* **112**, 127203 (2014).

[107] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *Journal of Research of the National Bureau of Standards* **45**, 255 (1950).

[108] J. Haegeman, KrylovKit.jl repository, <https://github.com/Jutho/KrylovKit.jl>.

[109] D. Yamamoto, A. Masaki, and I. Danshita, Quantum phases of hardcore bosons with long-range interactions on a square lattice, *Physical Review B* **86**, 054516 (2012).

[110] X. Qian and M. Qin, Absence of spin liquid phase in the $J_1 - J_2$ Heisenberg model on the square lattice, *Physical Review B* **109**, L161103 (2024).

[111] W.-Y. Liu, S.-S. Gong, Y.-B. Li, D. Poilblanc, W.-Q. Chen, and Z.-C. Gu, Gapless quantum spin liquid and global phase diagram of the spin-1/2 $J_1 - J_2$ square antiferromagnetic Heisenberg model, *Science Bulletin* **67**, 1034–1041 (2022).

[112] H.-C. Jiang, H. Yao, and L. Balents, Spin liquid ground state of the spin- $\frac{1}{2}$ square J_1-J_2 Heisenberg model, *Physical Review B* **86**, 024424 (2012).

[113] J. Richter and J. Schulenburg, The spin-1/2 J_1-J_2 Heisenberg antiferromagnet on the square lattice:Exact diagonalization for $N=40$ spins, *The European Physical Journal B* **73**, 117–124 (2010).

[114] P. Kopietz, L. Bartosch, and F. Schütz, *Introduction to the Functional Renormalization Group*. Springer Berlin, 2013.

[115] F. L. Büßen, *A Functional Renormalization Group Perspective on Quantum Spin Liquids in Three-Dimensional Frustrated Magnets*, Universität zu Köln, 2019.

[116] D. Kiese, *Effective field theories for strongly correlated fermions - Insights from the functional renormalization group*, Universität zu Köln, 2022.

[117] F. L. Buessen, The SpinParser software for pseudofermion functional renormalization group calculations on quantum magnets, *SciPost Physics Codebases*, 5 (2022).

[118] F. L. Buessen, V. Noculak, S. Trebst, and J. Reuther, Functional renormalization group for frustrated magnets with nondiagonal spin interactions, *Physical Review B* **100**, 125164 (2019).

[119] A. Revelli et al., Fingerprints of Kitaev physics in the magnetic excitations of honeycomb iridates, *Physical Review Research* **2**, 043094 (2020).

[120] F. L. Buessen and Y. B. Kim, Functional renormalization group study of the Kitaev-model on the honeycomb lattice and emergent incommensurate magnetic correlations, *Physical Review B* **103**, 184407 (2021).

- [121] V. Noculak, D. Lozano-Gómez, J. Oitmaa, R. R. P. Singh, Y. Iqbal, M. J. P. Gingras, and J. Reuther, Classical and quantum phases of the pyrochlore $S = \frac{1}{2}$ magnet with Heisenberg and Dzyaloshinskii-Moriya interactions, *Physical Review B* **107**, 214414 (2023).
- [122] D. Lozano-Gómez, V. Noculak, J. Oitmaa, R. R. P. Singh, Y. Iqbal, J. Reuther, and M. J. P. Gingras, Competing gauge fields and entropically driven spin liquid to spin liquid transition in non-Kramers pyrochlores, *Proceedings of the National Academy of Sciences* **121**, e2403487121 (2024).
- [123] R. Hedden, V. Meden, T. Pruschke, and K. Schönhammer, A functional renormalization group approach to zero-dimensional interacting systems, *Journal of Physics: Condensed Matter* **16**, 5279 (2004).
- [124] C. Honerkamp and M. Salmhofer, Temperature-flow renormalization group and the competition between superconductivity and ferromagnetism, *Physical Review B* **64**, 184516 (2001).
- [125] C. Platt, W. Hanke, and R. Thomale, Functional renormalization group for multi-orbital Fermi surface instabilities, *Advances in Physics* **62**, 453–562 (2013).
- [126] A. A. Katanin, Fulfillment of Ward identities in the functional renormalization group approach, *Physical Review B* **70**, 115109 (2004).
- [127] F. B. Kugler and J. von Delft, Derivation of exact flow equations from the self-consistent parquet relations, *New Journal of Physics* **20**, 123029 (2018).
- [128] F. B. Kugler and J. von Delft, Multiloop functional renormalization group for general models, *Physical Review B* **97**, 035162 (2018).
- [129] D. Kiese, T. Müller, Y. Iqbal, R. Thomale, and S. Trebst, Multiloop functional renormalization group approach to quantum spin systems, *Physical Review Research* **4**, 023185 (2022).
- [130] J. Thoenniss, M. K. Ritter, F. B. Kugler, J. von Delft, and M. Punk, Multiloop pseudofermion functional renormalization for quantum spin systems: Application to the spin- $\frac{1}{2}$ kagome Heisenberg model, [arXiv: 2011.01268](https://arxiv.org/abs/2011.01268).
- [131] I. T. "Diatlov, V. V. Sudakov, and K. A. Ter-Martirosian, Asymptotic Meson-Meson Scattering Theory, *Sov. Phys. JETP* **5**, 631–642 (1957).
- [132] I. D. Yu. A. Bychkov L.P. Gor'kov, Possibility of superconductivity type phenomena in a one-dimensional system, *Sov. Phys. JETP* **23**, 489 (1965).
- [133] J. Krieg and P. Kopietz, Exact renormalization group for quantum spin systems, *Physical Review B* **99**, 060403 (2019).
- [134] J. Reuther and R. Thomale, Functional renormalization group for the anisotropic triangular antiferromagnet, *Physical Review B* **83**, 024402 (2011).
- [135] N. Wentzell, G. Li, A. Tagliavini, C. Taranto, G. Rohringer, K. Held, A. Toschi, and S. Andergassen, High-frequency asymptotics of the vertex function: Diagrammatic parametrization and algorithmic implementation, *Physical Review B* **102**, 085106 (2020).
- [136] D. Roscher, F. L. Buessen, M. M. Scherer, S. Trebst, and S. Diehl, Functional renormalization group approach to SU(N) Heisenberg models: Momentum-space renormalization group for the large-N limit, *Physical Review B* **97**, 064416 (2018).
- [137] V. N. Popov and S. A. Fedotov, The functional-integrationmethod and diagram technique for spin systems, *Zh. Eksp. Teor. Fiz.* **94**, 183–194 (1998).

[138] B. Schneider, D. Kiese, and B. Sbierski, Taming pseudofermion functional renormalization for quantum spins: Finite temperatures and the Popov-Fedotov trick, *Physical Review B* **106**, 235113 (2022).

[139] B. Schneider, J. Reuther, M. G. Gonzalez, B. Sbierski, and N. Niggemann, Temperature flow in pseudo-Majorana functional renormalization for quantum spins, *Physical Review B* **109**, 195109 (2024).

[140] N. Niggemann, B. Sbierski, and J. Reuther, Frustrated Quantum Spins at finite Temperature: Pseudo-Majorana functional RG approach, *Physical Review B* **103**, 104431 (2021).

[141] N. Niggemann, J. Reuther, and B. Sbierski, Quantitative functional renormalization for three-dimensional quantum Heisenberg models, *SciPost Physics* **12**, 156 (2022).

[142] J. Beyer, F. Goth, and T. Müller, Better integrators for functional renormalization group calculations, *The European Physical Journal B* **95**, 116 (2022).

[143] P. Bogacki and L. F. Shampine, A 3(2) pair of Runge - Kutta formulas, *Applied Mathematics Letters* **2**, 321–325 (1989).

[144] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes* 3rd Edition: The Art of Scientific Computing. Cambridge University Press, 2007.

[145] J. S. Dai, Euler–Rodrigues formula variations, quaternion conjugation and intrinsic connections, *Mechanism and Machine Theory* **92**, 144–152 (2015).

[146] D. Kiese, Y. He, C. Hickey, A. Rubio, and D. M. Kennes, TMDs as a platform for spin liquid physics: A strong coupling study of twisted bilayer WSe₂, *APL Materials* **10**, 031113 (2022).

[147] Y. Iqbal, W.-J. Hu, R. Thomale, D. Poilblanc, and F. Becca, Spin liquid nature in the Heisenberg J1 – J2 triangular antiferromagnet, *Physical Review B* **93**, 144411 (2016).

[148] Y. Iqbal, T. Müller, P. Ghosh, M. J. P. Gingras, H. O. Jeschke, S. Rachel, J. Reuther, and R. Thomale, Quantum and Classical Phases of the Pyrochlore Heisenberg Model with Competing Interactions, *Physical Review X* **9**, 011005 (2019).

[149] B. Canals and C. Lacroix, Quantum spin liquid: The Heisenberg antiferromagnet on the three-dimensional pyrochlore lattice, *Physical Review B* **61**, 1149–1159 (2000).

[150] Y. Huang, K. Chen, Y. Deng, N. Prokof'ev, and B. Svistunov, Spin-Ice State of the Quantum Heisenberg Antiferromagnet on the Pyrochlore Lattice, *Physical Review Letters* **116**, 177203 (2016).

[151] X.-G. Wen, *Quantum Field Theory of Many-Body Systems: From the Origin of Sound to an Origin of Light and Electrons*. Oxford University Press, 2004.

[152] H. Yan, O. Benton, L. Jaubert, and N. Shannon, Theory of multiple-phase competition in pyrochlore magnets with anisotropic exchange with application to Yb₂Ti₂O₇, Er₂Ti₂O₇, and Er₂Sn₂O₇, *Physical Review B* **95**, 094422 (2017).

[153] K. T. K. Chung, J. S. K. Goh, A. Mukherjee, W. Jin, D. Lozano-Gómez, and M. J. P. Gingras, Probing Flat Band Physics in Spin Ice Systems via Polarized Neutron Scattering, *Physical Review Letters* **128**, 107201 (2022).

[154] S. Chillal et al., Evidence for a three-dimensional quantum spin liquid in PbCuTe₂O₆, *Nature Communications* **11**, 2348 (2020).

[155] I. Živković et al., Magnetic Field Induced Quantum Spin Liquid in the Two Coupled Trillium Lattices of K₂Ni₂(SO₄)₃, *Physical Review Letters* **127**, 157204 (2021).

- [156] C. Balz et al., Physical realization of a quantum spin liquid based on a complex frustration mechanism, *Nature Physics* **12**, 942–949 (2016).
- [157] V. Noculak and J. Reuther, Pseudo-fermion functional renormalization group with magnetic fields, *Physical Review B* **109**, 174414 (2024).
- [158] J. Potten, Y. Iqbal, R. Thomale, and T. Müller. Keldysh pseudo-fermion functional renormalization group for quantum magnetism. arXiv: 2503.11596.
- [159] Y. Schaden, M. G. Gonzalez, and J. Reuther, Phase diagram of the XXZ pyrochlore model from pseudo-Majorana functional renormalization group, *Physical Review B* **111**, 134442 (2025).
- [160] Y. Schaden and J. Reuther, Bilinear Majorana representations for spin operators with spin magnitudes $\$S>1/2\$$, *Physical Review Research* **5**, 023067 (2023).
- [161] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Reviews of Modern Physics* **68**, 13–125 (1996).
- [162] G. Rohringer, H. Hafermann, A. Toschi, A. A. Katanin, A. E. Antipov, M. I. Katsnelson, A. I. Lichtenstein, A. N. Rubtsov, and K. Held, Diagrammatic routes to nonlocal correlations beyond dynamical mean field theory, *Reviews of Modern Physics* **90**, 025003 (2018).
- [163] K. Held, A. A. Katanin, and A. Toschi, Dynamical Vertex Approximation: An Introduction, *Progress of Theoretical Physics Supplement* **176**, 117–133 (2008).
- [164] D. Vilardi, C. Taranto, and W. Metzner, Antiferromagnetic and d-wave pairing correlations in the strongly interacting two-dimensional Hubbard model from the functional renormalization group, *Physical Review B* **99**, 104501 (2019).
- [165] R. Goll, D. Tarasevych, J. Krieg, and P. Kopietz, Spin functional renormalization group for quantum Heisenberg ferromagnets: Magnetization and magnon damping in two dimensions, *Physical Review B* **100**, 174424 (2019).
- [166] D. Tarasevych and P. Kopietz, Critical spin dynamics of Heisenberg ferromagnets revisited, *Physical Review B* **105**, 024403 (2022).
- [167] A. Rückriegel, J. Arnold, R. Goll, and P. Kopietz, Spin functional renormalization group for dimerized quantum spin systems, *Physical Review B* **105**, 224406 (2022).
- [168] D. Tarasevych, A. Rückriegel, S. Keupert, V. Mitsiioannou, and P. Kopietz, Spin-functional renormalization group for the $J_1J_2J_3$ quantum Heisenberg model, *Physical Review B* **106**, 174412 (2022).
- [169] P. W. Anderson, Ordering and Antiferromagnetism in Ferrites, *Physical Review* **102**, 1008–1013 (1956).
- [170] R. M. Nandkishore and M. Hermele, Fractons, *Annual Review of Condensed Matter Physics* **10**, 295–313 (2019).
- [171] M. J. P. Gingras and P. A. McClarty, Quantum spin ice: A search for gapless quantum spin liquids in pyrochlore magnets, *Reports on Progress in Physics* **77**, 056501 (2014).
- [172] N. Niggemann, Quantum Effects on Unconventional Pinch Point Singularities, *Physical Review Letters* **130** (2023).
- [173] O. Benton, O. Sikora, and N. Shannon, Seeing the light: Experimental signatures of emergent electromagnetism in a quantum spin ice, *Physical Review B* **86**, 075154 (2012).

[174] K. A. Ross, L. Savary, B. D. Gaulin, and L. Balents, Quantum Excitations in Quantum Spin Ice, *Physical Review X* **1**, 021002 (2011).

[175] L. Pan, N. J. Laurita, K. A. Ross, B. D. Gaulin, and N. P. Armitage, A measure of monopole inertia in the quantum spin ice $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Nature Physics* **12**, 361–366 (2016).

[176] K. A. Ross, J. P. C. Ruff, C. P. Adams, J. S. Gardner, H. A. Dabkowska, Y. Qiu, J. R. D. Copley, and B. D. Gaulin, Two-Dimensional Kagome Correlations and Field Induced Order in the Ferromagnetic XY Pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review Letters* **103**, 227202 (2009).

[177] J. Gaudet, K. A. Ross, E. Kermarrec, N. P. Butch, G. Ehlers, H. A. Dabkowska, and B. D. Gaulin, Gapless quantum excitations from an icelike splayed ferromagnetic ground state in stoichiometric $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review B* **93**, 064406 (2016).

[178] J. D. Thompson, P. A. McClarty, D. Prabhakaran, I. Cabrera, T. Guidi, and R. Coldea, Quasiparticle Breakdown and Spin Hamiltonian of the Frustrated Quantum Pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$ in a Magnetic Field, *Physical Review Letters* **119**, 057203 (2017).

[179] A. Scheie, J. Kindervater, S. Zhang, H. J. Changlani, G. Sala, G. Ehlers, A. Heinemann, G. S. Tucker, S. M. Koohpayeh, and C. Broholm, Multiphase magnetism in $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Proceedings of the National Academy of Sciences* **117**, 27245–27254 (2020).

[180] A. Scheie, O. Benton, M. Taillefumier, L. D. C. Jaubert, G. Sala, N. Jalarvo, S. M. Koohpayeh, and N. Shannon, Dynamical Scaling as a Signature of Multiple Phase Competition in $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review Letters* **129**, 217202 (2022).

[181] O. Benton, L. D. C. Jaubert, H. Yan, and N. Shannon, A spin-liquid with pinch-line singularities on the pyrochlore lattice, *Nature Communications* **7**, 11572 (2016).

[182] D. A. Garanin and B. Canals, Classical spin liquid: Exact solution for the infinite-component antiferromagnetic model on the *kagomé* lattice, *Physical Review B* **59**, 443–456 (1999).

[183] B. Canals and D. A. Garanin, Spin-liquid phase in the pyrochlore anti-ferromagnet, *Canadian Journal of Physics* **79**, 1323–1331 (2001).

[184] J. Gaudet, D. D. Maharaj, G. Sala, E. Kermarrec, K. A. Ross, H. A. Dabkowska, A. I. Kolesnikov, G. E. Granroth, and B. D. Gaulin, Neutron spectroscopic study of crystalline electric field excitations in stoichiometric and lightly stuffed $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review B* **92**, 134420 (2015).

[185] J. G. Rau and M. J. P. Gingras, Frustration and anisotropic exchange in ytterbium magnets with edge-shared octahedra, *Physical Review B* **98**, 054408 (2018).

[186] L. D. C. Jaubert, O. Benton, J. G. Rau, J. Oitmaa, R. R. P. Singh, N. Shannon, and M. J. P. Gingras, Are Multiphase Competition and Order by Disorder the Keys to Understanding $\text{Yb}_2\text{Ti}_2\text{O}_7$? *Physical Review Letters* **115**, 267208 (2015).

[187] N. Francini, L. Janssen, and D. Lozano-Gómez, Higher-rank spin liquids and spin nematics from competing orders in pyrochlore magnets, *Physical Review B* **111**, 085140 (2025).

[188] H. Yan, O. Benton, A. H. Nevidomskyy, and R. Moessner, Classification of Classical Spin Liquids: Detailed Formalism and Suite of Examples. version 1.

[189] J. G. Rau, R. Moessner, and P. A. McClarty, Magnon interactions in the frustrated pyrochlore ferromagnet $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review B* **100**, 104423 (2019).

[190] D. F. Bowman et al., Role of defects in determining the magnetic ground state of ytterbium titanate, *Nature Communications* **10**, 637 (2019).

[191] J. Robert, E. Lhotel, G. Remenyi, S. Sahling, I. Mirebeau, C. Decorse, B. Canals, and S. Petit, Spin dynamics in the presence of competing ferromagnetic and antiferromagnetic correlations in $\text{Yb}_2\text{Ti}_2\text{O}_7$, *Physical Review B* **92**, 064425 (2015).

[192] H. Yan, O. Benton, L. D. C. Jaubert, and N. Shannon, Rank-2 $\text{U}(1)$ Spin Liquid on the Breathing Pyrochlore Lattice, *Physical Review Letters* **124**, 127203 (2020).

[193] S. Guitteny, S. Petit, E. Lhotel, J. Robert, P. Bonville, A. Forget, and I. Mirebeau, Palmer-Chalker correlations in the XY pyrochlore antiferromagnet $\text{er}_2\text{sn}_2\text{o}_7$, *Physical Review B* **88**, 134408 (2013).

[194] S. Hu, W. Zhu, S. Eggert, and Y.-C. He, Dirac Spin Liquid on the Spin- 1 / 2 Triangular Heisenberg Antiferromagnet, *Physical Review Letters* **123**, 207203 (2019).

[195] Y. Iqbal, F. Becca, S. Sorella, and D. Poilblanc, Gapless spin-liquid phase in the kagome spin- 1 / 2 Heisenberg antiferromagnet, *Physical Review B* **87**, 060405 (2013).

[196] Y.-C. He, M. P. Zaletel, M. Oshikawa, and F. Pollmann, Signatures of Dirac Cones in a DMRG Study of the Kagome Heisenberg Model, *Physical Review X* **7**, 031020 (2017).

[197] T. Fennell, J. O. Piatek, R. A. Stephenson, G. J. Nilsen, and H. M. Rønnow, Spangolite: An $s = 1/2$ maple leaf lattice antiferromagnet? *Journal of Physics: Condensed Matter* **23**, 164201 (2011).

[198] P. Schmoll, H. O. Jeschke, and Y. Iqbal, Tensor network analysis of the maple-leaf antiferromagnet spangolite, *Communications Materials* **6**, 178 (2025).

[199] S. J. Mills, A. R. Kampf, A. G. Christy, R. M. Housley, G. R. Rossman, R. E. Reynolds, and J. Marty, Bluebellite and mojaveite, two new minerals from the central Mojave Desert, California, USA, *Mineralogical Magazine* **78**, 1325–1340 (2014).

[200] Y. Haraguchi, A. Matsuo, K. Kindo, and Z. Hiroi, Quantum antiferromagnet bluebellite comprising a maple-leaf lattice made of spin- 1 / 2 Cu^{2+} ions, *Physical Review B* **104**, 174439 (2021).

[201] R. Makuta and C. Hotta, Dimensional reduction in quantum spin- 1 / 2 system on a 1 / 7 -depleted triangular lattice, *Physical Review B* **104**, 224415 (2021).

[202] P. Ghosh, T. Müller, Y. Iqbal, R. Thomale, and H. O. Jeschke, Effective spin-1 breathing kagome Hamiltonian induced by the exchange hierarchy in the maple leaf mineral bluebellite, *Physical Review B* **110**, 094406 (2024).

[203] A. R. Kampf, S. J. Mills, R. M. Housley, and J. Marty, Lead-tellurium oxysalts from Otto Mountain near Baker, California: VIII. Fuettererite, $\text{Pb}_3\text{Cu}_{62}\text{Te}_6\text{O}_6(\text{OH})_7\text{Cl}_5$, a new mineral with double spangolite-type sheets, *American Mineralogist* **98**, 506–511 (2013).

[204] F. Olmi, C. Sabelli, and R. Trosti-Ferroni, The crystal structure of sabelliite, *European Journal of Mineralogy* **7**, 1331–1338 (1995).

[205] D. Schmalfuß, P. Tomczak, J. Schulenburg, and J. Richter, The spin- $\frac{1}{2}$ Heisenberg antiferromagnet on a $\frac{1}{7}$ -depleted triangular lattice: Ground-state properties, *Physical Review B* **65**, 224405 (2002).

[206] B. Sriram Shastry and B. Sutherland, Exact ground state of a quantum mechanical antiferromagnet, *Physica B+C* **108**, 1069–1070 (1981).

[207] J. Y. Lee, Y.-Z. You, S. Sachdev, and A. Vishwanath, Signatures of a Deconfined Phase Transition on the Shastry-Sutherland Lattice: Applications to Quantum Critical SrCu_2 (BO_3) 2 , *Physical Review X* **9**, 041037 (2019).

[208] T. Senthil, A. Vishwanath, L. Balents, S. Sachdev, and M. P. A. Fisher, Deconfined Quantum Critical Points, *Science* **303**, 1490–1494 (2004).

[209] T. Senthil, L. Balents, S. Sachdev, A. Vishwanath, and M. P. A. Fisher, Quantum criticality beyond the Landau-Ginzburg-Wilson paradigm, *Physical Review B* **70**, 144407 (2004).

[210] J. Yang, A. W. Sandvik, and L. Wang, Quantum criticality and spin liquid phase in the Shastry-Sutherland model, *Physical Review B* **105**, L060409 (2022).

[211] L. L. Viteritti, R. Rende, A. Parola, S. Goldt, and F. Becca, Transformer wave function for two dimensional frustrated magnets: Emergence of a spin-liquid phase in the Shastry-Sutherland model, *Physical Review B* **111**, 134411 (2025).

[212] W.-Y. Liu, J. Hasik, S.-S. Gong, D. Poilblanc, W.-Q. Chen, and Z.-C. Gu, Emergence of Gapless Quantum Spin Liquid from Deconfined Quantum Critical Point, *Physical Review X* **12**, 031039 (2022).

[213] V. Kalmeyer and R. B. Laughlin, Equivalence of the resonating-valence-bond and fractional quantum Hall states, *Physical Review Letters* **59**, 2095–2098 (1987).

[214] C. Hickey, L. Cincio, Z. Papić, and A. Paramekanti, Haldane-Hubbard Mott Insulator: From Tetrahedral Spin Crystal to Chiral Spin Liquid, *Physical Review Letters* **116**, 137202 (2016).

[215] C. Hickey, L. Cincio, Z. Papić, and A. Paramekanti, Emergence of chiral spin liquids via quantum melting of noncoplanar magnetic orders, *Physical Review B* **96**, 115115 (2017).

[216] T. Momoi, K. Kubo, and K. Niki, Possible Chiral Phase Transition in Two-Dimensional Solid ${}^3\text{He}$, *Physical Review Letters* **79**, 2081–2084 (1997).

[217] K. Kubo and T. Momoi, Ground state of a spin system with two- and four-spin exchange interactions on the triangular lattice, *Zeitschrift für Physik B Condensed Matter* **103**, 485–489 (1997).

[218] S. D. Yi, S. Onoda, N. Nagaosa, and J. H. Han, Skyrmions and anomalous Hall effect in a Dzyaloshinskii-Moriya spiral magnet, *Physical Review B* **80**, 054416 (2009).

[219] J.-H. Park and J. H. Han, Zero-temperature phases for chiral magnets in three dimensions, *Physical Review B* **83**, 184406 (2011).

[220] T. Shimokawa, T. Okubo, and H. Kawamura, Multiple- q states of the $J_1 - J_2$ classical honeycomb-lattice Heisenberg antiferromagnet under a magnetic field, *Physical Review B* **100**, 224404 (2019).

[221] M. Mohylna, F. A. Gómez Albarracín, M. Žukovič, and H. D. Rosales, Spontaneous anti-ferromagnetic skyrmion/antiskyrmion lattice and spiral spin-liquid states in the frustrated triangular lattice, *Physical Review B* **106**, 224406 (2022).

[222] T. Cookmeyer, J. Motruk, and J. E. Moore, Four-Spin Terms and the Origin of the Chiral Spin Liquid in Mott Insulators on the Triangular Lattice, *Physical Review Letters* **127**, 087201 (2021).

[223] J.-C. Domenge, P. Sindzingre, C. Lhuillier, and L. Pierre, Twelve sublattice ordered phase in the $J_1 - J_2$ model on the kagomé lattice, *Physical Review B* **72**, 024433 (2005).

- [224] L. Messio, C. Lhuillier, and G. Misguich, Lattice symmetries and regular magnetic orders in classical frustrated antiferromagnets, *Physical Review B* **83**, 184401 (2011).
- [225] M. Gembé, H.-J. Schmidt, C. Hickey, J. Richter, Y. Iqbal, and S. Trebst, Noncoplanar magnetic order in classical square-kagome antiferromagnets, *Physical Review Research* **5**, 043204 (2023).
- [226] W.-J. Hu, W. Zhu, Y. Zhang, S. Gong, F. Becca, and D. N. Sheng, Variational Monte Carlo study of a chiral spin liquid in the extended Heisenberg model on the kagome lattice, *Physical Review B* **91**, 041124 (2015).
- [227] S.-S. Gong, W. Zhu, L. Balents, and D. N. Sheng, Global phase diagram of competing ordered and quantum spin-liquid phases on the kagome lattice, *Physical Review B* **91**, 075112 (2015).
- [228] A. Wietek, A. Sterdyniak, and A. M. Läuchli, Nature of chiral spin liquids on the kagome lattice, *Physical Review B* **92**, 125122 (2015).
- [229] F. Oliviero, J. A. Sobral da Silva, E. Andrade, and R. G. Pereira, Noncoplanar magnetic orders and gapless chiral spin liquid on the kagome lattice with staggered scalar spin chirality, *SciPost Physics* **13**, 050 (2022).
- [230] T. Liu, W. Li, A. Weichselbaum, J. von Delft, and G. Su, Simplex valence-bond crystal in the spin-1 kagome Heisenberg antiferromagnet, *Physical Review B* **91**, 060403 (2015).
- [231] H. J. Changlani and A. M. Läuchli, Trimerized ground state of the spin-1 Heisenberg antiferromagnet on the kagome lattice, *Physical Review B* **91**, 100407 (2015).
- [232] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Rigorous results on valence-bond ground states in antiferromagnets, *Physical Review Letters* **59**, 799–802 (1987).
- [233] R. K. Kaul, Spin Nematics, Valence-Bond Solids, and Spin Liquids in $SO(N)$ Quantum Spin Models on the Triangular Lattice, *Physical Review Letters* **115**, 157202 (2015).
- [234] S. Pujari, T. C. Lang, G. Murthy, and R. K. Kaul, Interaction-Induced Dirac Fermions from Quadratic Band Touching in Bilayer Graphene, *Physical Review Letters* **117**, 086404 (2016).
- [235] Y. Nomura and M. Imada, Dirac-Type Nodal Spin Liquid Revealed by Refined Quantum Many-Body Solver Using Neural-Network Wave Function, Correlation Ratio, and Level Spectroscopy, *Physical Review X* **11**, 031034 (2021).
- [236] A. Keleş and E. Zhao, Rise and fall of plaquette order in the Shastry-Sutherland magnet revealed by pseudofermion functional renormalization group, *Physical Review B* **105**, L041115 (2022).
- [237] J. Sonnenschein, A. Maity, C. Liu, R. Thomale, F. Ferrari, and Y. Iqbal, Candidate quantum spin liquids on the maple-leaf lattice, *Physical Review B* **110**, 014414 (2024).
- [238] D. J. J. Farnell, R. Darradi, R. Schmidt, and J. Richter, Spin-half Heisenberg antiferromagnet on two archimedean lattices: From the bounce lattice to the maple-leaf lattice and beyond, *Physical Review B* **84**, 104406 (2011).
- [239] J. Beck, J. Bodky, J. Motruk, T. Müller, R. Thomale, and P. Ghosh, Phase diagram of the $J - J_d$ Heisenberg model on the maple leaf lattice: Neural networks and density matrix renormalization group, *Physical Review B* **109**, 184422 (2024).

[240] P. Schmoll, J. Naumann, E. L. Weerda, J. Eisert, and Y. Iqbal. Bathing in a sea of candidate quantum spin liquids: From the gapless ruby to the gapped maple-leaf lattice. arXiv: 2407.07145.

[241] A. Smerald and N. Shannon, Theory of NMR $1/T_1$ relaxation in a quantum spin nematic in an applied magnetic field, *Physical Review B* **93**, 184419 (2016).

[242] P. Sindzingre, N. Shannon, and T. Momoi, Nematic order in square lattice frustrated ferromagnets, *Journal of Magnetism and Magnetic Materials* **310**, 1340–1342 (2007).

[243] P. Sindzingre, L. Seabra, N. Shannon, and T. Momoi, Phase diagram of the spin-1/2 J1-J2-J3 Heisenberg model on the square lattice with ferromagnetic J1, *Journal of Physics: Conference Series* **145**, 012048 (2009).

[244] P. Sindzingre, N. Shannon, and T. Momoi, Phase diagram of the spin-1/2 J1-J2-J3 Heisenberg model on the square lattice, *Journal of Physics: Conference Series* **200**, 022058 (2010).

[245] R. Shindou, S. Yunoki, and T. Momoi, Projective studies of spin nematics in a quantum frustrated ferromagnet, *Physical Review B* **84**, 134414 (2011).

[246] Y. Iqbal, P. Ghosh, R. Narayanan, B. Kumar, J. Reuther, and R. Thomale, Intertwined nematic orders in a frustrated ferromagnet, *Physical Review B* **94**, 224403 (2016).

[247] A. Wietek and A. M. Läuchli, Valence bond solid and possible deconfined quantum criticality in an extended kagome lattice Heisenberg antiferromagnet, *Physical Review B* **102**, 020411 (2020).

[248] T. Momoi, P. Sindzingre, and N. Shannon, Octupolar Order in the Multiple Spin Exchange Model on a Triangular Lattice, *Physical Review Letters* **97**, 257204 (2006).

[249] P. Ghosh, T. Müller, F. P. Toldin, J. Richter, R. Narayanan, R. Thomale, J. Reuther, and Y. Iqbal, Quantum paramagnetism and helimagnetic orders in the Heisenberg model on the body centered cubic lattice, *Physical Review B* **100**, 014420 (2019).

[250] D. Kiese et al., Pinch-points to half-moons and up in the stars: The kagome skymap, *Physical Review Research* **5**, L012025 (2023).

[251] K. Hukushima and K. Nemoto, Exchange Monte Carlo Method and Application to Spin Glass Simulations, *Journal of the Physical Society of Japan* **65**, 1604–1608 (1996).

[252] R. H. Swendsen and J.-S. Wang, Replica Monte Carlo Simulation of Spin-Glasses, *Physical Review Letters* **57**, 2607–2609 (1986).

[253] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov, Electric Field Effect in Atomically Thin Carbon Films, *Science* **306**, 666–669 (2004).

[254] S. Manzeli, D. Ovchinnikov, D. Pasquier, O. V. Yazyev, and A. Kis, 2D transition metal dichalcogenides, *Nature Reviews Materials* **2**, 17033 (2017).

[255] C. R. Dean et al., Boron nitride substrates for high-quality graphene electronics, *Nature Nanotechnology* **5**, 722–726 (2010).

[256] Y. Cao et al., Correlated insulator behaviour at half-filling in magic-angle graphene superlattices, *Nature* **556**, 80–84 (2018).

[257] Y. Cao, V. Fatemi, S. Fang, K. Watanabe, T. Taniguchi, E. Kaxiras, and P. Jarillo-Herrero, Unconventional superconductivity in magic-angle graphene superlattices, *Nature* **556**, 43–50 (2018).

- [258] X. Lu et al., Superconductors, orbital magnets and correlated states in magic-angle bilayer graphene, *Nature* **574**, 653–657 (2019).
- [259] M. Yankowitz, S. Chen, H. Polshyn, Y. Zhang, K. Watanabe, T. Taniguchi, D. Graf, A. F. Young, and C. R. Dean, Tuning superconductivity in twisted bilayer graphene, *Science* **363**, 1059–1064 (2019).
- [260] G. W. Burg, J. Zhu, T. Taniguchi, K. Watanabe, A. H. MacDonald, and E. Tutuc, Correlated Insulating States in Twisted Double Bilayer Graphene, *Physical Review Letters* **123**, 197702 (2019).
- [261] C. Shen et al., Correlated states in twisted double bilayer graphene, *Nature Physics* **16**, 520–525 (2020).
- [262] X. Liu et al., Tunable spin-polarized correlated states in twisted double bilayer graphene, *Nature* **583**, 221–225 (2020).
- [263] Y. Cao, D. Rodan-Legrain, O. Rubies-Bigorda, J. M. Park, K. Watanabe, T. Taniguchi, and P. Jarillo-Herrero, Tunable correlated states and spin-polarized phases in twisted bilayer–bilayer graphene, *Nature* **583**, 215–220 (2020).
- [264] G. Chen et al., Evidence of a gate-tunable Mott insulator in a trilayer graphene moiré superlattice, *Nature Physics* **15**, 237–241 (2019).
- [265] G. Chen et al., Signatures of tunable superconductivity in a trilayer graphene moiré superlattice, *Nature* **572**, 215–219 (2019).
- [266] G. Chen et al., Tunable correlated Chern insulator and ferromagnetism in a moiré superlattice, *Nature* **579**, 56–61 (2020).
- [267] H. Zhou et al., Half- and quarter-metals in rhombohedral trilayer graphene, *Nature* **598**, 429–433 (2021).
- [268] J. Yang et al., Spectroscopy signatures of electron correlations in a trilayer graphene/hBN moiré superlattice, *Science* **375**, 1295–1299 (2022).
- [269] C. Jin et al., Observation of moiré excitons in WSe₂/WS₂ heterostructure superlattices, *Nature* **567**, 76–80 (2019).
- [270] E. C. Regan et al., Mott and generalized Wigner crystal states in WSe₂/WS₂ moiré superlattices, *Nature* **579**, 359–363 (2020).
- [271] L. Wang et al., Correlated electronic phases in twisted bilayer transition metal dichalcogenides, *Nature Materials* **19**, 861–866 (2020).
- [272] C. Jin et al., Stripe phases in WSe₂/WS₂ moiré superlattices, *Nature Materials* **20**, 940–944 (2021).
- [273] M. A. Ruderman and C. Kittel, Indirect Exchange Coupling of Nuclear Magnetic Moments by Conduction Electrons, *Physical Review* **96**, 99–102 (1954).
- [274] T. Kasuya, A Theory of Metallic Ferro- and Antiferromagnetism on Zener’s Model, *Progress of Theoretical Physics* **16**, 45–57 (1956).
- [275] Y. Lian, A. Rosch, and M. O. Goerbig, SU(4) Skyrmiions in the $\theta = \pm 1$ Quantum Hall State of Graphene, *Physical Review Letters* **117**, 056806 (2016).
- [276] Y. H. Kwan, G. Wagner, T. Soejima, M. P. Zaletel, S. H. Simon, S. A. Parameswaran, and N. Bultinck, Kekulé Spiral Order at All Nonzero Integer Fillings in Twisted Bilayer Graphene, *Physical Review X* **11**, 041063 (2021).

- [277] W.-Y. Liao et al., Intervalley Coherent Order in Rhombohedral Tetralayer Graphene on MoS₂, *Physical Review Letters* **135**, 046202 (2025).
- [278] W. Feng, J.-P. Hanke, X. Zhou, G.-Y. Guo, S. Blügel, Y. Mokrousov, and Y. Yao, Topological magneto-optical effects and their quantization in noncoplanar antiferromagnets, *Nature Communications* **11**, 118 (2020).
- [279] M. J. Calderón, A. Camjayi, and E. Bascones, Mott correlations in ABC graphene trilayer aligned with hBN, *Physical Review B* **106**, L081123 (2022).
- [280] Y.-H. Zhang, D. Mao, Y. Cao, P. Jarillo-Herrero, and T. Senthil, Nearly flat Chern bands in moiré superlattices, *Physical Review B* **99**, 075127 (2019).
- [281] N. D. Mermin and H. Wagner, Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models, *Physical Review Letters* **17**, 1133–1136 (1966).
- [282] V. L. Berezinskii, Destruction of Long-range Order in One-dimensional and Two-dimensional Systems having a Continuous Symmetry Group I. Classical Systems, *Sov. Phys. JETP* **32**, 493–500 (1971).
- [283] J. M. Kosterlitz and D. J. Thouless, Ordering, metastability and phase transitions in two-dimensional systems, *Journal of Physics C: Solid State Physics* **6**, 1181–1203 (1973).
- [284] L. Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199–208 (2010).
- [285] A. S. Patri and T. Senthil, Strong correlations in ABC-stacked trilayer graphene: Moiré is important, *Physical Review B* **107**, 165122 (2023).
- [286] L. V. Keldysh, Diagram Technique for Nonequilibrium Processes, *Sov. Phys. JETP* **20** (1965).
- [287] S. G. Jakobs, M. Pletyukhov, and H. Schoeller, Nonequilibrium functional renormalization group with frequency-dependent vertex function: A study of the single-impurity Anderson model, *Physical Review B* **81**, 195109 (2010).
- [288] A. Ge, N. Ritz, E. Walter, S. Aguirre, J. Von Delft, and F. B. Kugler, Real-frequency quantum field theory applied to the single-impurity Anderson model, *Physical Review B* **109**, 115128 (2024).

Acknowledgements

I would like to begin by thanking my supervisor, Simon Trebst, who has supported me ever since my master's thesis. He gave me considerable freedom to explore ideas I found interesting, while always nudging me in the right direction and helping me focus on what truly matters. I am also grateful to Matteo Rizzi for kindly agreeing to serve as the second examiner of this thesis, and to Markus Braden for acting as chair during my defense.

I owe much to the excellent scientists I had the pleasure of working with, who made the projects of this thesis possible in the first place. I am grateful to Michael Scherer, who co-supervised the projects on SU(4) models and moiré materials. His expertise in both the material and FRG literature was invaluable. Likewise, I sincerely thank Yasir Iqbal, who initiated all the projects on maple-leaf and pyrochlore materials. Yasir is a constant pool of ideas, with a truly broad knowledge of both frustrated magnetism and numerical techniques. I am particularly thankful to him for inviting me to work in his group at IIT Madras for three months, which was a truly great experience. I also thank Daniel Lozano-Gómez for the very nice collaboration on the pyrochlore project.

I gratefully acknowledge funding from the German Research Foundation through the CRC 1238. The numerical simulations of this work were performed on the JUWELS cluster at Forschungszentrum Jülich and the Noctua2 cluster at the Paderborn Center for Parallel Computing (PC2).

My heartfelt thanks also go to all the members of the institute in Cologne who made the PhD a genuinely enjoyable experience. The traditional 12:00 Mensa lunches and subsequent coffee breaks were always a highlight of my day. I thank my former office mate Dominik, who first introduced me to the world of pf-FRG, and Sagar, whose vast knowledge of IT and programming solved countless problems. I very much enjoyed your company. I also thank Yoshito, Aprem, and Sagar for proofreading this thesis, and Christoph for providing the L^AT_EX template on which this thesis is based.

Finally, I want to thank my wife Lisa, who kept me sane during the stressful times of the PhD and, more importantly, makes me a much happier person every day I am with her.

Magnets have fascinated people for centuries and profoundly shaped society—from compass needles that guided early explorers to permanent magnets in electric motors that power modern technology. In conventional magnetic materials, known as ferromagnets, the spins of the electrons align collectively, producing a stable magnetic field. Depending on the atomic ingredients and crystal structure, however, the interactions between electron spins can favor very different, and sometimes competing, patterns of alignment. When this competition is particularly strong, theory predicts that the quantum nature of electrons can give rise to more unusual forms of magnetism, in which spins fail to order altogether, even at absolute zero. Instead, they remain in a fluctuating, highly entangled state known as a *quantum spin liquid*. Despite decades of theoretical and experimental effort—driven by the exotic nature of these phases, their potential relevance for quantum computing, and their possible connection to high-temperature superconductivity—an unambiguous realization of a quantum spin liquid in real materials remains elusive.

In this thesis, we numerically investigate three classes of magnetic materials that are considered promising candidates for hosting quantum spin liquid behavior: moiré materials, maple-leaf magnets, and pyrochlore rare-earth oxides. Using effective theoretical models in the form of frustrated spin Hamiltonians, we study how quantum fluctuations modify their ground-state phase diagrams—by stabilizing new ordered phases, shifting phase boundaries, or, in the most favorable scenario, suppressing classical magnetic order altogether in favor of quantum spin liquid behavior. To systematically assess the impact of quantum effects, we employ and extend the pseudo-fermion functional renormalization group—a particularly versatile many-body method for describing quantum fluctuations in frustrated magnets—and complement it with classical and semiclassical approaches.