2025

Static and dynamic properties of quantum magnets

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

STATIC AND DYNAMIC PROPERTIES OF QUANTUM MAGNETS

by

GABE SCHUMM

B.A., University of California, Berkeley, 2019

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2025

Approved by

First Reader

Anders W. Sandvik, Ph.D.
Professor of Physics

Second Reader

Pankaj Mehta, Ph.D.
Professor of Physics

Acknowledgments

This research was supported by the Simons Foundation under Grant No. 511064. Computational resources were provided by the Shared Computing Cluster managed by Boston University's Research Computing Services and by the Flatiron Institute Scientific Computing Center.

I first and foremost must thank my research advisor, Anders Sandvik, for his mentorship and guidance over the past five years. Anders has been able to strike an incredible balance as an advisor, pushing me to learn, step out of my comfort zone, and produce high-quality research, while never introducing unnecessary stress or strain into the Ph.D. experience. I feel fortunate for the opportunities made possible by his support; conferences in Portugal and Colorado, a Summer school in Germany, and a pre-doctoral fellowship at the Flatiron Institute in New York City. These experiences have had a lasting impact on me as a young physicist and as a young adult. I do not take for granted the comfortability Anders' advising style has provided. While one of the leading experts in our field, his humbleness and down-to-earth personality made talking about many-body physics just as easy as talking about gardening or what to order for lunch.

I would also like to thank my research mentor during my pre-doc fellowship at the Flatiron Institute, Shiwei Zhang. He is truly one of the kindest and most brilliant physicists I have had the pleasure of working with. From day one, he made it clear that his goal was to make the pre-doc experience as beneficial as possible for me as a young scientist, and he achieved that, and then some.

I've been able to learn, teach for, and discuss physics with some amazing professors while at Boston University; my thesis committee members Pankaj Mehta, Ami Katz, Anatoli Polkovnikov, and Shyam Erramilli, as well as Claudio Chamon, Wangzheng Hu, Liam Fitzpatrick, and Anushya Chandran. I am also grateful for the relationships

I've made with my cohort members, officemates, and the CMT department as a whole. When I begin working in Anders' research group, I came in with very little previous experience in our field, and the encouragement and mentorship of Bowen Zhao, Kai-Hsin Wu and Sibin Yang were invaluable. I aimed to pay that forward to the members that have joined since, Ryan Flynn, Zhi Lin, and Nikhil Nair, and am thankful the opportunity to both guide and collaborate with them.

I, of course, must express my deep appreciation for my family and friends. My parents, Dan and Ellen, have paved the path for me to be where I am today; they have investigated so much time and effort into my success and it has more than payed off. And I cannot imagine the past five years without the love and support of my partner, Molly Hydorn. Moving across the country during such an unknown and unstable time, navigating through the highs and lows of a Ph.D. program, and maturing into the adult I am today would not have been the same without her. I will be forever grateful for the imprint she has made on this experience.

Gabe Schum

Ph.D. Candidate

Department of Physics

STATIC AND DYNAMIC PROPERTIES OF QUANTUM MAGNETS

GABE SCHUMM

Boston University, Graduate School of Arts and Sciences, 2025

Major Professor: Anders W. Sandvik, PhD, Professor of Physics

ABSTRACT

Quantum spin models provide the basis for understanding the complex phenomena associated with quantum magnetism. By tuning parameters describing the interactions between these spins, the onset of quantum fluctuations can drive a spin system into an exotic and complex "highly-correlated" state of matter, often displaying unusual, non-classical properties. These quantum phases typically involve the delicate interplay of an exponentially large number of degrees of freedom, posing substantial difficulties to studying using purely analytical tools. Numerical simulation of spin systems thus provides a key piece of this puzzle, allowing for the systematic study of quantum magnetism and the discovery of new quantum phenomena.

To study quantum magnetism, physicists rely on lattice-models, regular grids hosting particles that interact with one another via quantum exchanges. These simplified models are used to study the low energy behavior of real quantum magnets and are typically more amenable to both analytic and numerical calculations, the latter of which being the main approach discussed in my thesis. Using quantum Monte Carlo simulations, I have studied various spin models with three main goals in mind: to explain observations made by experimental physicists, to predict how novel quantum systems may behave, and to provide benchmark calculations for the wide array of

quantum devices and emulators that are entering the fray. In this thesis, I discuss the development and implementation of various numerical methods which I have used to contribute to all three of these branches of computational physics research. In particular, my thesis focuses on the calculation of static and dynamic properties of quantum magnetic materials, the different numerical tools used to do so, and the insights we gain by considering these two classes of observables in tandem.

This thesis focuses of two models in particular: the Fully Frustrated Transverse Field Ising Model on the square-lattice and the two-dimensional Hubbard Model. I investigated both the ground-state properties and out-of-equilibrium dynamics of the former using a novel approach involving two order parameters, which capture different symmetries of the ordered phase. My work on the latter has focused on resolving the single-particle spectral function and the density of states by utilizing recent advancements in the stochastic analytic continuation method. I discuss my contributions to the development of the stochastic analytic continuation method and the application of this method to the Hubbard Model.

Contents

List of Figures

List of Abbreviations

AF	....	Antiferromagnetic
AFQMC	....	Auxiliary-Field Quantum Monte Carlo
BA	....	Bethe Ansatz
BZ	....	Brillouin Zone
DOS	....	Density of States
FFTFIM	....	Fully Frustrated Transverse Field Ising Model
KT	....	Kosterlitz-Thouless
KZ	....	Kibble-Zurek
MEM	....	Maximum Entropy Method
QA	....	Quantum Annealing
QCP	....	Quantum Critical Point
QLRO	....	Quasi Long-Range Order
QMC	....	Quantum Monte Carlo
QPT	....	Quantum Phase Transition
RBA	....	Rigid Band Approximation
SA	....	Simulated Annealing
SAC	....	Stochastic Analytic Continuation
SQA	....	Simulated Quantum Annealing
SSE	....	Stochastic Series Expansion

Chapter 1

Introduction

This dissertations explores the numerical study of quantum magnets with a focus on the methods used to measure both their static and dynamic properties. In this context, the quantifier "static" refers to a system in its equilibrium state, at either finite or zero temperature. Such a system is typically characterized by measurable quantities such as internal energy, spatial correlations (e.g., the static structure factor), and susceptibility to an external field, among others. There have been decades of development of the numerical tools used to calculate this class of properties, and, in particular, the collection of QMC methods available for measuring them with high accuracy is vast [115].

The particular approach taken in this thesis is the development of numerical techniques to better establish the connection between lattice models and the underlying continuum field theories that describe their criticality, particularly through scaling dimensions. By analyzing the critical, power-law decaying correlations in a lattice model, this relation can be made in a precise and methodical manner, providing a deeper understand of the system's criticality [118].

Calculating dynamical properties of quantum magnets is typically a more challenging task [117]. By "dynamics," we refer to the response of a strongly interacting quantum many-body system to some time-dependent external probe—what modes are excited, or equivalently, what higher energy states become occupied upon interaction with an operator coupled to the system. A helpful, yet highly simplified, classical analog is the response of a tuning fork to a mallet's strike. The interaction with this stimulus excites the natural modes of the tuning fork, which we detect as a ringing at the corresponding natural frequency. If one were to capture the sound made by the fork and Fourier transform the result, this would correspond to measuring a type of "spectral function" for this system. Along with a dominant peak at the natural frequency of the tuning fork, one would also likely observe peaks at an array of frequencies corresponding to harmonics, as well as other modes excited by the mallet. This feedback provides important information about the system—in this example, about the geometry of the tuning fork, the material it is made of, and any impurities or structural imperfections present.

In a quantum many body system, things are much more complicated, but the underlying principle still applies—one can gain a lot of insight by "striking" a quantum magnet with a "mallet." If this figurative mallet is a neutron, then spin-excitations (magnons) are probed; the dynamic response function in this case is typically referred to as the dynamic structure factor, $S(q, \omega)$. If instead one uses photons, then electrons are excited subsequently ejected from the lattice; the response function is instead referred to as the single-particle spectral function, $A(k, \omega)$. These response functions have immense importance in the connection between experiment and theory, as many of the spectroscopy tools used to study quantum magnets access them directly. From the theoretical side, a response function can be calculated using the exact time evolution of the quantum system, as the Fourier transform of the time-dependent correlation function, but this is notoriously computationally expensive. To utilize QMC methods, we have to take another route—*imaginary* time evolution. While distinctly different, many of the dynamical properties contained in the real time correlations of quantum system are also contained in the imaginary time correlations [62].

The key challenge in this approach is the "analytic continuation" back to the real-frequency axis, an infamously challenging inverse problem. In this thesis, we explore and further develop the Stochastic Analytic Continuation (SAC) method, a powerful tool that allows for the resolution of spectral functions with unprecedented fidelity [124]. We discuss our recent developments and application to one of the quintessential models of correlated electron physics: the 2D Hubbard Model.

1.1 The Fully Frustrated Transverse Field Ising Model

In Chapter 1, we focus on a large-scale quantum Monte Carlo (QMC) study of a particular quantum magnet, the Fully Frustrated (Villain) Ising Model on the square lattice. We studied this model at both finite temperature (T) and at its zero temperature quantum critical point (QCP) driven by the tuning of a transverse magnetic field Γ. Here, we explore the static properties of this system using a novel approach: considering both a primary and secondary order parameter. These two order parameters lead to the same finite temperature phase diagram, which we've mapped out in $T-\Gamma$ plane, but capture different symmetries of the ground state ($\mathbb{Z}_8$ and $\mathbb{Z}_4$, for the primary and secondary, respectively), and fascinatingly, scale differently at the two critical regimes of the model. In the finite temperature critical phase, which is associated with the emergence of bound topological defects, the correlations of the primary order parameter decays as a power law, with the conventional exponent predicted by a mapping to the Classical 2D 8-state Clock Model. The secondary order parameter, on the other hand, decays four times as fast, highlighting its secondary nature.

The quantum critical point, which we located using finite-size scaling analysis ($\Gamma_c/J = 1.5768(9)$; uncertainty estimated as described in Sec. 2.6.1), is in the 3D XY universality class, which implies that the scaling of the primary order parameter is instead dictated by the critical exponent, η_{3DXY}. However, the scaling of secondary order parameter can no longer be simply related to that of the primary order parameter. This universality class has a field theoretic description, the 3D O(2) conformal field theory, which allows us to draw direct connection between the lattice operators associated with these order parameters and operators defined within the field theory. By evoking this connection, we relate the critical exponents of both the primary and secondary operators to scaling dimensions of the relevant field operators, Δ_ϕ and Δ_t. We draw physical connection between these two classes of operators to motivate the observed scaling relations.

As an extension of this work, we investigate the out-of-equilibrium dynamics of the Fully Frustrated Transverse Field Ising Model (FFTFIM), exploring quantum annealing through the QCP, to the classical limit ($\Gamma = 0$). Recent studies by the superconducting qubit quantum annealing company D-Wave Quantum have examined this geometrically frustrated quantum magnet as a benchmark for their state-of-the-art devices. Our study of the FFTFIM's static properties were instrumental in shaping the design and analysis of their experiments, and our follow-up study provides a valuable point of comparison for D-Wave's results.

1.2 Stochastic Analytic Continuation—Principles and Implementation

In Chapter 2, we introduce the numerical analytic continuation problem and layout the framework behind the SAC method. After presenting the problem, we briefly discuss the traditional method used to solve it, the Maximum Entropy Method (MEM). We then present our implementation of the SAC procedure, beginning with the standard "unconstrained" sampling scheme.

After describing the stochastic sampling procedure that largely differentiates SAC from MEM, we discuss the class of "constrained" SAC parameterizations that represent the true advantage that SAC offers. By imposing restrictions on the sampling process, one can resolve spectra with sharp features that are fundamentally inaccessible using MEM and other standard analytic continuation methods. We discuss the implementation of existing constrained SAC parameterizations, as well as new constrained parameterizations developed as part of this thesis. We also refer to our GitHub repository, where a user-friendly Julia implementation of the SAC program was written as part of this thesis is available.

1.3 Single-Particle Dispersion and Density of States of the Half-Filled 2D Hubbard Model

In Chapter 3, we apply the SAC method to to calculate the single-particle dispersion relation and to resolve the density of states (DOS) of one of the quintessential models of correlated electron physics—the two-dimensional (2D) Hubbard Model at half-filling. While its electronic properties are typically the focus of numerical studies, at any non-zero interaction strength, $U/t > 0$, the half-filled model is a Mott insulator, displaying antiferromagnet order. Spin correlations thus play a large role in the physics of the Hubbard Model, emphasized by the inherent (π, π) nesting of the non-interacting Fermi surface [105].

The 2D Hubbard Model also presents the perfect grounds to explore the impact that the developments in SAC will have on our understanding of the fundamental models of condensed matter physics. Extracting the single-particle dispersion relation and density of states of the half-filled model remain long-standing challenges in the field, due in large part to the the difficulties surrounding the numerical analytic continuation problem. The advancements of SAC are uniquely capable of addressing this problem, in light of the sharp features expected to be present in both the single-particle spectral function $A_{\mathbf{k}}(\omega)$ and the DOS.

Using numerically exact auxiliary-field quantum Monte Carlo (AFQMC) simulations, at intermediate interactions of $U/t = 4, 6$, we find quadratic dispersion around the gap minimum at wave-vectors $\mathbf{k} = (\pm\pi/2, \pm\pi/2)$ (the Σ points). In contrast to the previously held understanding of this model, we find saddle points at $\mathbf{k} = (\pm\pi, 0), (0, \pm\pi)$ (the X points) where the dispersion is quartic, leading to a sharp DOS maximum above the almost flat ledge arising from the states close to Σ. The fraction of quasi-particle states within the ledge is $n_{\text{ledge}} \approx 0.15$, which we determine using a novel constrained SAC parameterization. Upon doping, within the rigid-band approximation, these results support Fermi pockets around the Σ points, with states around the X points becoming filled only at doping fractions $x \geq n_{\text{ledge}}$. The high density of states and the associated onset of (π, π) scattering may be an important clue for a finite minimum doping level for superconductivity in the cuprates, emphasizing the impactful insights one gains by investigating the dynamic properties of highly-correlated, quantum many-body systems using numerical techniques.

1.4 Cross Validation in Stochastic Analytic Continuation

When performing numerical analytical continuation, it is common to not know what exact sharp features, if any, are present in the spectral function *a priori*. Due to the ill-posed nature of the analytic continuation problem, multiple spectral representations may be acceptable, and the goodness-of-fit χ^2 cannot be used alone to determine which representation is best. In chapter 4, we borrow from the machine learning and statistics literature and explore the implementation of a cross validation technique to provide an unbiased method to identify the most likely spectrum amongst a set obtained with different spectral parameterizations and imposed constraints.

We begin by discussing how cross validation can be used to detect over-fitting—an issue that can significantly affect the results of numerical analytic continuation due to the noise inherent in QMC-generated data. We then present a range of examples demonstrating how cross validation can be used as a tool for "model selection." By comparing the *validation* χ^2 values associated with different SAC parameterizations, we are able to determine, with impressive accuracy, which spectrum most accurately represents the true underlying spectral function. We apply this framework to three distinct classes of problems.

The first purely serves as a proof of concept: given an artificial spectrum and corresponding synthetic data aimed to mimic that generated by QMC simulations, can cross validation correctly identify the most accurate SAC spectrum?

The second case involves real QMC-generated data for a system whose dynamic structure factor has a known analytic form, which can be reproduced by one of the constrained SAC parameterizations (a power law edge).

The third, and most impactful, case also involves QMC-generated data, but for a system where the exact form of the spectral function is unknown, and furthermore, is expected to evolve as model parameters are tuned.

1.5 Conclusions

The work presented in this thesis represents several advancements in the field of computational condensed matter theory. From a methods standpoint, we explore novel approaches to studying the static properties of quantum magnets, emphasizing the connection between lattice models, effective field theories, and analogous classical systems. Our results have provided important quantum annealing benchmarks, leading to promising advancements in the field of quantum computing. The latter portion of this thesis is dedicated to the development of the SAC method—improving, extending, and tailoring the program for applicability in an even larger class of numerical analytic continuation problems. These advancements have directly led to new discoveries, providing insights on the nature of the superconducting transition in the cuprates.

This thesis provides several demonstrations of the importance of numerical method development in the pursuit of scientific understanding. While “designer Hamiltonians” and simplified models have their limitations, they are crucial to advancing the field of computational physics and, in many cases, serve as powerful tools for understanding the fundamental mechanisms that govern complex quantum systems.

Chapter 2

The Fully Frustrated Transverse Field Ising Model

2.1 Introduction

The concept of an order parameter is key to quantitative descriptions of phase transitions. In some systems it is natural to define more than one order parameter, either in some trivial way or using emergent degrees of freedom originating from some mapping to an effective low-energy model. The relationships between different order parameters may be non-trivial, e.g., unexplained behavior of a "parasitic" ferromagnetic order parameter in a system with primarily antiferromagnetic order was reported [47, 48, 133]. Here we consider a quantum spin model that very clearly illustrates two different order parameters that not only exhibit different scaling behaviors but the relationships between the critical exponents of the order parameters are also different at temperature $T = 0$ and $T > 0$.

We study the two-dimensional (2D) square-lattice Fully Frustrated Transverse Field Ising Model (FFTFIM), also referred to as the Villain model, with Hamiltonian

$$H = \sum_{\langle ij \rangle} J_{ij} \sigma_i^z \sigma_j^z - \Gamma \sum_j \sigma_j^x, \tag{2.1}$$

where σ_i^x and σ_i^z are Pauli operators. The couplings J_{ij} are equal in magnitude but the number of antiferromagnetic (AF) couplings around any elementary plaquette is odd [137], here with $J_{ij} = +J$ (AF) on every second column and $J_{ij} = -J$ on all other bonds as depicted in Fig. 2·1. The classical model at $\Gamma = 0$ hosts a large ground state degeneracy that is lifted by the transverse field via an "order-by-disorder" mechanism [139, 90, 50].

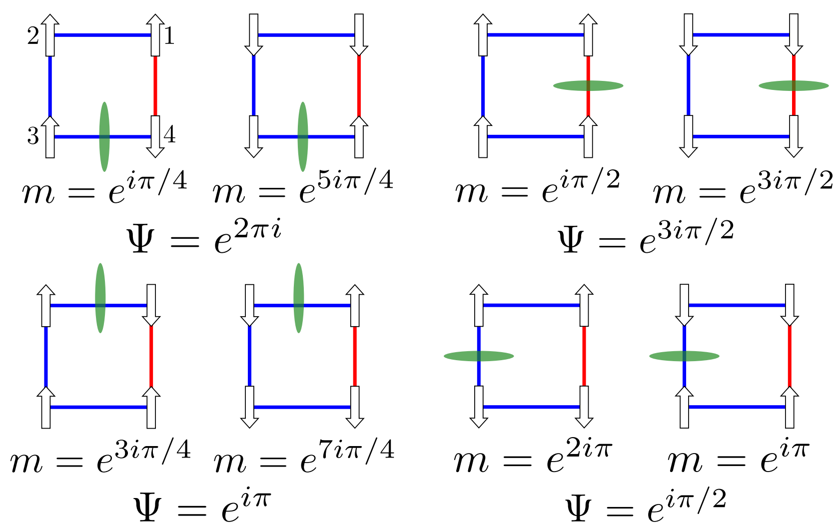

Figure 2·1: The degenerate ground states represented by elementary plaquettes. Blue and red bonds show $J_{ij} = -J$ and $J_{ij} = +J$, respectively. The direction of the sublattice magnetizations m_s, $s = 1, \dots, 4$, are indicated by the arrows, and dimers (green ovals) are defined on the frustrated bonds. The values of the magnetization and dimer order parameters (see Sec. 2.2 for definitions) shown correspond to the ground state sublattice magnetizations of the stacked, classical ($\Gamma = 0$) model [14], $\sin(\pi/8)$ for the sites sharing a frustrated bond and $\cos(\pi/8)$ for the two others [see Eqs. (3) and (5)].

The first studies of the FFTFIM considered the stacked version of the classical model using a using Landau-Ginsburg-Wilson approach [14] as well as Monte Carlo (MC) simulations [61]. This study predicted an eight-fold degenerate ground state, which was more precisely characterized by Ref. [61] as a $\mathbb{Z}_4$ symmetry breaking phase, corresponding to 90° rotations of the lattice, paired with a global spin-flip symmetry ($\mathbb{Z}_2$) .

The model was later treated using quantum MC (QMC) simulations [143], where spin and dimer order was found at $T = 0$ for $\Gamma < \Gamma_c \sim 1.578$. In this phase the frustrated bonds (mapped to dimers as in Fig. 2·1) align along alternating columns or rows; the $\mathbb{Z}_4$ symmetry breaking phase identified in the stacked model. The order parameters used to measure this phase in previous studies detected $\mathbb{Z}_4$ symmetry and $\mathbb{Z}_2$ spin-reflection symmetry separately, but a proper $\mathbb{Z}_8 = \mathbb{Z}_4 \times \mathbb{Z}_2$ order parameter was never considered. The reconciliation of $\mathbb{Z}_8$ versus $\mathbb{Z}_4$ breaking was touched on briefly by Coletta et al. [30], but the the relationship between the respective order parameters was not explored.

Here we define a proper $\mathbb{Z}_8$ symmetric spin order parameter and demonstrate that the $\mathbb{Z}_4$ dimer order parameter should be considered as *secondary.* While both order parameters lead to the same phase diagram, see Fig. 2·8, the $\mathbb{Z}_4$ order parameter exhibits faster decaying critical correlations. Both order parameters, when correctly defined, exhibit emergent U(1) symmetry in the critical phase as well as at the quantum phase transition, stemming from the irrelevance of the discrete symmetry-breaking terms at criticality [65]. However, the distinction between primary and secondary order is made quantitative by considering the critical scaling in these two regimes. This is only made possible by connecting these order parameters to operators in the relevant field theories, at $T = 0$ and $T > 0$. Finite-size scaling of QMC (stochastic series expansion [114]) results of the full FFTFIM Hamiltonian, Eq. (2.1), support the predicted scaling, emphasizing the utility of this novel approach to studying quantum magnetism.

Secondary order parameters have previously been used to describe higher harmonic contributions to spatial modulation in density wave systems, e.g., liquid crystals [148, 2, 46, 96, 20]. A secondary order parameter can clearly be defined also in the FFTFIM, but the different scaling forms of the spin and dimer order parameter in the FFTFIM have not been addressed. This work provides a framework for secondary order not just in the FFTFIM, but in the entire class of frustrated Ising models with effective dimer degrees of freedom, e.g., the antiferromagnet on the triangular lattice [15, 16, 60, 91, 127].

2.2 Order Parameters

To construct a proper primary order parameter, we follow standard procedures [14, 60, 30], using an effective Hamiltonian for the amplitude m and phase θ of critical modes:

$$H_{\text{eff}} = \sum_{\vec{q}}(r+q^2)|m|^2+u_4|m|^4+u_6|m|^6+(u_8+v_8/32)|m|^8-(v_8/32)|m|^8\cos(8\theta). \quad (2.2)$$

The eight-state clock anisotropy implies an eight-fold degenerate ground state, characterized by sublattice magnetizations (m_1, m_2, m_3, m_4); see Fig. 2·1. Each state corresponds to one frustrated bond in a plaquette, where the magnitude of the sublattice magnetizations of the sites sharing this bond are smaller than the other two, $\sin(\pi/8)$ and $\cos(\pi/8)$ respectively [14]. An overall spin-flip transformation gives a total of eight degenerate states.

Based on the low-energy behavior of the stacked model, as well as the semiclassical analysis of Ref. [30], we define the primary order parameter as the complex number

$$m = m_x + im_y = \frac{1}{2}\left(m_1 e^{i\frac{\pi}{8}} + m_2 e^{i\frac{3\pi}{8}} + m_3 e^{i\frac{5\pi}{8}} + m_4 e^{i\frac{7\pi}{8}}\right) \quad (2.3)$$

where

$$m_s = \frac{4}{N}\sum_{j\in s}\sigma_j^z. \quad (2.4)$$

The eight ground state configurations of the stacked model correspond to $m = e^{in\pi/4}$, $n = 1, 2, \ldots, 8$.

The problem can also be mapped onto that of dimer coverings, where a dimer is assigned across each frustrated bond [91], and we define a secondary order parameter

Ψ with this mapping in mind. This order parameter is also complex number, defined in terms of the dimer density modulation on the dual lattice [143]:

$$\Psi = \Psi_x + i\Psi_y = 2\tilde{d}_x(0,\pi) + 2i\tilde{d}_y(\pi,0), \tag{2.5}$$

where

$$\tilde{d}_\alpha(\boldsymbol{q}) = \frac{1}{N}\sum_i e^{i\boldsymbol{q}\cdot\boldsymbol{r}_i} d_{i,\alpha}, \tag{2.6}$$

is the Fourier transformed dimer density

$$d_{i,\alpha} = 1 + \frac{J_{i,j_\alpha}}{J}\sigma_i^z\sigma_{j_\alpha}^z, \tag{2.7}$$

and j_α is the index of nearest neighbor to site i in the α direction. A more transparent representation of Ψ is in terms of the bond energy densities, ε_{ij}, which measure the average value of the bond energy between sites on sublattice i and j (bond (i,j)),

$$\varepsilon_{ij} = \frac{4}{N}\sum_a \varepsilon_{ij}^a, \tag{2.8}$$

where

$$\varepsilon_{ij}^a = \begin{cases} +1, & \text{if bond } (i,j) \text{ in plaquette } a \text{ is frustrated} \\ -1, & \text{if bond } (i,j) \text{ in plaquette } a \text{ is satisfied} \end{cases} \tag{2.9}$$

and the sum is over all plaquettes on the lattice. Using these variables, Ψ can be defined,

$$\Psi = \frac{1}{2}\left(\varepsilon_{34} + \varepsilon_{23}e^{i\pi/2} + \varepsilon_{12}e^{i\pi} + \varepsilon_{14}e^{3i\pi/2}\right), \tag{2.10}$$

Long-range ordering is associated with $|\Psi|$ taking a finite value, while the specific (symmetry-broken) ordering pattern is identified by the phase. In the columnar phases, one of ε_{ij} is equal to 1 (frustrated), while the rest are equal to -1 (satisfied), so Ψ takes one out of the values $e^{in\pi/2}$, $n = 1, 2, \ldots, 4$. The connection between the columnar states and the sublattice magnetizations is illustrated in Fig. 2·1.

2.3 $\mathbb{Z}_8$ versus $\mathbb{Z}_4$ Symmetry Breaking

We detect the order-parameter symmetries by plotting the probability distributions of m and Ψ accumulated during QMC simulations. In the ordered phase, we expect eight (four) δ-functions at the eight (four) values corresponding to the columnar spin (dimer) states in Fig. 2·1. These δ-functions smear for finite systems, appearing as highly peaked Gaussian distributions for large systems.

In the critical phase, we observe the emergent U(1) symmetry expected from the mapping to the height model [63, 73, 51]. We assign height differences to neighboring spins based whether or not the bond they share is frustrated (i.e., crossing a dimer) [73], as detailed in Sec 2.4.1. In the ordered phase, the height profile is "flat," with the height values bounded from above. In the critical and disordered phases, the model is in its "rough" phase, with a logarithmically diverging height profile. This behavior can be described by an effective elastic free energy with a periodic "locking" potential favoring the eight flat height configurations, i.e., the columnar states in Fig. 2·1. This effective free energy is precisely that of the 2D XY model, where the locking potential corresponds to an $q = 8$ state clock anisotropy term. This connection allows us to apply the renormalization group analysis of Ref. [65] to understand the observed behavior.

Jose et al. [65] first showed that the classical 2D q-state clock model is characterized by three temperature regimes if $q > 4$. At temperatures below the lower critical temperature T_{c1}, the clock term is relevant and the system orders into the $\mathbb{Z}_q$ clock phase. At T above T_{c1} but below the upper critical temperature T_{c2}, the clock term is irrelevant and the free energy reduces to that of the XY model in the Kosterlitz-Thouless (KT) phase [74]. In this phase, the system fluctuates between the flat height configurations as a consequence of the proliferation of bound topological defects (so called "screw dislocations" in the height field, which map to spin vortices in the 2D XY model), thus resulting in the U(1) symmetric distributions that we observe. Finally, above T_{c2} the critical phase melts into the disordered phase as defects proliferate.

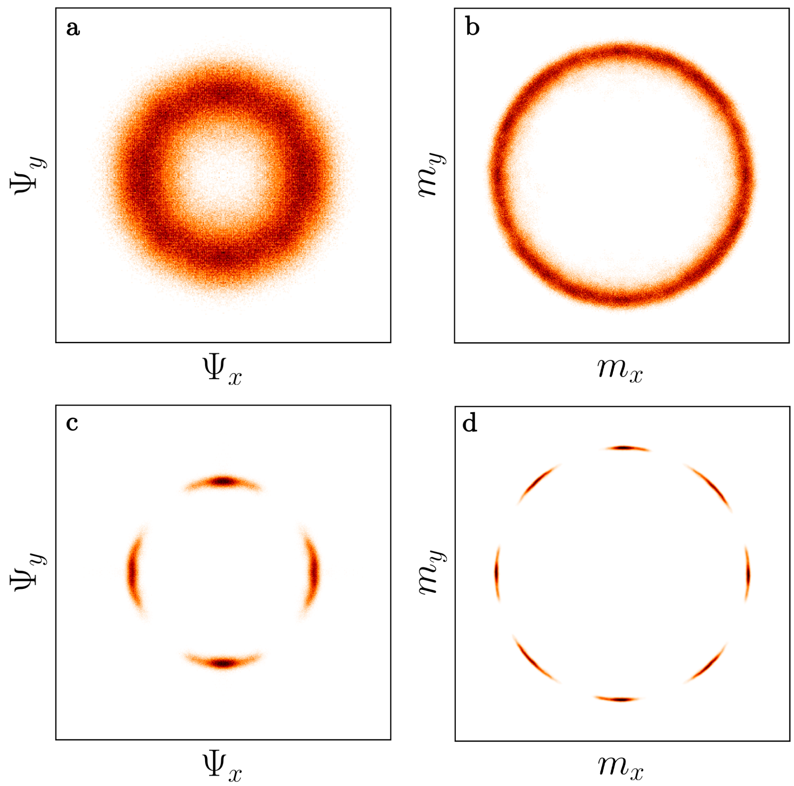

Figure 2·2: Distributions of (Ψ_x, Ψ_y) in (a,c) and (m_x, m_y) in (b,d), collected in several independent simulations (each initialized in one of the four columnar states, to prevent trapping in states with topological defects) and symmetrized using lattice rotations and reflections. The system size is $L = 96$ and $\Gamma/J = 0.43$, with $T/J = 0.21$ in (a,b) (in the critical phase) and $T/J = 0.014$ in (c,d) (in the ordered phase).

An example of the symmetry reduction in the ordered phase is shown in Fig. 2·2, where (a) and (b) are collected from simulations in the critical phase, while (c) and (d) are from the ordered phase. While we detect emergent U(1) symmetry in both order parameters, a U(1) phase would not be expected for a primary $\mathbb{Z}_4$ dimer order parameter at $T > 0$, given the $q > 4$ criterion in the clock model [65]. However, with the spin order parameter corresponding to $q = 8$, the critical phase is expected.

2.4 Scaling at $T > 0$

As a quantitative characterization of the primary and secondary natures of the two order parameters, we compare the scaling of their respective correlation functions in the critical phase. Within the q-state clock-model description, the spin-spin correlations should decay algebraically with a scaling exponent η that varies continuously with the temperature [97]. The value of η at the upper and lower critical temperatures are known, $\eta = 1/4$ and $\eta = 4/q^2$, respectively [65, 51]. To extract η, we examine the magnitudes of both order parameters, which in this phase should scale with the lattice length L as

$$|m|^2 \propto L^{-\eta_m}, \qquad |\Psi|^2 \propto L^{-\eta_d}, \tag{2.11}$$

where we leave open the possibility that $\eta_m \neq \eta_d$. In Fig. 2·3(a) and 2·3(b), the order parameters are plotted versus system size for a range of temperatures between $T_{\mathrm{c}1}$ and $T_{\mathrm{c}2}$. As predicted, they scale algebraically with L, and we extract $\eta_m(T)$ and $\eta_d(T)$ by fitting data to Eq. (2.11). Results are shown versus T in Fig. 2·3(c).

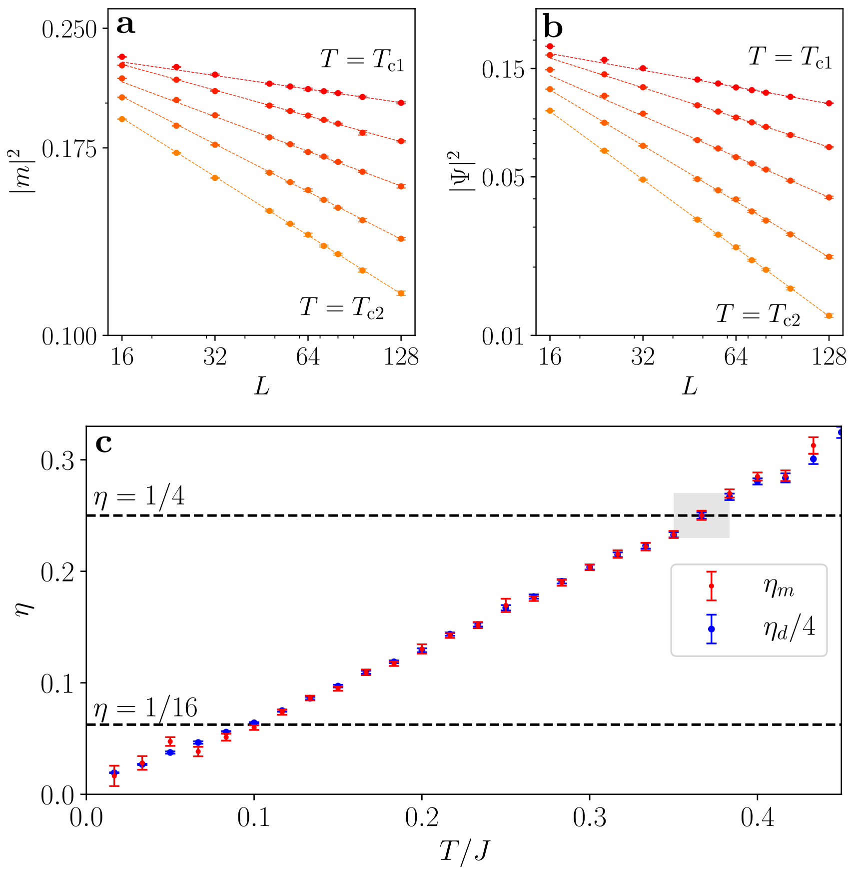

Figure 2·3: Log-log plots of the magnitude of the primary (a) and secondary (b) order parameter versus the system size L for a range of temperatures at $\Gamma/J = 0.67$. The dashed lines are power-law fits to the largest eight system sizes, with T increasing with color brightness (red to orange). (c) Anomalous dimensions $\eta_{m,d}$ versus T for the primary (red) and secondary (blue) order parameters, extracted from fitting data to Eq. (2.11). The exponents align if η_d is rescaled by a factor 1/4. The gray box denotes $T_{c2} \pm \sigma$ from Binder crossing results. The dashed lines are at the predicted values at the phase boundaries.

The primary order parameter scales with the expected exponent $\eta_m = 1/4$ at the upper critical temperature T_{c2} extracted using the crossing of the Binder cumulant.

We define the cumulants for both order parameters [12]:

$$U_d = 2 - \frac{\langle |\Psi|^4 \rangle}{\langle |\Psi|^2 \rangle^2}, \quad U_m = 2 - \frac{\langle |m|^4 \rangle}{\langle |m|^2 \rangle^2}. \tag{2.12}$$

At the critical temperature, the Binder cumulant should be system-size independent, and we indeed find that curves of U versus T for different L values intersect each other close to a point $T_{\rm c2}$. We extract the temperatures at which the cumulants for pairs of system sizes L and $2L$ intersect, and extrapolate these to the infinite-size $T_{\rm c2}$ using the expected finite-size scaling form (shift of a definition of the critical point with the system size) at a KT transition [55]:

$$T_{c2}(L) = T_{c2}(\infty) + \frac{a}{\log^2(aL)}, \tag{2.13}$$

An example of this procedure is shown in Fig. 2·4 for $\Gamma/J = 1.50$. The fits to the primary (red) and secondary (blue) binder crossing temperatures give infinite-size T_{c2} valuess that agree well with each other to within error bars estimated using bootstrapping of the data.

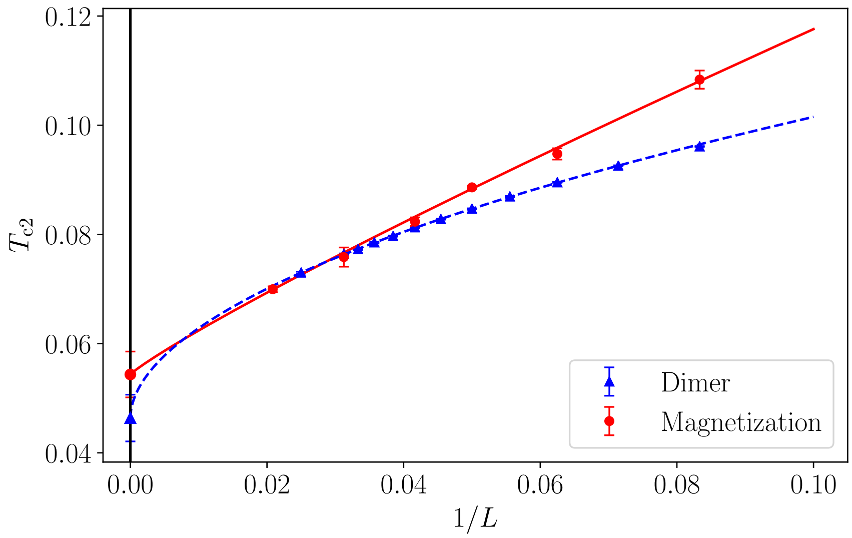

Figure 2·4: Finite-size scaling of the Binder crossing temperature at $\Gamma/J = 1.50$. The temperatures for the magnetization (red) and dimer (blue) order parameters were extracted for pairs of systems of size $(L, 2L)$. The error bars on the extrapolated $T_{c2}(\infty)$ values were estimated by a bootstrapping procedure.

Below this temperature, the exponent linearly decreases, but we do not observe any further crossing of the Binder cumulant. However, we do find that the statistical quality of the fits deteriorates below the temperature at which $\eta_m = 1/16$, which is the predicted value at the lower transition point T_{c1} [65], where the system orders. Figure 2·5 shows the reduced χ^2 for these fits versus temperature compared to the reference value corresponding to the 95% confidence level (black dashed line). The results are indeed consistent with the power law form no longer being applicable below T_{c1} and above T_{c2}.

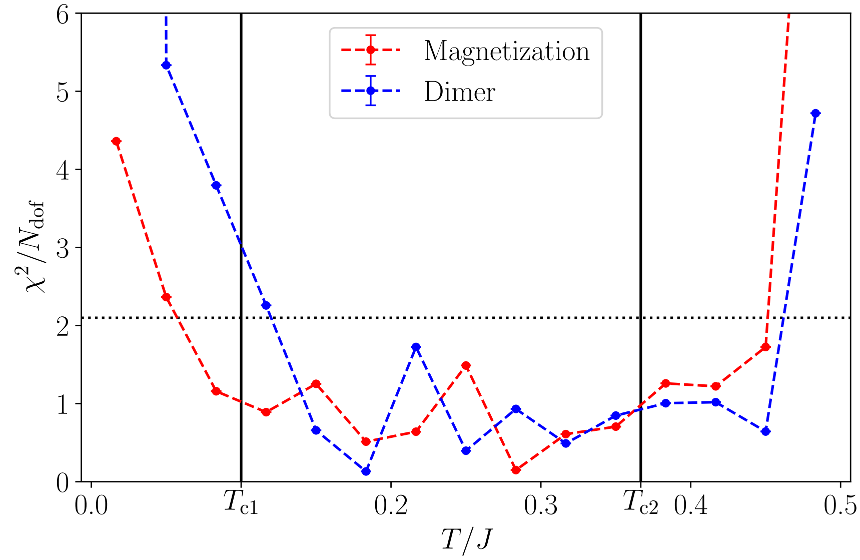

Figure 2·5: Reduced χ^2 values from the power law fits used to extract η_m (red) and η_d (blue) in Fig. 3 of the main text. The black dashed line denotes the value corresponding to the 95% confidence level for the number of degrees of freedom of the fits. The reduced χ^2 begins to exceed the 95% confidence level below T_{c1} and above T_{c2}, indicating that the order parameters no longer scale as power laws.

While the linear decrease in η appears to continue below T_{c1}, the deterioration of the power-law fit used to extract the exponents below this temperature implies that this trend should not be given any credence. Thus, our numerical results are consistent with the theory, and we can use $\eta_m = 1/4$ and $\eta_m = 1/16$ to set more precise upper and lower boundaries.

The same behavior is observed for the dimer order parameter, except that the value of the η_d is consistently approximately four times larger than η_m. After rescaling η_d by a factor of four, the results match within statistical errors, suggesting the relation $\eta_d = 4\eta_m$. This relationship between η_m and η_d can be explained by mapping to the height model. By representing the spin operators in terms of height variables, one can relate the spin-spin and dimer-dimer correlation functions to the logarithmically diverging height difference profile; the observed factor of four relating η_d and η_m then emerges, as we detail below.

2.4.1 Height Model Mapping

Here, we make explicit the mapping between the spins, dimers, and height variables of the FFTFIM. Following the prescription in Ref. [63, 73], we start by identifying the 2-to-1 mapping between the ground states of the classical fully frustrated Ising model and those of the purely kinetic dimer model on the square lattice [138]. These degenerate states are defined by the configurations where each plaquette contains a single frustrated bond. The dimer model on the square lattice has been studied extensively [35, 110, 4] and its height representation is known [51]. We repeat the mapping here for the specific sake of clarifying the relationship between the spin and dimer order parameters, which was not emphasized in previous treatments.

The mapping proceeds as follows: begin by dividing the dual lattice into two sublattices A and B, where dimers connect dual lattice sites from sublattice A to B. Assign an arbitrary height value to a site on the spin lattice and move clockwise around a dual lattice site. If the dual lattice site is on sublattice A and a dimer is crossed, the height value at the following spin lattice site is increased by 3, otherwise it is reduced by 1; see Fig. 2·6. If the dual lattice site is on sublattice B, the same rules apply with signs reversed; -3 and +1, respectively.

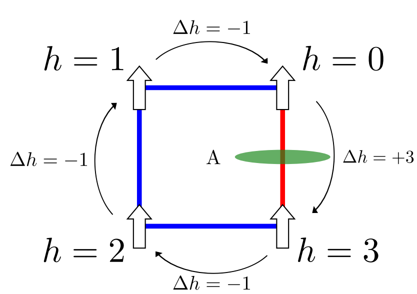

Figure 2·6: Height model mapping rules for a plaquette at dual sublattice A. Height variables h are labeled at the spin sites.

In the ordered phase, the dimers are ordered along columns or rows, corresponding to a height profile that is "flat," i.e., the height value at any given lattice site is bounded from above. In the critical and disordered phases, plaquettes with three frustrated bonds begin to populate the system in pairs; these are the aforementioned screw dislocations (topological defects). The height model is then "rough", with a logarithmically diverging height profile. This behavior can be described by an effective elastic free energy with a periodic "locking" potential favoring flat height configurations [19, 4];

$$F(\{h(\vec{r})\}) = \int d\vec{x}\left[\frac{K(T)}{2}\left|\vec{\nabla}h(\vec{r})\right| + V\cos\left(2\pi h(\vec{r})\right)\right], \tag{2.14}$$

where $h(\vec{r})$ is the coarse-grained height field (height variables averaged over a plaquette), $K(T)$ is the temperature dependent stiffness constant, and $V > 0$ is the strength of the locking potential.

As discussed above, for $T < T_{c1}$ the locking potential is a relevant perturbation and the system orders [65, 51]. For $T_{c1} < T < T_{c2}$, screw dislocations begin to emerge in pairs throughout the system. For $T > T_{c2}$, the interaction between these defects becomes sufficiently weak for the pairs to unbind, spreading throughout the system and disordering the critical phase [63, 64, 73].

In the critical (rough) phase, the locking potential is irrelevant, and the system is described by a free Gaussian theory. The long distance height difference correlations can be easily calculated [97, 51, 65]:

$$\frac{1}{2}\langle\left|h(\vec{r}) - h(\vec{0})\right|^2\rangle \propto \frac{\ln(r)}{2\pi K}. \tag{2.15}$$

We relate this form to the spin-spin and dimer-dimers correlations by noting that the spins are periodic in the height variables, with a period of 8. This can be seen by turning the spins around a plaquette twice as depicted in Fig. 2·7.

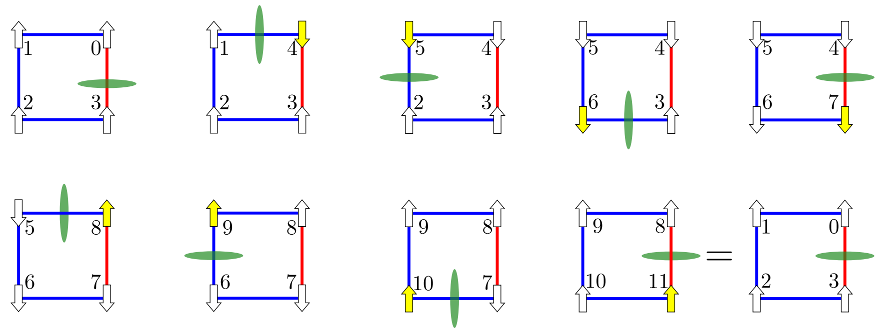

Figure 2·7: Demonstration of the periodicity of the spin variables with the height variables. Upon traversing the plaquette, each spin (marked in yellow) is flipped sequentially (left to right, upper then lower row) and the height variables are computed according to the rules laid out in the text. The final spin configuration is identical to the initial one, but with each height increased by 8, implying that the spins must be periodic in the heights with period 8.

In light of this periodicity, it is convenient to recast the height field $h(\vec{r})$ such that instead of being defined modulo 8, it is defined modulo 2π,

$$\theta(\vec{r}) = \frac{2\pi}{8} h(\vec{r}), \tag{2.16}$$

and

$$F(\{\theta(\vec{r})\}) = \int d\vec{x} \left[\frac{K(T)}{2} \left(\frac{8}{2\pi} \right)^2 \left| \vec{\nabla}\theta(\vec{r}) \right| + V \cos\left(8\theta(\vec{r})\right) \right], \tag{2.17}$$

Under this transformation, the free energy is now identical to that of the 2D classical 8-state clock model, and the origin of the 8-state clock term is now transparent.

The periodicity also implies that the spins can be written as a Fourier series:

$$\sigma^z(\vec{r}) = \sum_G \tilde{\sigma}^z_G e^{iGh(\vec{r})} \tag{2.18}$$

where $G = \frac{2\pi}{8} n$ with $n = 1, 2, 3, \ldots$. The spin-spin correlations for a single Fourier mode is given by

$$\langle e^{iGh(\vec{r})} e^{-iGh(\vec{0})} \rangle = e^{-\frac{1}{2} G^2 \langle |h(\vec{r}) - h(\vec{0})|^2 \rangle}, \tag{2.19}$$

and clearly the smallest G value in the expansion in Eq. (2.18) will dominate. Thus, we can express the coarse-grained spin and dimers as

$$\begin{aligned} \sigma^z(\vec{r}) &\propto e^{i\frac{2\pi}{8}h(\vec{r})} \\ d(\vec{r}) &\propto \sigma^z(\vec{x})\sigma^z(\vec{r}+\hat{x}) \approx (\sigma^z(\vec{r}))^2 \propto e^{i\frac{4\pi}{8}h(\vec{r})}. \end{aligned} \tag{2.20}$$

Plugging in the height difference correlations in the rough phase, Eq. (2.15), we can extract the anomalous dimensions for the spins and dimers:

$$\begin{aligned} \langle \sigma^z(\vec{x})\sigma^z(\vec{0})\rangle &\propto e^{-\frac{1}{2}\left(\frac{2\pi}{8}\right)^2 \langle |h(\vec{r})-h(\vec{0})|^2\rangle} \\ &\propto r^{-\pi/32K} \propto r^{-\eta_m} \\ \langle d(\vec{r})d(\vec{0})\rangle &\propto e^{-\frac{1}{2}\left(\frac{24pi}{8}\right)^2 \langle |h(\vec{r})-h(\vec{0})|^2\rangle} \\ &\propto r^{-\pi/8K} \propto r^{-\eta_d} \end{aligned} \tag{2.21}$$

This mapping implies that $\eta_d = 4\eta_m$, which we have confirmed numerically in Fig. 3c in the main text.

The above derived relationship between η and $K(T)$ allows us to determine the values of anomalous dimension at the boundaries of the critical phase. Because the effective free energy in Eq. (2.17) is exactly that of the 2D classical XY model with an eight-state clock anisotropy term, we can use the renormalization group analysis preformed in Ref. [65] to determine the critical values of the coupling $K(T)$ separating the critical phase from the ordered and disordered phases.

The lower and upper critical temperatures correspond to $K(T_{c1}) = \pi/2$ and $K(T_{c2}) = \pi/16$, between which the system is critical (cf. with Ref. [97]). The values of η_m at the phase boundaries are thus

$$\eta_m(T_{c1}) = 1/16, \quad \eta_m(T_{c2}) = 1/4. \tag{2.22}$$

2.5 Finite Temperature Phase Diagram

The finite temperature behavior of this model can be organized into a phase diagram, which we plot in Fig. 2·8. In summary, at temperatures above $T_{c2}(\Gamma)$, or transverse field values above Γ_c, the system is in the paramagnetic (PM) phase, while for $\Gamma < \Gamma_c$ and $T_{c1}(\Gamma) < T < T_{c2}(\Gamma)$, it is in a critical KT phase. For $\Gamma < \Gamma_c$ and $T < T_{c1}(\Gamma)$, the system is ordered, exhibiting Z_8 and Z_4 symmetry breaking in the primary and secondary order parameters, respectively.

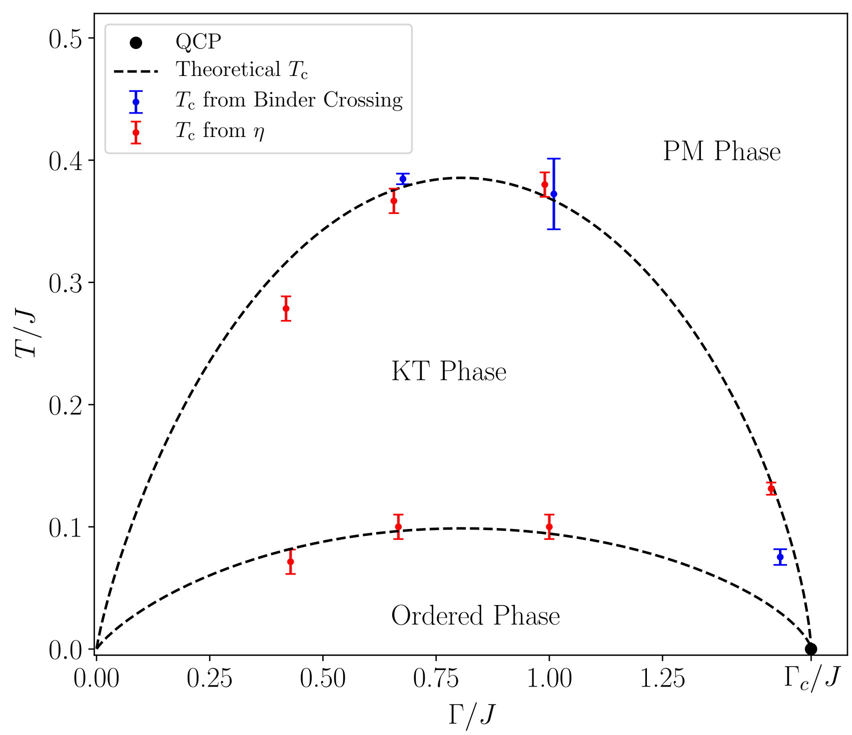

Figure 2·8: Phase diagram of the square-lattice FFTFIM in the plane of transverse field Γ and temperature T. The QCP is shown as the black circle at $T = 0$. The black dashed lines are the theoretical predictions for the phase boundaries made in Ref. [63, 73], and the red and blue points are the transition temperatures identified using our QMC simulations. The theoretical phase boundaries have overall unknown factors that we have adjusted for best agreement at the maximums.

The phase diagram was constructed in the following way: The black point at $T = 0$

in Fig. 2·8 is the quantum critical point (QCP), which was identified in Ref. [143]; we further refined its location in the present work as detailed below in Sec. 2.6.1. The black dashed lines show the theoretical scaling of T_{ci} versus Γ predicted by Ref. [63, 73], in which an overall factor has been adjusted to fit our data at the maximums. These phase boundaries are formally only valid for $\Gamma_c - \Gamma \ll \Gamma_c$ (near the QCP), but we include them here all the way to $\Gamma = 0$ as a visual aid. The blue points were identified using the Binder crossing analysis shown in Fig. 2·4 and the red points were determined by identifying the temperatures where the correlations in the critical phase decay as a power law with the expected exponents (e.g. Fig. 2·3).

2.6 Scaling at $T = 0$

At the quantum critical point, there must be a different relationship between η_d and η_m. Given the irrelevance in 2+1 dimensions of $\mathbb{Z}_4$ or $\mathbb{Z}_8$ perturbations to a U(1) order parameter, the scaling of the primary order parameter will be dictated by the conventional 3D XY critical exponent $1 + \eta_{3DXY}$ [21], not the 2D clock model as is the case at finite temperature. This was indeed confirmed in previous studies where a spin-based order parameter was used [143]. However, to explain the critical scaling of the dimer order parameter, we must reference other aspects of the field theory. Here we analyze both order parameters using simulations at $T = 1/L$ (i.e. in the ground state).

At the QCP, we extract η_m from the asymptotic long distance critical spin-spin correlation function, which is expected to scale as

$$\lim_{r\to\infty} C_M(\vec{r}) = \lim_{r\to\infty} \langle \sigma^z_{m_i}(\vec{r})\sigma^z_{m_i}(0)\rangle \sim \frac{1}{r^{1+\eta_m}}, \tag{2.23}$$

where $\sigma^z_{m_i}(\vec{r})$ is the spin operator at a site $\vec{r}$ on sublattice i and we expect $\eta_m = \eta_{3DXY}$. In connection to the relevant field theory, $1+\eta_{3DXY} = 2\Delta_\phi$, where $\Delta_\phi \approx 0.519088$ [28]

is the scaling dimension of the order parameter of the 3D O(2) model. This is often referred to as a charge-1 (or spin-1) operator [28, 46], indicating that it corresponds to a perturbation field, e.g., $h\cos(\theta)$, inducing order in a single direction in the O(2) space, so that the degeneracy is completely lifted. A corresponding perturbation in FFTFIM would be one that fully breaks the $\mathbb{Z}_8$ symmetry in the ordered state, favoring one of the eight columnar spin configurations.

A perturbation that couples an external field to the secondary order parameter would not fully break the $\mathbb{Z}_8$ symmetry of the ground state, but would favor the two spin configurations of given columnar dimer state. Accordingly, in the low-energy U(1) theory the perturbation should be charge-2 (or spin-2 traceless symmetric) of the form $h\cos(2\theta)$, which can also be accomplished with products of ϕ components, e.g., $h\phi_1\phi_2$. This operator has scaling dimension often referred to as Δ_t [28, 46],

To test the scaling form $r^{-(1+\eta_d)} = r^{-2\Delta_t}$, with $\Delta_t \approx 1.23629$ [28], we analyze the oscillating part $C_D(r)$ of the dimer-dimer correlation function. The raw dimer-dimer correlation function is defined as

$$D_x(\vec{r}) = \langle d_x(\vec{r}) d_x(0) \rangle, \tag{2.24}$$

where $d_x(\vec{r})$ is the x-oriented dimer operator at a site $\vec{r}$. However, due to the staggered ordering of the dimers in the columnar phase, the order parameter of interest is the difference between the modulated correlations;

$$D_x^*(\vec{r}) = D_x(\vec{r}) - \frac{D_x(\vec{r}-\hat{y}) + D_x(\vec{r}+\hat{y})}{2}. \tag{2.25}$$

Averaging over both horizontal and vertical dimer orientations, the expected scaling is

$$\lim_{r\to\infty} C_D(\vec{r}) = \lim_{r\to\infty} \frac{D_x^*(\vec{r}) + D_y^*(\vec{r})}{2} \sim \frac{1}{r^{1+\eta_d}} \tag{2.26}$$

For both $C_M(r)$ and $C_D(r)$, we evaluate the correlation fucntion at the largest separation for a periodic, finite system of linear size L, $\vec{r} = (L/2, L/2)$, which will pick up the desired scaling $\sim L^{-(1+\eta)}$. Since the connection between the primary order parameter and Δ_ϕ is well known, we first used its scaling behavior to refine the value of Γ_c reported in Ref. [143]. We then calculated $C_M(L/2)$ and $C_D(L/2)$ at this value of Γ_c; their scaling behaviors are shown in Fig. 2·9. The results match very well when using the scaling dimensions obtained in recent numerical conformal bootstrap calculations [28]: $\Delta_\phi \approx 0.519088$ and $\Delta_t \approx 1.23629$.

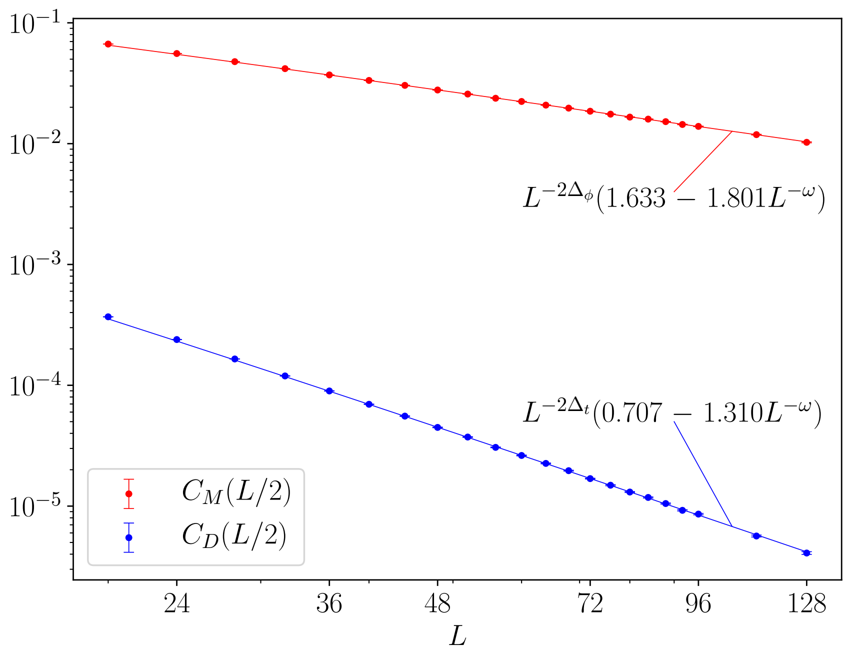

Figure 2·9: Dimer-dimer (blue) and spin-spin (red) correlation functions at the quantum critical point, $\Gamma/J = 1.5768$. The lines are fits of the $L \geq 44$ data to the expected scaling form $\propto L^{-2\Delta_{s,t}}(a + bL^{-\omega})$, with $\Delta_\phi = 0.519088$, $\Delta_t = 1.23629$ [28], and the correction exponent $\omega = 0.789$ [45] in the 3D O(2) universality class. The reduced χ^2 values are ~ 0.8 and ~ 0.9 for the spin and dimer correlations respectively.

2.6.1 Determining Γ_c

Here, we briefly discuss how we use the known power law scaling of $C_M(L/2)$ to determine Γ_c. We measured the correlations function for various system sizes and

Γ values in a small range around the previously reported Γ_c and compared the χ^2 from power law fits of the form $C_M(L/2) \sim r^{-2\Delta_\phi}$, to the data at these different Γ values. To reduce the error bars on the measured data, we interpolated the correlation function data for a given system size using second-order polynomials for several data sets in the Γ range $[1.576, 1.578]$.

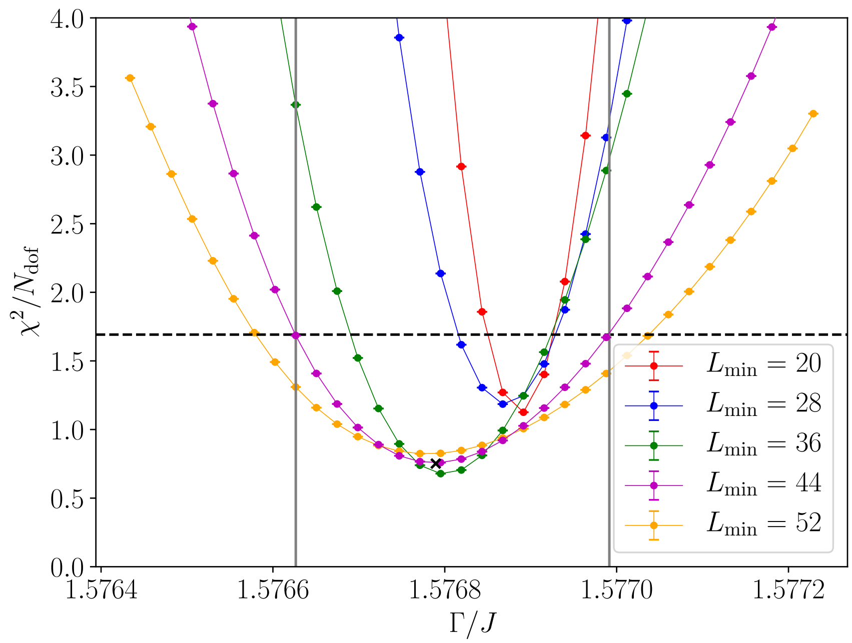

Figure 2·10: Reduced χ^2 from power-law fits to the spin-spin correlation function data versus Γ. The χ^2 minimum converges as the smallest L value used in the fit is increased. Using the number of degrees of freedom for the $L_{\rm min} = 44$ fit, we determine the Γ values within the 95% confidence interval of the χ^2 distribution, indicated here by the two vertical black lines. The error bars were computed by bootstrapping the raw binned data sets.

Reduced χ^2 versus Γ for the spin-spin correlation function is shown in Fig. 2·10, where we see convergence in the location of the χ^2 minimum as the minimum system size used for the fit is increased. We judge that any remaining systematic errors from scaling corrections are negligible for $L_{\rm min} = 44$, and the resulting critical point is then

$\Gamma_c = 1.57680 \pm 0.00009$. To determine the error bar on Γ_c, we located Γ values where the reduced χ^2 is equal to the the 95% confidence level for the number of degrees of freedom used, denoted by the two vertical black lines.

2.6.2 The Classical 3D 8-state Clock Model

When analyzing the critical behavior of quantum systems, it is often insightful to consider a classical analog, typically defined in one higher dimension. In this case, the simplest classical model exhibiting the same critical scaling behavior as the FFTFIM is the 3D 8-state clock model. The generalized 3D classical q-state clock model is defined by the Hamiltonian

$$H_q = -\sum_{\langle i,j \rangle} \cos\left(\theta_i - \theta_j\right) - h_q \sum_i \cos\left(q\theta_i\right), \tag{2.27}$$

with $\theta_i \in [0, 2\pi)$ defining the orientation of a two-component spin on cubic lattice site i. At fixed h_q, below $T = T_c$, the systems orders into a Z_q symmetry breaking phase. The Z_q perturbation for $q \geq 4$ is known to be irrelevant at the critical point, so that the 3D O(2) universality class applies (though changing the symmetry broken in the ordered phase).

For $q = 8$, the microscopic symmetries of the model exactly match those of the square-lattice FFTFIM, and we here study this clock model as a bench-mark case of primary and secondary order parameters. Just as in the FFTFIM, we can define charge-1 and charge-2 scalar order parameters for the clock models:

$$m_x = \sum_i \cos(l\theta_i), \quad m_y = \sum_i \sin(l\theta_i), \quad m_l^2 = \langle m_x^2 \rangle + \langle m_y^2 \rangle, \tag{2.28}$$

where l labels the charge.

Although our primary interest here is in the $q = 8$ case, we carried out simulations at the critical points of three different versions of the model, namely, the XY model (with no clock field, corresponding to $q \to \infty$), the 8-state hard clock model (with the spins restricted to the q discrete states, corresponding to $h_8 \to \infty$), and the 4-state soft clock model with anisotropic field $h_4 = 1$. We expect the same scaling dimensions in all cases on account of the irrelevance of the clock perturbations. The critical points of the first two models can be found in Ref. [45], while $T_c = 2.20465(1)$ of the third model was reported in [122]. In an improved analysis, we found this value to be about two error bars too low and here report results obtained with $T = 2.20467$.

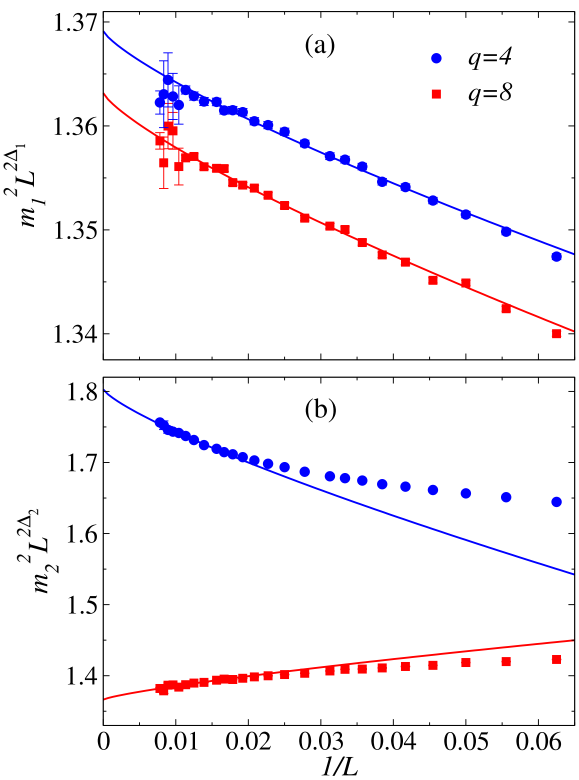

Figure 2·11: Scaling behaviors of m_l^2, with charge $l = 1$ (a) and $l = 2$ in (b), at the estimated critical temperatures for the $q = 4$ and $q = 8$ clock models (using $1/T_c = 0.45416467$ for $q = 8$ [45] and $T_c = 2.20467$ for $q = 4$. The curves show the form Eq. (2.29) with only the constants a and b adjusted for best fits and with results for the smaller system sizes exccluded systematically until a statistically good reduced χ^2 value is obtained.

In Fig. 2·11 we fit our data for $q = 4$ and $q = 8$ to the expected scaling form with the leading finite-size correction included;

$$m_l^2 L^{2\Delta_l} = a + bL^{-\omega} \tag{2.29}$$

with $\Delta_1 = \Delta_\phi = 0.519088$, $\Delta_2 = \Delta_t = 1.23629$ [28], and $\omega = 0.789$ [45]. Only the prefactors a and b were fitted. In the case of the XY model, the results fall so close to those for $q = 8$ that we have not included them in the figure for clarity (we also note that the error bars are larger for the XY model after simulations of comparable length).

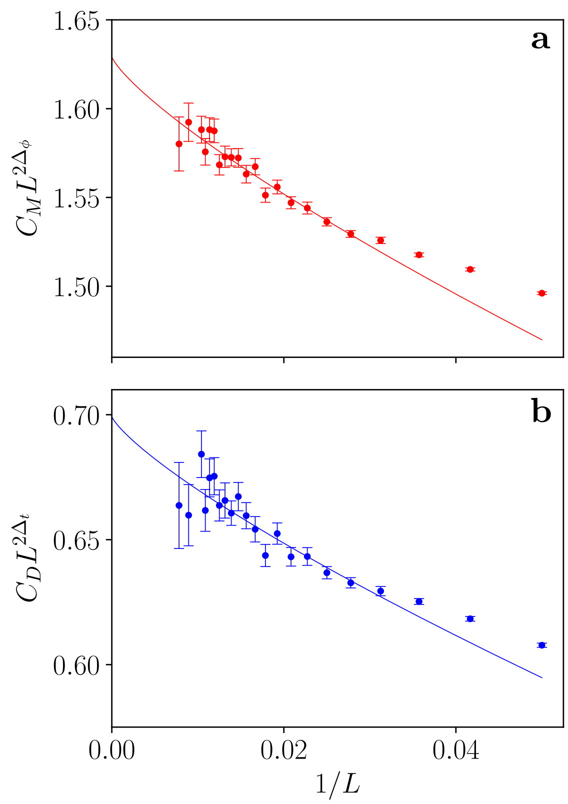

Figure 2·12: Scaling behaviors of $C_M(L/2)$ and $C_D(L/2)$, plotted as in Fig. 2·11.

To compare the clock results more directly to the FFTFIM results shown in Fig. 2·9, in Fig. 2·12 we show the FFTIFM data analyzed in the same way as the clock results in Fig. 2·11, with the expected leading power laws of L divided out and fitting only to the expected correction given by Eq. (2.29). Though the sign of the (non-universal) amplitude of the correction to the dimer correlations in Fig. 2·12(b) is different from that in the $q = 8$ result in Fig. 2·11(b), the overall behavior of the corrections (the leading term fitted to and the higher-order corrections causing deviations from the fit) is similar. The primary order parameter of the FFTFIM in Fig. 2·12(a) seems to have more significant subleading corrections than the clock results in Fig. 2·11, but the large-L asymptotic form is again similar and consistent with the leading correction.

The excellent agreement with the expected exponents in both the FFTFIM and clock model confirms without doubt the emergent U(1) symmetry and the primary and secondary nature of the order parameters in the FFTFIM.

2.7 Quantum Annealing of the FFTFIM

Our interest in this system was partially motived by discussions with the superconducting qubit quantum annealing company D-Wave Quantum, as they were in the early stages of implementing this model on their devices and wanted a theoretical benchmark to compare with their results [5]. The focus of their study was the out-of-equilibrium dynamics of the system, which they explore through quantum annealing (QA) and analyze using extensions of the Kibble-Zurek (KZ) mechanism [67, 159].

In 1976, Tom Kibble developed a theory on the formation of topological defects during the cooling of the early universe [67]. Describing the universe by a gauge theory exhibiting spontaneous symmetry breaking (the Higgs field), Kibble proposed that one should expect domain walls, monopoles, and vortices to populate as the universe passed through a critical temperature T_c. Throughout his paper, he references similar phenomena in condensed matter systems, specifically superconductors and superfluids. Roughly a decade later, Wojciech Zurek took this analogy one step further, applying Kibble's theory directly to the cooling of ^{4}He through its superfluid transition [159]. In subsequent papers, Zurek proposed tabletop experiments where Kibble's theory could be empirically confirmed [160, 161]; one could explore the cosmos in a condensed matter laboratory.

In the following years, this so called KZ-mechanism was rigorously tested and extended using numerics [77, 154] and experiment [29, 92]. Such extensions have been used to describe the physics of out-of-equilibrium classical systems as they approach their critical points [27, 31], making this a valuable tool to use in dynamic scaling analysis for nonequilibrium quenches, i.e. simulated classical annealing (SA). In their 2014 paper, Liu, Polkovnikov, and Sandvik developed a KZ ansatz describing the scaling of physical observables measured at a classical phase transition as a function of system size and annealing velocity [81]. Their methods allow one to extract important information about the dynamics of critical systems, such as the dynamic scaling exponent z, which also provides insight into the efficicacy of the annealing routine.

In the following section, we will lay out the groundwork for this KZ ansatz, drawing heavily from Ref. [81].

2.7.1 The KZ Mechanism

The main idea behind Kibble and Zurek's argument was that a system undergoing a quench (or anneal) to the critical point can only stay adiabatic if the annealing time to reach the critical temperature (τ_a) is of the order of the relaxation time of the system (τ_r),

$$\tau_r \sim \tau_a. \tag{2.30}$$

For the following derivations, we will assume we are annealing to the critical point T_c, not past it. In this case, the annealing time at a given temperature T is simply

$$\tau_a = \frac{T - T_c}{v}, \tag{2.31}$$

where v is the annealing velocity. The relaxation time is the characteristic time scale of the microscopic dynamical properties of the system, which can be related to the correlation-length, the characteristic *length* scale of the system, by the dynamical exponent z:

$$\tau_r(T) \sim (\xi_T)^z \sim |T - T_c|^{-z\nu}, \tag{2.32}$$

where ν is the equilibrium correlation-length exponent. At a fixed temperature, we can define a velocity $v_{\mathrm{KZ}}(T)$ separating the adiabatic and diabatic regimes by plugging in Eq. 2.31 and Eq. 2.32 into the relation in Eq. 2.30 (i.e. at this temperature, what velocity would one have to take, if annealing to the critical temperature, such that the two time scales are equivalent):

$$v_{\mathrm{KZ}}(T) \sim |T - T_c|^{1+z\nu} \tag{2.33}$$

From this expression one can immediately see that in the thermodynamic limit, it is impossible to stay adiabatic during the full extent of the anneal; the correlation-length, and thus the relaxation time, diverges as the critical point is approached, requiring an infinitesimally slow velocity for adiabaticity to be maintained. This is nicely visualized with the schematic in Fig. 2·13, taken from Ref. [31].

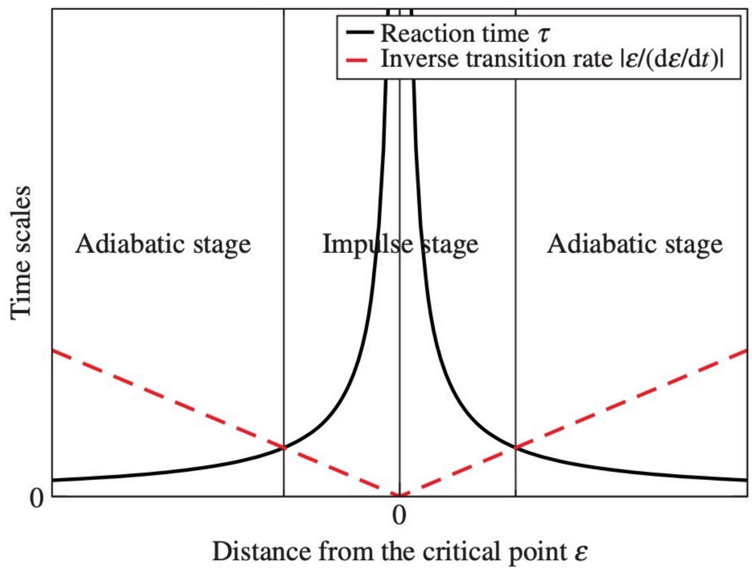

Figure 2·13: Schematic displaying the relation between the diverging relaxation time (labeled "reaction time" in the legend) and the linearly decreasing annealing time (labeled "inverse transition rate" in the legend). When these two time scales equal one another, the system leaves the adiabatic regime and enters the frozen out diabatic regime. Figure taken from Ref. [31]

However, for finite systems this is not the case. We can see this by instead relating $v_{\mathrm{KZ}}(T)$, Eq. (2.33), to the corresponding diverging length scale ξ_T,

$$\xi_T \sim |T - T_c|^{-\nu}, \tag{2.34}$$

by plugging Eq. (2.34) into Eq. (2.33),

$$v_{\mathrm{KZ}}(T) \sim \xi_T^{-(z+1/\nu)}. \tag{2.35}$$

For finite systems in d dimensions with volume L^d, the largest length scale is L, so one can define a lower bound on the characteristic velocity separating the adiabatic and diabatic regimes by setting $\xi_T = L$, just as in the the standard finite-size scaling procedure [129]:

$$v_{\mathrm{KZ}}(L) \sim L^{-(z+1/\nu)}. \tag{2.36}$$

For a finite system of linear length L, if $v < v_{KZ}(L)$ the nonadiabtic response of the system will be very small, and thus can be treated as just a perturbation. On the other hand, if $v > v_{KZ}(L)$, adiabaticity breaks down.

2.7.2 Generalized KZ Finite-Size Scaling

Using the scaling relation derived above, one can define a finite-size scaling type description of an out-of-equilibrium system, quenched towards its critical point. Here, we will only reproduce the relevant formalism, and refer the reader to Ref. [81] for a complete derivation. In the flavor of the finite-size scaling hypothesis, one can describe the scaling of an observable A by a nonsingular function f of the system size, the distance (in temperature) to the critical point, and the annealing velocity,

$$A(L, T, v) = L^{\kappa/\nu} f(L/\xi_T, v/v_{\mathrm{KZ}}) = L^{\kappa/\nu} F\left(\left(T - T_c\right) L^{1/\nu}, vL^{z+1/\nu}\right), \tag{2.37}$$

where κ is an exponent depending on the universality class of the transition and the specific observable A, and F is a simple rescaling of f. Assuming one anneals to the critical point, this scaling form simplifies further,

$$A(L, T_c, v) = L^{\kappa/\nu} F\left(vL^{z+1/\nu}\right). \tag{2.38}$$

This relation is a valuable tool when analyzing the results of annealing protocols, as it gives a direct, reliable means of extracting the dynamical exponent z. This is done as follows: one performs annealing runs to the critical point for a series of system sizes, using a series of annealing velocities, and measures the value of an observable of interest at the end of each run. The results can then be analyzed via a “data-collapse” [115] according to the scaling form in Eq. (2.38). For a simple example we consider the 2D Ising model and the observable m^2, the magnetization squared. The critical exponent κ describing the scaling of the magnetization $|m|$ is typically referred to as β, so the scaling relation reads,

$$\langle m^2 \rangle = L^{-2\beta/\nu} F\left(vL^{z+1/\nu}\right) \to \langle m^2 \rangle L^{2\beta/\nu} = F\left(vL^{z+1/\nu}\right). \tag{2.39}$$

For a series of system sizes $\{L\}$ and velocities $\{v\}$, we can define the variables

$$x_L = vL^{z+1/\nu}, \quad y_L = \langle m^2 \rangle L^{2\beta/\nu} \tag{2.40}$$

and the relation in Eq. (2.39) will now read

$$y_L = F(x_L). \tag{2.41}$$

These rescaled points x_L and y_L, measured in simulation, will fall a single curve (with corrections [115]), dictated by the scaling function F. Assuming one knows critical exponents ν, and β, this data collapse procedure allows one to determine z in a systematic and methodical way by optimization (i.e. tuning z until the data points collapse in a statistically acceptable manner).

The KZ mechanism, and the finite-size scaling procedure outlined above, does not only apply to classical phase transitions, but also *quantum* phase transitions [162, 100, 101]. Using the exact same formalism, these tools were applied to quantum systems passing through their QCP's [82] (in this case, the diverging length scale is associated with a QPT driven by tuning a parameter in the Hamiltonian, such as a transverse magnetic field). This gives rise to an interesting question: does simulated *quantum* annealing (SQA) provide an advantage over SA, i.e. how does the scaling of v_{KZ} compare for a quantum system and its classical analog? And with the recent advancements of quantum computing technology, a perhaps even more scintillating prospect is the extension of this to real QA. This has been one of the focuses of exploration in the field of quantum computing, with D-Wave acting as a leader in the effort to understand the viability of QA [69, 68, 71, 5, 70].

2.7.3 Quantum Annealing of the FFTFIM

The scaling dimensions extracted in our work on the FFTIM are crucial in applying the KZ formalism presented above to both QA and SQA runs of the model. The key insight is that the ratio of exponents κ/ν appearing in the scaling relation Eq. (2.38) can be directly related to Δ. Just as in the 2D Ising model example discussed in the previous section, for the primary order parameter m in the 3D XY universality class, κ is typically referred to as β and can simply related to the anomalous dimension $\eta_{3\mathrm{D}XY}$, and thus Δ_ϕ, via the hyper-scaling relations [23],

$$2\beta/\nu = 1 + \eta_{3\mathrm{D}XY} = 2\Delta_\phi. \tag{2.42}$$

The KZ finite-size scaling relation then reads,

$$\langle |m|^2 \rangle = L^{-2\Delta_\phi} F\left(vL^{z+1/\nu}\right), \tag{2.43}$$

where $|m|^2$ is the squared magnitude of the complex order parameters. This can be generalized to the secondary order parameter as well, where the definition of κ in terms of a critical exponent is not as transparent,

$$\langle |\Psi|^2 \rangle = L^{-2\Delta_t} G\left(vL^{z+1/\nu}\right), \tag{2.44}$$

where G is a potentially different scaling function than F. There are some subtleties that arise due to the fact that $2\Delta_t > d = 3$ for this 2D quantum system, but the relation in Eq. (2.43) has been used to extract the dynamical exponent z in some QA experiments [5].

2.7.4 Simulated Annealing of the 3D Classical 8-State Clock Model

As touched on in Sec. 2.6.2, it is often insightful to consider classical analogs to quantum systems, particularly from a numerical standpoint. The algorithms used to study classical systems are nearly always much simpler and typically allow for more flexibility. So before considering the SQA of the quantum system, it may be useful to first consider SA of a model with the same ground state symmetries, hosting a phase transition in the same universality class: the classical 3D 8-state clock model, Eq. 2.27 with $q = 8$.

A thermal anneal of this model through the critical temperature T_c will share many properties with the quantum anneal of the FFTFIM through the critical transverse field Γ_c. While dictated by different dynamic exponents ($z = 1$ in the quantum case, while z is MC simulation dependent for the classical case), the evolution of the $\mathbb{Z}_8$ and $\mathbb{Z}_4$ order parameters as the annealing velocity is changed is expected to be similar for these two models. This is of particular interest, as the order parameter probability distribution is one of the few observables accessible in D-Wave's quantum annealers.

In Fig. 2·14, we show the evolution of the probability distributions for the two order parameters in the hard (i.e. $h_q = \infty$) clock model as the annealing time is increased (here, n_{anneal} is directly proportional to the annealing time). One can clearly see the $\mathbb{Z}_8/\mathbb{Z}_4$ order stabilize as the velocity is decreased.

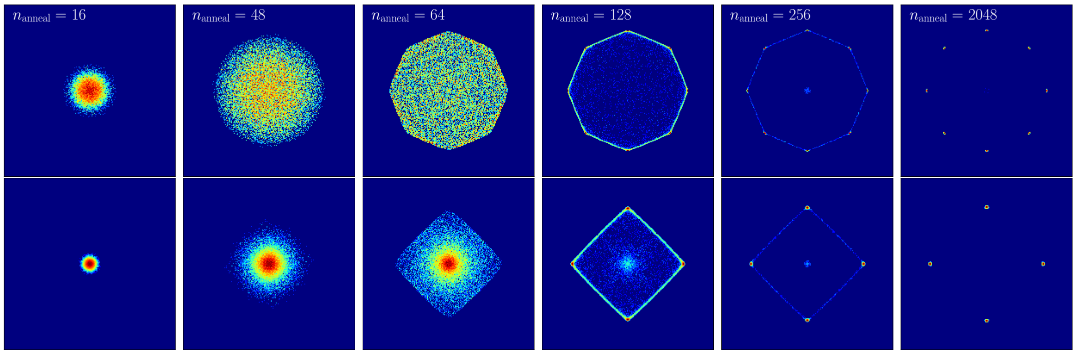

Figure 2·14: Order parameter distributions for increasing simulated annealing times in the 3D classical 8-state clock model. The top row shows the evolution of the $\mathbb{Z}_8$ order parameter distributions as the number of annealing steps (inversely proportional to the annealing velocity) is increased. The bottom row shows the corresponding $\mathbb{Z}_4$ order parameter distributions. This is for an $L = 12$ system that was annealed from an initial temperature of $2T_c$ to a final temperature of $0.5T_c$.

These results unexpectedly contrasted with the QA results for the FFTFIM (not included in this thesis, but shared via private communication with the authors of Ref. [5]). At short and intermediate times the distributions look comparable to those of the clock model (Gaussian about zero and $\mathbb{Z}_8/\mathbb{Z}_4$ with thin "tunnelling" lines at intermediate times, respectively), but at long times, the quantum system instead displays U(1) symmetry. The first expected culprit was noise, which is known to impact the results of longer anneals[72, 70].

Noise induced $U(1)$ symmetry

To determine if noise could induce the behavior seen in D-Wave's experiments, we implemented a noise protocol in our simulated annealing program that was aimed to mimic the behavior of D-Wave's Advantage quantum annealer. Noise effects are known to kick in at longer times, as thermal excitations emerge and decoherence begins to take effect [72]. Implementing a wide array of noise protocols is straightforward in the classical MC algorithm we use to simulate the 3D clock model; "noise steps" are simply added in between the local, Metropolis updates that evolve the classical spin configurations in the clean version of the model. To mimic the behavior observed by D-Wave, we implement a "noise-ramp," where after each MC sweep, the each spin is randomly rotated either clockwise or counterclockwise by one clock angle, with a probability P_{noise} that increases according to a simple protocol (i.e. as some power of the simulation time). In In Fig. 2·15, we show results for the evolution of the probability distributions of the $\mathbb{Z}_8$ order parameter as the annealing time is increased, but with noise that increases quartically with the simulation time.

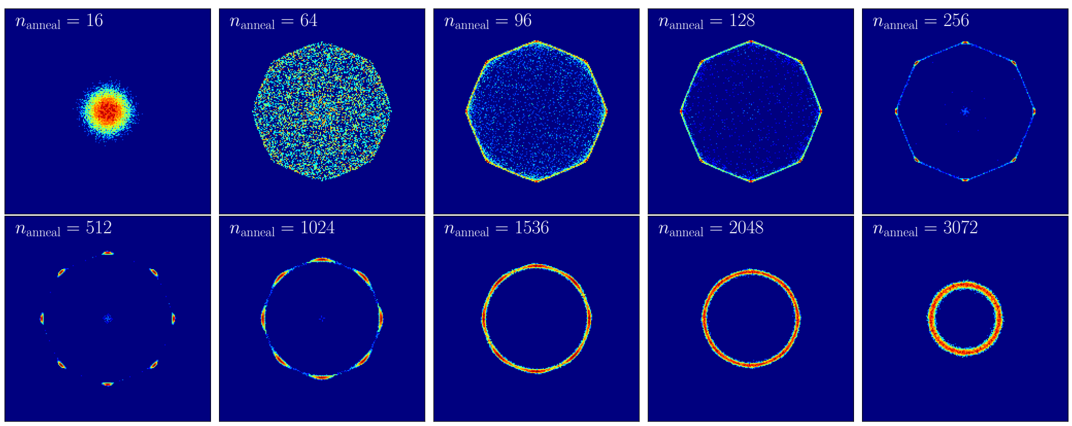

Figure 2·15: Order parameter distributions for increasing simulated annealing times in the noisy 3D classical 8-state clock model. This is for an $L = 12$ system that was annealed from an initial temperature of $2T_c$ to a final temperature of $0.5T_c$, but with a single spin rotation noise that increases in likelihood as the fourth power of the simulation time.

By including this simple noise protocol, the observed behavior is indeed replicated—as the total annealing time increases, the order isn't completely destroyed, but rather a $U(1)$ symmetry emerges. While this is of course a far cry from that is happening within D-Wave's quantum anneal, the point of this simple test was to see if noise has the ability to only partially destroy the $\mathbb{Z}_8$ order of the ground state. This test thus provides important insight into the real life QA of the FFTFIM.

2.8 Summary and Discussion

We have clarified the nature of the two order parameters, based on spins and dimers, in the square-lattice FFTFIM. The spin-based order parameter is primary, as it scales with the leading critical exponents and displays the full eight-fold degeneracy of the ordered phase, while the $\mathbb{Z}_4$ dimer order parameter is secondary with faster decaying critical correlations. Our QMC results at $T = 0$ and $T > 0$ confirm the scaling exponents predicted from the respective low-energy field theories.

Mapping to dimer models is a powerful tool in the study of frustrated spin systems, and the problem of lattice coverings by hard-core dimers is an interesting topic in itself. The connection between the secondary dimer order parameter and the primary spin order parameter then provides a crucial link between the models that had not been previously drawn in this context. Beyond the particular FFTFIM considered here, the triangular and kagome lattice AF Ising models [95, 91], the fully-frustrated honeycomb lattice Ising model [91], and the fully frustrated 4-8 lattice Ising models [49] are all well studied systems where the secondary order parameter prescription could also be applied.

Our insights also explain the critical scaling of a so-called "parasitic" order parameter studied in the AF 3-state Potts model on the diamond lattice [47, 48, 133]. While this system primarily orders antiferromagnetically in the ground state, it was shown that the presence of this order induces a finite ferromagnetic moment, which is captured by a secondary order parameter. The observed scaling of this order parameter had previously eluded explanation, but it is now clear that this "parasitic" order parameter also has scaling dimension Δ_t in that system.

The relevance of the secondary order parameter operator ($\Delta_t < 3$) in the FFT-

FIM implies that a perturbation favoring one of the dimer states, accomplished by appropriately modulating the Ising couplings, will induce a $\mathbb{Z}_2$ symmetry breaking phase in the plane of Γ and the dimer field (modulation) strength h_d. The phase boundary between the paramagnetic and ordered phases should have the asymptotic form $h_{d,c} \sim |\Gamma_c(h_{d,c}) - \Gamma_c(0)|^{\nu/\nu_d}$, where ν is the 3D O(2) correlation-length exponent and $\nu_d = (3 - \Delta_t)^{-1}$.

The FFTFIM, with $h_d = 0$ and $h_d > 0$, can be implemented on current D-Wave quantum annealing devices, as discussed in Sec. 2.7. The frustrated AFM triangular AF Ising model has already been studied in depth [68, 69], and more recently was the FFTFIM implemented on such a device [5]. In this recent study, both the spin and dimer order parameters were investigated, in light of the results of this work. The results were promising, but it appears that noise may impact the ability of slower anneals to truly remain adiabatic. We have investigated this phenomenon using a classical analog, the 3D 8-state clock model. While there are limitations to the mapping between this model and the full quantum treatment, we do gain insights on the observed emergence of U(1) symmetry when the quantum system is subject to a noisy annealing procedure. With the advancement of quantum annealing software, it will be fascinating to see how the results of D-Wave's QA hardware improve and the experiments they are able to run evolve.

Chapter 3

Stochastic Analytic Continuation—Principles and Implementation

In this chapter, I will introduce the numerical analytic continuation problem, present the principles behind the stochastic analytic continuation (SAC) method, our tool of choice for addressing this problem, and detail its implementation. The final section of this chapter will reference the GitHub repository where a user-friendly Julia implementation of the SAC method was developed as part of this thesis is available.

3.1 Preliminaries

We consider the imaginary-time correlation function of an operator O at inverse temperature $1/T = \beta$, evaluated at time points $\tau \in [0, \beta]$,

$$G(\tau) = \langle O^\dagger(\tau) O(0) \rangle. \tag{3.1}$$

Here, O can be either a bosonic operator, such as S^z, or a fermionic operator, such as $c^\dagger$, with only minor adjustments to the SAC procedure. In either case, the imaginary-time dependence is defined in the Heisenberg picture, taking $\hbar = 1$,

$$O(\tau) = e^{\tau H} O e^{-\tau H}. \tag{3.2}$$

The corresponding spectral function can be defined in terms of the eigenstates $|n\rangle$ and energy eigenvalues E_n of the Hamiltonian describing the system,

$$S(\omega) = \frac{\pi}{Z} \sum_{m,n} e^{-\beta E_n} |\langle m|O|n\rangle|^2 \delta\left(\omega - [E_m - E_n]\right), \tag{3.3}$$

where Z is the partition function. With this definition of the spectral function, the relationship between $S(\omega)$ and $G(\tau)$ is given by the equation

$$G(\tau) = \int_{-\infty}^{\infty} d\omega\, S(\omega) K(\tau, \omega), \tag{3.4}$$

where the kernel $K(\tau, \omega)$ is defined as

$$K(\tau, \omega) = \frac{1}{\pi} e^{-\tau\omega}. \tag{3.5}$$

Bosonic spectral functions defined using the convention in Eq. (3.3) obey the condition

$$S_B(-\omega) = e^{-\beta\omega} S_B(\omega), \tag{3.6}$$

which allows us to consider $S_B(\omega)$ on just the positive frequency axis when converting to $G(\tau)$:

$$G_B(\tau) = \int_{0}^{\infty} d\omega\, S_B(\omega) \tilde{K}_B(\tau, \omega), \tag{3.7}$$

with

$$\tilde{K}_B(\tau, \omega) = \frac{1}{\pi}\left(e^{-\tau\omega} + e^{-(\beta-\tau)\omega}\right). \tag{3.8}$$

The tilde will be dropped in the remainder of this chapter.

In some contexts it may instead be convenient to define the bosonic spectral function using a slightly different convention:

$$A_B(\omega) = \left(1 - e^{-\beta\omega}\right) S_B(\omega), \tag{3.9}$$

which instead satisfies the condition $A_B(-\omega) = -A_B(\omega)$. Under this convention, the so called Lehman representation of $A_B(\omega)$ (Eq. (3.3)) is modified to:

$$A_B(\omega) = \frac{\pi}{Z}\sum_{m,n}\left(e^{-\beta E_n} - e^{-\beta E_m}\right)|\langle m|O|n\rangle|^2\delta\left(\omega - [E_m - E_n]\right). \tag{3.10}$$

and the kernel now reads

$$K_B(\tau,\omega) = \frac{1}{\pi}\frac{e^{-\tau\omega} + e^{-(\beta-\tau)\omega}}{1 - e^{-\beta\omega}}. \tag{3.11}$$

In the fermionic case, one can define the spectral function as in Eq. (3.3) and the kernel as in Eq. (3.5), but it is the standard convention to use,

$$A_F(\omega) = \frac{\pi}{Z}\sum_{m,n}\left(e^{-\beta E_n} + e^{-\beta E_m}\right)|\langle m|O|n\rangle|^2\delta\left(\omega - [E_m - E_n]\right). \tag{3.12}$$

With $O = c^\dagger$, this definition has a very direct connection to the addition of a particle ($E_m > E_n$ and the first exponential dominates) or subtraction of a particle ($E_m < E_n$ and the second exponential dominates) from the system. The kernel relating $A_F(\omega)$ to the corresponding "single-particle Green's function" $G_F(\tau)$ is

$$K_F(\tau,\omega) = \frac{1}{\pi}\frac{e^{-\tau\omega}}{1 + e^{-\beta\omega}}. \tag{3.13}$$

Unlike in the bosonic case, there is no inherent relation between $A_F(-\omega)$ and $A_F(\omega)$. In some models, such as the Hubbard model at half-filling, the Hamiltonian is particle-hole symmetric, meaning that the addition of a particle can be mapped exactly to the addition of a hole (removal of particle) using a symmetry transformation. If this is the case, then $A_F(-\omega) = A_F(\omega)$ for select spectral functions (such as the density of states), but generally this cannot be assumed.

It should be emphasized that for both the bosonic and fermionic cases, the conventions used only determine the kernel used, and consequently, the precise definition of the spectral function. The frequency information (i.e. peak locations) will be similar, if not identical, no matter the kernel, but it is important to be aware of which formalism is being used when comparing results to those in the literature or to experiment. The SAC procedure, however, remains the same, so for consistency, in the remainder of this thesis we will use the spectral function defintion in Eq. (3.3) and the kernel in Eq. (3.8) when considering bosonic spectral functions and Eq. (3.12), Eq. (3.13) when considering fermionic spectral functions.

3.2 The Maximum Entropy Method

The goal of numerical analytic continuation is to the invert the integral in Eq. (3.4), solving for $S(\omega)$. However, this is what is considered an ill-posed problem [62]. The "ill-conditioned" kernel exponentially suppresses the information contained in $S(\omega)$, so given incomplete $G(\tau)$ data, i.e. limited imaginary-time resolution and/or the QMC-associated error, there is an infinite set of solutions for $S(\omega)$ that satisfy the integral.

The tradition tool to address this is the Maximum Entropy Method (MEM), which has produced many important results in the last 30 years [126, 125, 43, 62, 11]. This method is rooted in Bayes theorem and produces spectra that minimize the goodness-of-fit χ^2, while maximizing their entropy. Here, the definition of entropy depends on the way in which the spectrum is parameterized, but when $S(\omega)$ is treated as a continuous function of ω, a modified Shannon information entropy is typically what is used [62].

While considered one of the more reliable numerical analytic continuation methods, MEM has its limitations, being able only to reproduce broad spectral features, unless significant prior information about more complicated structure is provided to the algorithm. An alternative approach is to instead rely on the stochastic averaging of many solutions for the spectral function that fit the QMC-generated data well. This is the principle behind the SAC (or average spectrum) method [144, 113, 10, 134, 109, 130, 39, 107, 40, 41, 40, 42]. While SAC is still hampered by the over-favoring of smooth spectra, and, in fact, can be mapped directly to MEM in certain limits [10, 124], recent extensions of the method have shown that this shortcoming can in many cases be overcome. This is done by appropriately constraining the sampling space to favor sharp features, e.g., peaks and edges that often appear in ground-state spectral functions of quantum many-body systems [116, 123, 124, 119, 152, 120]. This represents the true advantage that SAC offers, which will be explored in detail in the follows sections.

3.3 SAC Procedure

I will now summarize the implementation of the SAC method and refer the reader to Ref. [124] for a more detailed and expansive explanation. In the basic formulation, the spectrum $S(\omega)$ is parametrized by a large number N_ω (typically 10^3 or larger) of δ-functions carrying weights a_i at energies ω_i, which can take continuous values (in practice values on a very dense frequency grid, $\Delta\omega \sim 10^{-5}$), as illustrated in Fig. 3·1(a) in the case of uniform amplitudes and Fig. 3·1(b) for variable amplitudes. Panels (c) - (e) correspond to "contrained" SAC parameterization, which will be discussed in the later sections.

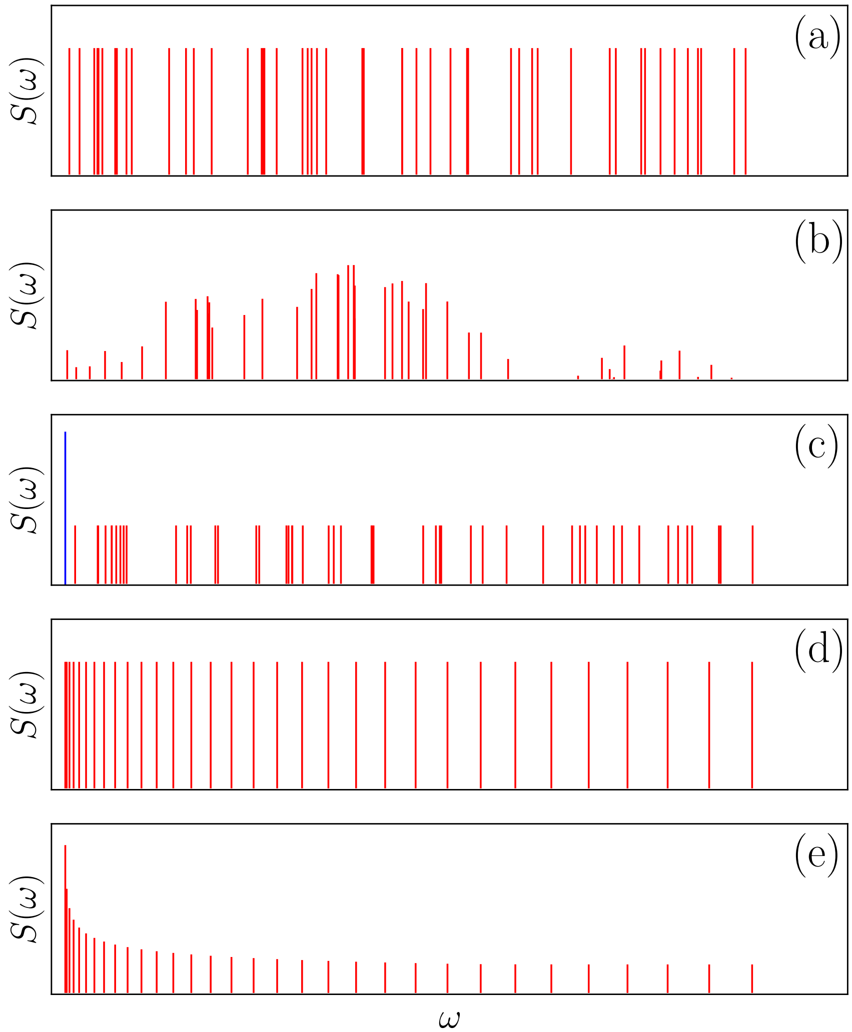

Figure 3·1: Depictions of the parameterizations of the spectrum $S(\omega)$: (a) Equal amplitudes and sampled frequencies in the continuum. (b) Sampled amplitudes and sampled frequencies in the continuum. (c) A "macroscopic" δ-function with amplitude A_0 at ω_0, followed by a large number N_ω of "continuum" δ-functions at positions $\omega_i > \omega_0$ with equal amplitudes $A_i = (1 - A_0)/N_\omega$. The amplitude A_0 is fixed in a sampling run, but the frequency ω_0 is sampled. (d) Equal amplitudes with monotonically increasing spacing $d_i = \omega_{i+1} - \omega_i$. Just as in (b), the lowest frequency ω_1 is sampled along with all other frequencies, with the constraint $d_{i+1} > d_i$. (e) Varying amplitudes with monotonically increasing spacing, as in (c). In all cases, the final spectrum is the mean amplitude density of the δ-functions, which is accumulated in a histogram during the sampling process.

The energies and, optionally, the amplitudes, are importance-sampled according to a Boltzmann-like probability distribution

$$P(S) \propto \exp[-\chi^2(S)/2\Theta] \tag{3.14}$$

and the resulting spectral configurations are accumulated in a histogram. With a large number of sampled δ-function's and a many ($\sim 10^4$) sampling steps, sufficiently smooth histograms can be produced, with bin spacing $\Delta\omega_h$, typically one or two orders of magnitude larger than $\Delta\omega$.

The goodness-of-fit χ^2 is calculated with respect to the QMC-generated Green's function $\bar{G}(\tau_i)$, $i = 1, \ldots, N_\tau$, with $\bar{G}$ denoting the average over N_B bins:

$$\bar{G}(\tau_i) = \frac{1}{N_B} \sum_{b=1}^{N_B} \bar{G}^b(\tau_i), \tag{3.15}$$

$\bar{G}^b(\tau_i)$ being the value from bin measurement b at the imaginary time point τ_i. Because the statistical errors of different τ points are correlated, the full characterization of the error, and thus the calculation of χ^2, requires the covariance matrix [62]. We use the bootstrap method to calculate the covariance matrix,

$$C_{ij} = \frac{1}{M} \sum_{m=1}^{M} (G^m(\tau_i) - \bar{G}(\tau_i))(G^m(\tau_j) - \bar{G}(\tau_j)), \tag{3.16}$$

where G^m represents the average of a bootstrap sample of N_B randomly chosen bins and M is the total number of bootstrap samples, typically much larger than N_B.

For a given SAC spectrum S,

$$S = \sum_{i}^{N_\omega} a_i \delta(\omega_i), \tag{3.17}$$

the corresponding $G(\tau)$ is calculated using Eq. (3.4), or the appropriate variant, and the goodness-of-fit is given by

$$\chi^2 = \sum_{i,j}^{N_\tau} \left(G(\tau_i) - \bar{G}(\tau_i)\right) C_{ij}^{-1} \left(G(\tau_j) - \bar{G}(\tau_j)\right). \tag{3.18}$$

Alternatively, one can compute χ^2 in the eigenbasis of the covariance matrix:

$$\chi^2 = \sum_{i=1}^{N_\tau} \left(\frac{G'(\tau_i) - \bar{G}'(\tau_i)}{\sqrt{\sigma_i}}\right)^2, \tag{3.19}$$

where σ_i are the eigenvalues of the covariance matrix,

$$\sigma = U^T C U, \tag{3.20}$$

and both $G(\tau_i)$ and $\bar{G}$ have now been transformed to this eigenbasis,

$$G' = U^T G, \quad \bar{G}' = U^T \bar{G}. \tag{3.21}$$

In practice, G is calculated directly in this basis by instead transforming the kernel once at the beginning of the simulation,

$$K'(\omega) = U^T K(\omega), \quad G(\tau)' = \int_{-\infty}^{\infty} d\omega S(\omega) K'(\tau,\omega) = \sum_i a_i K'(\tau,\omega_i) \tag{3.22}$$

It's been shown that the covariance may become unstable (its inverse may become noise dominated) if too many τ points are used, or equivalently, not enough QMC-generated bins are used [62, 124]. Empirically, it has been found that $N_b > 2N_\tau$ is typically sufficient to avoid such instabilities, but in some cases, even more, $N_b > 10N_\tau$, is necessary. One strategy is to generate a very large number of bins, composed of a relatively small number of QMC sweeps (but not too small, as to ensure that each bin is still statistically independent). The balance of statistical independence and number of bins can then be adjusted by re-binning the data, post-simulation.

However, it is the best practice to use a large enough number of QMC sweeps per bin, such that the autocorrelation time is far exceeded by the bin length (in number of sweeps) [115].

In the acceptance probability, Eq. (3.14), we have introduced a fictitious temperature Θ, which regulates the balance of entropy and goodness-of-fit. In the limit $\Theta \to 0$, the sampled spectrum purely minimizes χ^2, which will not reproduce the most probable spectrum when noisy data are used, while in the limit $\Theta \to \infty$, configurational entropy will dominant, and smooth, broad, and flat spectra will be favored. Purely Bayesian arguments imply that the temperature should be fixed to $\Theta = 1$ [144, 130]. However, for $S(\omega)$ parameterized by a large number of sampled degrees of freedom, entropy will ruin the goodness-of-fit at this fixed value of Θ [124]. Instead, it has been proposed to use a temperature determined by a simple criterion to give the optimal balance of entropy and goodness-of-fit, corresponding to the value where

$$\langle \chi^2(\Theta) \rangle = \chi^2_{\min} + a\sigma_{\chi^2}, \tag{3.23}$$

where $\sigma_{\chi^2} = \sqrt{2\chi^2_{\min}}$ is used as a proxy for the standard deviation of the χ^2 distribution that the value of $\langle \chi^2 \rangle$ follows, and a is an order-one number (typically $a = 1/2$ is used). This criterion, motivated by the properties of the χ^2 distribution, has been shown to produce reliable results, reproducing spectra generated from synthetic QMC $G(\tau)$ data with high fidelity [124]. This criterion has been further supported using cross validation, as will be explored in Chapter 5.

To determine $\chi^2_{\min}$, a simulated annealing type procedure is run, where Θ is gradually reduced until $\langle \chi^2(\Theta) \rangle$ (averaged over the sampled spectral configurations) has converged, as shown in Fig. 3·2. After this first anneal, another is performed, terminating when χ^2 satisfies the criterion in Eq. (3.23). At this stage, a much longer sampling process is carried out at the temperature where the anneal was terminated, and the accumulated spectrum is output to file.

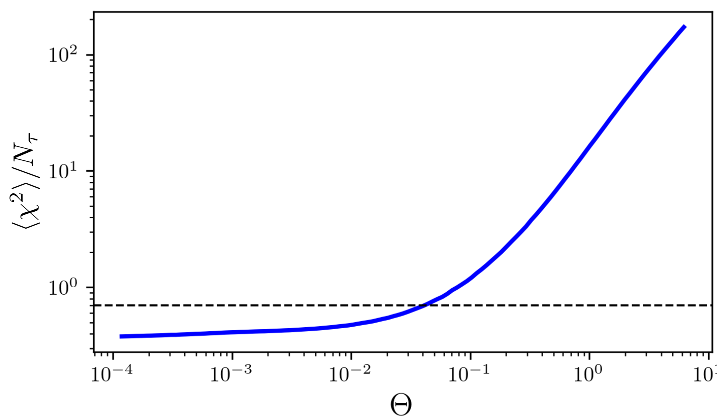

Figure 3·2: The average χ^2 versus Θ during a SAC annealing run. The black dashed line corresponds to the criterion in Eq. (3.23) with $a = 1/2$.

An illustration of this sampling process is depicted in Fig. 3·3, showing the evolution of the histogram as configurations are accumulated. The three rows are snapshots at three different points in the sampling process (first sampled configuration, midway through the sampling process, and the final sampled configuration). In the left column, we plot the individual δ-function configurations at each of these particular stages of the sampling process (only 40 of the 2,000 total sampled δ-functions are shown, for clarity). In the middle column, we plot the corresponding histograms, with bin widths denoted by the blue ticks on the ω axis. The bin width is much larger than what is typically used, which was chosen for illustrative clarity ($\Delta\omega_h = 0.1$). Using the equal amplitude parameterization, the histogram height corresponds to the density of δ-functions within each bin. When the amplitudes are sampled as well, the histogram instead represented a weighted density. The column on the right shows the accumulated histogram, which uses a much finer bin width than the middle column ($\Delta\omega_h = 0.005$). As more and more configurations are sampled, the accumulated histogram becomes smoother, as can be seen in the evolution from the top row to the bottom row (right column).

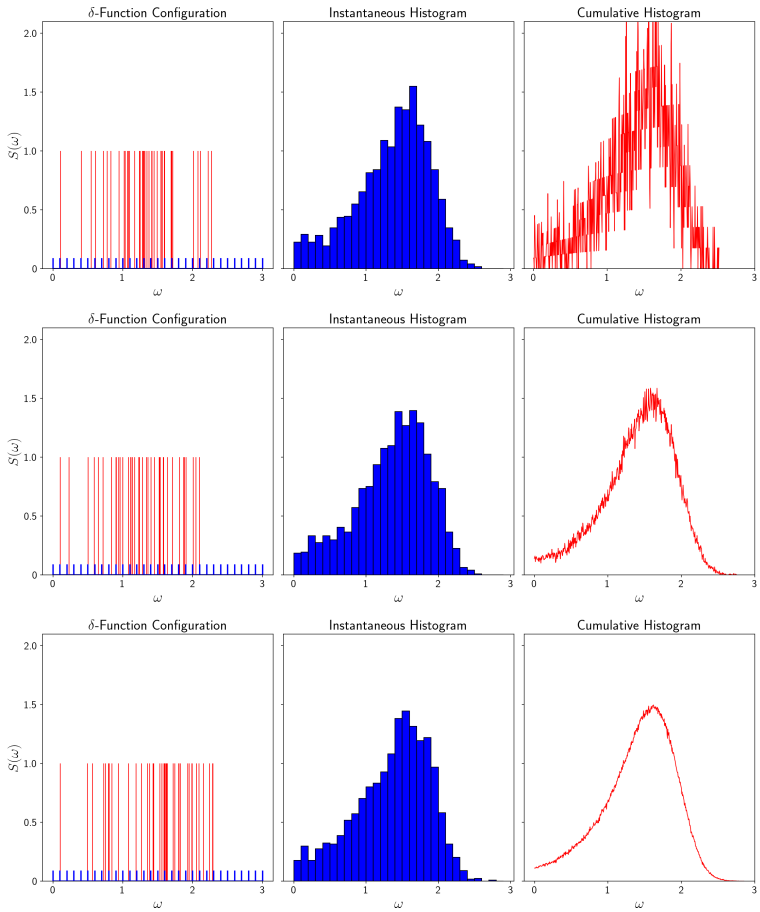

Figure 3·3: Demonstration of the SAC sampling process. Each row represents a different point in the sampling process at a fixed temperature Θ (first configuration, midway though, and final configuration, for the top, middle, and bottom, respectively). In the left column, 40 $\delta - function$ samples of the individual δ-function configurations are plotted. In the middle column, the corresponding histograms for these configurations are plotted, with a bin width denoted by the blue ticks on the ω-axis in the left column ($\Delta\omega_h = 0.1$. which is much larger than what is typically used). In the right column, the accumulated histogram throughout the course of the sampling process is plotted, which uses a much finer bin width than what is used in the middle column ($\Delta\omega_h = 0.005$).

3.4 Constrained Parameterizations

The true advantage that SAC has over other numerical analytic continuation methods is the recently developed class of "constrained" parameterizations. If certain restrictions are imposed on the locations and amplitudes of the δ-functions, the SAC method can reproduce sharp spectral features present only at low temperatures, such as narrow quasi-particle peaks and power-law edge singularities [116, 123, 124]. The key insight allowing for the resolution of such features is that imposing constraints can drastically change the entropic pressures felt within the sampling procedure, resulting in dramatic changes in the shape of the output spectra. The two constrained parameterizations discussed in this thesis are shown in Fig. 3·1(c) - (d), and will be described below.

3.4.1 Delta Peak Constrained Parameterization

A spectrum containing a "macroscopic" δ-peak followed by a smooth continuum can be reproduced using the parameterization depicted in Fig. 3·1(c) [123]

$$S_{\text{peak}}(\omega) = a_0\delta(\omega - \omega_0) + \sum_{i}^{N_\omega} a_i\delta(\omega - \omega_i), \tag{3.24}$$

where $a_0 + \sum_i a_i = 1$ and the constraint $\omega_i > \omega_0$ for all i is imposed during the sampling process. With this parameterization, the location of the peak, ω_0, is sampled within the program, along with all of the other "microscopic" δ-functions, but the weight of the macroscopic δ-peak, a_0, is fixed and optimized using a simple scan to find a χ^2 minimum, as detailed in the following section. Unless the optimal a_0 is very small, ω_0 fluctuates very little once it has equilibrated to its data-determined position.

This type of spectral function is appropriate for describing the spectrum of the $S = 1/2$ AFM Heisenberg model in two-dimensions (2D) [22], where linear spin-wave theory predicts the presence of a macroscopic, "single-magnon" peak [58], while an incoherent continuum at higher energy can be attributed to either multimagnon excitations or partially deconfined spinon excitations (both of which cannot be captured by convention spin-wave theory [123]). The SAC parameterization in Eq. (3.24) has been used to produce results for the spectral function of the aforementioned square-lattice Heisenberg model that agreed very well with results from inelastic neutron-scattering experiments on a material considered the best physical realization of this model [123].

Optimizing a_0

To determine the optimal value of a_0, we track how $\langle \chi^2 \rangle$ changes as a_0 is increased from zero at a fixed value of the sampling temperature Θ. We typically compare the location of the optimal a_0 at a few decreasing values of Θ. At a higher temperatures, where $\langle \chi^2 \rangle$ does not represent a statistically acceptable fit, the $\langle \chi^2 \rangle$ minimum will be more pronounced, flattening out as Θ is reduced. We determine the optimal value by tracking the location of the $\langle \chi^2 \rangle$ minimum as it approaches the value dictated by the simple criterion typically used to fix Θ when performing SAC [124]. The exact location of the minimum will typically depend on the sampling temperature to some extent, but in most cases this dependance is very minor. Furthermore, any value of a_0 chosen around the χ^2 minimum will produce nearly identical spectra, with perhaps slightly varied peak locations ω_0 (fluctuating within 1%). This variation allows us to determine an approximate error bar on ω_0.

An example of a scan over a_0 for the $S = 1/2$ AFM Heisenberg chain is shown in Fig. 3·4, for $q = \pi/2$, panel (a), and $q = 3\pi/4$, panel (b). While the 1D model is not expected to host spin-waves, and thus should not contain a dominate δ-function peak, its spectral function is expected to contain a sharp, power-law decaying edge. This parameterization may be able to reproduce this edge feature to some extent, particularly for finite-sized systems.

In this case, we use a value of Θ that is slightly elevated from that corresponding to the criterion Eq. (3.23) for an unrestricted sampling run performed before these scans, as discussed in Ref. [124]; at this temperature, a clear minimum is present, but the χ^2 value is still at an acceptable value. In more recent works, we have opted to instead choose a scan temperature such that the value of χ^2 at the minimum versus a_0 satisfies the criterion (with $\chi^2_{\rm min}$ determined by the delta peak scans, not an unrestricted sampling run). Again, it should be emphasized that this choice only marginally changes the optimal a_0, and therefore only marginally changes the shape of the output spectra, but may impact the exact location of the edge, which is important when extracting dispersion relations using SAC.

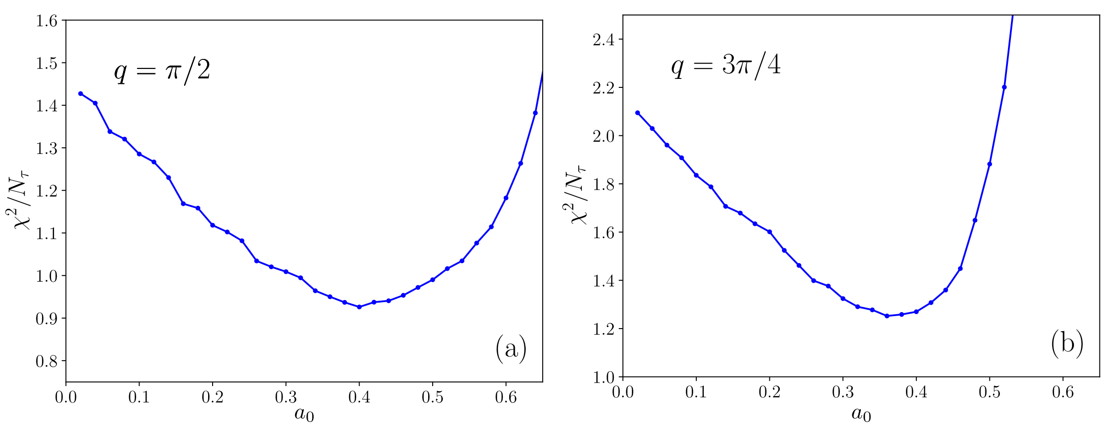

Figure 3·4: Results of a_0 scan for the Heisenberg chain. Panels (a) and (b) show the χ^2 value versus a_0 ($\Delta a_0 = 0.01$) for $q = \pi/2$ and $q = 3\pi/4$, respectively.

3.4.2 Power-Law Edge Constrained Parameterization

Another example of a sharp spectral feature that can only be resolved using a constrained SAC sampling parameterization is an edge singularity. In quantum many-body systems with fractionalized excitations, spectral weight may be spread over a range of energies, often with a power-law distribution that diverges at some frequency ω_q, where q is the total momentum of the fractionalized quasiparticle [124]. One of the most well know examples of this is the spectral function of the operator $O = S^z_q$ in the aforementioned $S = 1/2$ AFM Heisenberg chain. Using the Bethe ansatz (BA) solution, it was shown that in the thermodynamic limit, this spectral function diverges as $(\omega - \omega_q)^{-1/2}$, with a logarithmic correction (except for at $q = \pi$, where the divergence is faster) [24, 25, 99].

It was discovered that a very simple lower bound (i.e. minimum frequency) constraint on the sampling space of the δ-functions can produce a spectrum this with this type of feature [116]. This parameterization was later improved with an alternative constraint: the locations of the N_ω, sampled δ-functions are restricted such that the distance between adjacent δ-function's monotonically increases [124]:

$$S_{\text{edge}}(\omega) = \sum_i^{N_\omega} a_i \delta(\omega - \omega_i), \quad \omega_{i+1} - \omega_i > \omega_i - \omega_{i-1}, \tag{3.25}$$

as depicted in Fig. 3·1(d). Due to entropic pressure, this parameterization naturally produces a spectrum where the mean amplitude density of the δ-functions forms precisely an edge singularity with the asymptotic behavior $(\omega - \omega_0)^{-1/2}$ [124], where we have defined the edge location $\omega_0 \equiv \langle \omega_1 \rangle$. This edge location adapts naturally to the data during the sampling process, as does the shape away from the edge, smoothly transitioning from the asymptotic power-law form to an arbitrary monotonically decaying tail.

It is also possible to resolve an edge diverging with any power p, $(\omega - \omega_0)^p$ with $p < 0$, by properly adjusting the amplitudes a_i [124]. For an edge that diverges faster/slower than $S(\omega \to \omega_0) \propto (\omega - \omega_0)^{-1/2}$, i.e. p less/greater than $-1/2$, the amplitudes of the δ-functions increase/decrease as they approach the edge. This is depicted in Fig 3·1(e) for an edge that decays with an exponent $p = -3/4$. Just as with a_0 in the delta peak parametrization, p can be scanned over to determine an optimal power-law exponent. In practice, determining a precise optimal exponent is quite challenging, as in most cases, the power law divergence is an *asymptotic* feature, with the form $(\omega-\omega_0)^p$ only applying in a small region close to ω_0. An added complication is that the adjusted amplitudes a_i are typically allowed to crossover to an equal amplitude form, with the crossover point sampled within the program as well. If the asymptotic form of the divergence was an exponent p that is close to $-1/2$, then this crossover will naturally adjust to a very small value, such that the majority of the δ-functions will have a constant amplitude (and the resulting χ^2 versus p curve will thus be very flat). Nevertheless, a scanning procedure can still be performed, and if a clear χ^2 minimum versus p can be identified, this typically will represent an optimal power law divergence form. In Sec. 5.5.1, we will explore these subtleties, using QMC-generated data for the $S = 1/2$ AFM Heisenberg chain, as an example case where the expected exponent is known.

A note about the output spectra for this parameterization: a better alternative to collecting spectral weight in a histogram, the mean spectral density can in this case be defined for $i = 1, \ldots, N_\omega$ as

$$S(\omega_{i+1/2}) = \frac{1}{2}\frac{a_i + a_{i+1}}{\langle\omega_{i+1} - \omega_i\rangle}, \tag{3.26}$$

where $\omega_{i+1/2} = \langle\omega_i + \omega_{i+1}\rangle/2$ defines the self-generated grid upon which the spectral function is evaluated. This method avoids artificial oscillatory behavior that may appear in the tail of the spectrum [124].

As an example, in Fig. 3·5 we plot spectra produced by the unconstrained (a, d), delta peak (b, e) and power-law edge (c, f) parameterizations for the dynamic spin structure factor of the $S = 1/2$ AFM Heisenberg chain at $q = \pi/2$ and $q = 3\pi/4$. For each parameterization, the same QMC-generated data was used (for a system of size

$L = 512$ at inverse temperature $\beta \geq L$), yet the output spectra vary dramatically—a result of the different imposed constraints. Furthermore, each parameterization produced a results with an acceptable value of the goodness-of-fit, emphasizing the ill-posed nature of the numerical analytic continuation problem. In each panel, we also plot the numerical BA solution for a system of size $L = 500$, which has been subject to artificial smoothing, but otherwise is expected to represent the exact spectrum well [1]. While the delta peak spectrum can approximate the sharp edge feature fairly well (locating the true edge location within 1%), the power-law edge spectrum matches the BA result nearly exactly. The unconstrained spectrum, on the other hand, produces a result that is significantly broader, with substantial spectral weight within the energy gap.

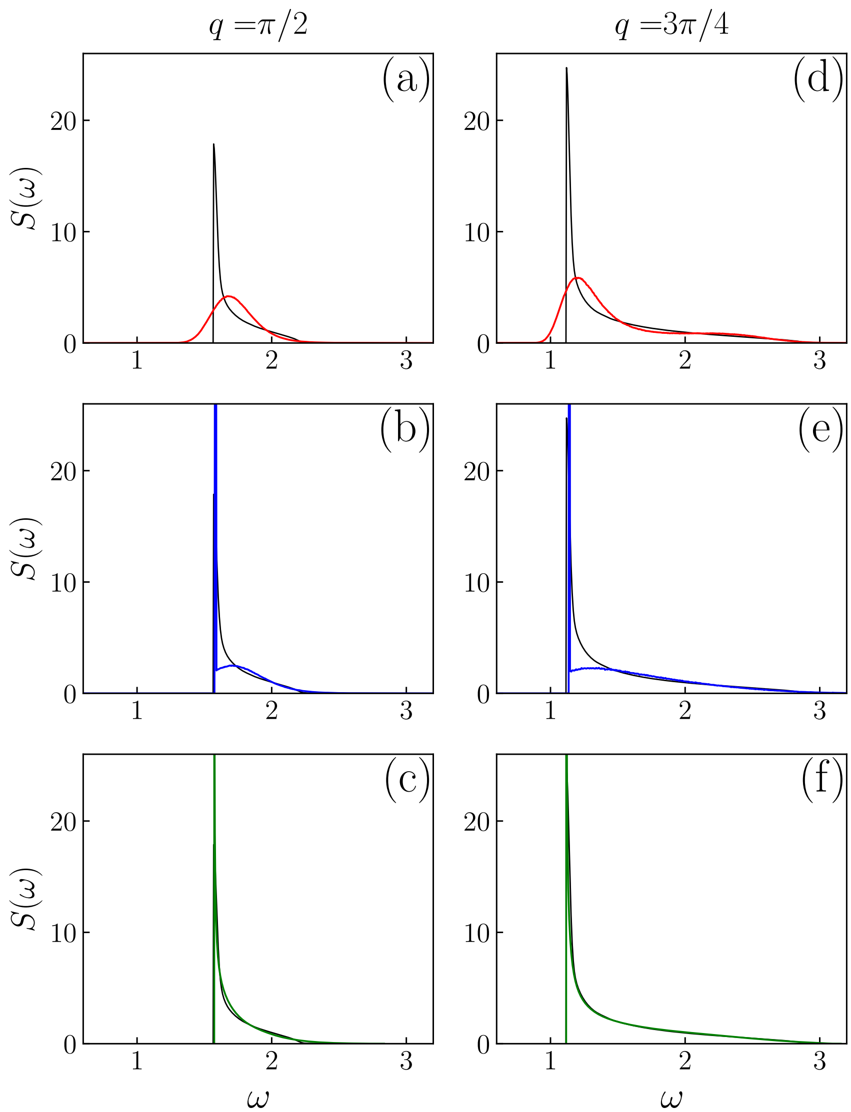

Figure 3·5: Spectra produced ny the unconstrained (a, d), the delta peak (b, e), and the power-law edge (c, f) parameterizations for the $S = 1/2$ AFM Heisenberg chain (QMC-generated data, system size $L = 512$). For each parameterization, the spectra are sampled at the temperatures corresponding to the criterion in Eq. (3.23), using their respective $\langle \chi^2 \rangle$ curves. For the unconstrained sampling parameterization, we use $N_\omega = 2,000$ δ-functions. For the single-peak parameterization we use leading δ-function weights a_0 determined by scans, as shown in Fig. 3·4, and $N_\omega = 2,000$ continuum δ-functions. For the edge parameterization, we use $N_\omega = 640$ and the exponents $p = -0.5$ and $p = -0.6$ for $q = \pi/2$ and $q = 3\pi/4$, respectively (as determined by scans). In each panel, we also plot the BA results for a chain of length 500 (black curves), which have also been subject to smoothing [1], resulting in more rounded peaks.

3.4.3 Combined Constrained Parameterizations

It is possible, and in general quite simple, to combine these constrained parameterizations with the standard unconstrained sampling scheme first introduced in Sec. 3.3, as well as with one another. In all cases, one represents the spectrum $S(\omega)$ as a sum of multiple contributions,

$$S(\omega) = a_A S_A(\omega) + a_B S_B(\omega), \tag{3.27}$$

with A and B corresponding to two different parameterizations, and $a_A + a_B = 1$. The relative weights, a_A and a_B, must be optimized using a scan, just as a_0 in the delta peak parameterization.

These two contributions to the spectrum, A and B, are sampled independently using the same updating schemes one would use in their standard prescriptions (see Ref. [124] for details), carrying along any parameterization specific constraints in S_A that may impact S_B (and vice versa). An example of this carry-over, we consider combining the power-law edge parameterization (S_A) and the unconstrained parameterization (S_B). In this scenario, the minimum frequency δ-function in S_B cannot be at a lower energy than the edge of S_A, ω_1:

$$\min\{\omega_i^B\} > \omega_1^A. \tag{3.28}$$

This can be easily encoded in the sampling process, with any update that violates this condition immediately being rejected. This combined spectrum can reproduce power law edges with non-monotonically decaying tails (i.e. peak structure(s) in the tail), as explored in Ref. [124].

Another example of a combined parameterization is a spectrum with two power-law edges, one at lower energy, decaying to the right and one at higher energy, decaying to the left, which is predicted to describe the single hole dynamics in the 1D

t-J model [17]. In this case, both spectra A and spectra B are parameterized as in Eq. (3.25), with the restriction that the rightward decaying spectrum, S_A, does not extend past the edge of the leftward decaying spectrum, S_B, (and vice versa):

$$\omega^B_{N_\omega} > \omega^A_1, \quad \omega^A_{N_\omega} < \omega^B_1 \tag{3.29}$$

An example of using this combined parameterization to resolve the single-hole spectral function at $k = \pi/2$ (i.e. $O = c_{k=\pi/2}$) in the 1D t-J model ($J/t = 0.4$) is shown in Fig. 3·6, with $a_A = a_B = 0.5$ [153]. This ratio was determined via a scan, which displayed a broad minimum between $a_A = 0.4$ and 0.6, representing the inherent flexibility of the combined parameterization. Since the tails of each contribution extend to the edges of their counterparts, a range of relative weight distributions may produce fairly similar spectra.

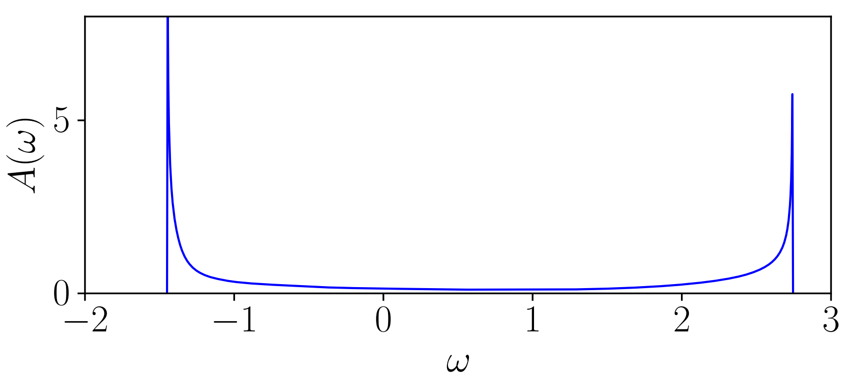

Figure 3·6: Double power law edge spectrum for the single-hole spectral function at $k = \pi/2$ in the 1D t-J model with $J/t = 0.4$. The relative weight ratio of the two contributions is $a_A = a_B = 0.5$

3.4.4 A Note on Fermionic Spectral Functions

In general, the detailed balance condition in Eq. (3.6) makes resolving bosonic spectral functions using constrained SAC much easier than resolving fermionic ones. This is because the sharp edge features require that the ratio of the total spectral weight that lies on the positive and negative frequency axes must either be fixed and optimized using a scan, or adjusted as part of the updating process (as opposed to the case where no sharp edges are present, and spectral weight can be distributed and sampled freely over the entire frequency axis). In practice, it is quite difficult to shift spectral weight between the two sides of the frequency axis after the spectrum has adequately equilibrated, so careful initialization of the δ-function configurations is necessary. Once can use basic information stored in $\bar{G}(\tau)$ data, such as the ratio of $\bar{G}(\tau \sim 0)$ to $\bar{G}(\tau \sim \beta)$, to roughly approximate the proper weight distribution, but this is not a failsafe approach. We note, however, that this difficulty to transfer weight is a reflection of the stability of the relative weights once equilibrium has been achieved, rather than an inherent inefficiency of the sampling procedure.

While it is still possible to produce fermionic spectral functions that are asymmetric about $\omega = 0$, i.e. two δ-function peaks with different weights $a_0^{\pm}$ and locations $\omega_0^{\pm}$, it is substantially easier to resolve symmetric fermionic spectra. If symmetry is inherent in the model, then one only has to sample on the positive frequency axis, using a modified kernel,

$$K_F^{\mathrm{symm}}(\tau, \omega) = \frac{1}{\pi} \left(\frac{e^{-\tau\omega}}{1 + e^{-\beta\omega}} + \frac{e^{\tau\omega}}{1 + e^{\beta\omega}} \right). \tag{3.30}$$

Examples of both symmetric and asymmetric constrained SAC spectra will be explored in detail in Chapter 4.

3.5 GitHub Repository

A comprehensive Julia SAC library can be found in the GitHub repository `https://github.com/gabes135/sac`. This library contains code to run SAC using all of parameterizations discussed in this thesis. It also contains code that can be used to prepare QMC-generated data to be read in by the SAC program (i.e. calculation of the covariance matrix), as well as programs that can be used to generate synthetic data for testing purposes.

The library supports both bosonic or fermionic spectral functions, fermionic spectra with and without the imposed symmetry $A(-\omega) = A(\omega)$, and various combined constrained parameterizations, as detailed in the README files in the repository.

The library also contains a utility program that can be used to run cross validation for model selection, the subject of Chapter 5.

Chapter 4

Single-Particle Dispersion and Density of States of the Half-Filled 2D Hubbard Model

4.1 Introduction

The Hubbard model [56, 66, 44] serves as the simplest and most essential model for the physics of correlated electrons, with the two-dimensional (2D) square-lattice case being of particular interest in the context of the unresolved puzzle of superconductivity in the cuprates [80, 37, 103, 158]. We focus here on half filling, where auxiliary-field quantum Monte Carlo (AFQMC) simulations can access the ground state of relatively large systems. Significant progress has been made on static observables [53, 54, 146, 147, 94, 93, 136, 106], but accessing dynamics, e.g., the important single-particle spectral function, is much more challenging [105]. While AFQMC simulations [36, 157] can be employed to calculate the wave-vector ($\mathbf{k}$) resolved imaginary-time Green's function [53, 145, 94, 117, 18, 140], the corresponding real frequency spectral function $A_{\mathbf{k}}(\omega)$ has been difficult to extract because of the ill-posed analytic continuation problem. Though some key aspects of the dispersion relation have been obtained [18, 102, 7, 8], significant uncertainties remain, and a precise characterization is still lacking. Within the rigid band approximation (RBA) [128, 33, 32], detailed knowledge of $A_{\mathbf{k}}(\omega)$ is required to understand the manner in which the quasiparticle states of the half-filled system are occupied upon doping and how important scattering channels emerge. This is essential for a coherent understanding of the model and its connection to high-temperature superconductivity.

The specific technical challenge of computing $A_{\mathbf{k}}(\omega)$ is that narrow quasiparticle peaks and associated sharp features in the density of states (DOS) cannot be reproduced by analytic continuation of AFQMC data with the conventional maximum-entropy method (MEM) [126, 125, 43, 62, 11] or the related stochastic analytic continuation (SAC) (or average spectrum) method [144, 113, 10, 134, 109, 130, 39, 107, 40, 41, 40, 42]. However, recent extensions of SAC have shown that this shortcoming can in many cases be overcome by appropriately constraining the sampling space to favor sharp features, e.g., peaks and edges that often appear in ground-state spectral functions of quantum many-body systems [116, 123, 124, 119, 152].

In this chapter, we implement constrained SAC to extract the dispersion relation and the DOS of the half-filled Hubbard model with sufficient precision to uncover features of key significance for understanding the emergence and initial evolution of the Fermi sea upon doping. Focusing on intermediate values $U/t = 4, 6$ of the Hubbard repulsion, we find a clear separation between the lowest single-particle energy ω_Σ at the four equivalent wave-vectors $\mathbf{k} = (\pm\pi/2, \pm\pi/2)$ (the Σ points) and almost dispersionless excitations at $\omega \approx \omega_{\mathrm{X}}$ close to $\mathbf{k} = (\pm\pi, 0), (0, \pm\pi)$ (the X points). With the dispersion being quadratic for $\mathbf{k}$ close to the Σ points, there is a ledge of almost constant DOS in the range $\omega \in [\omega_\Sigma, \omega_{\mathrm{X}})$, followed by a sharp edge with inverse square-root divergence at $\omega = \omega_{\mathrm{X}}$ from the almost flat (quartic) dispersion around the X points. We are able to determine the DOS to high precision and extract the total fraction n_{ledge} of quasiparticle states in the ledge, with $n_{\mathrm{ledge}} = 0.13$ and $n_{\mathrm{ledge}} = 0.20$ for $U/t = 4$ and $U/t = 6$, respectively.

Interpreting our results within the RBA, the density of quasiparticle states below ω_{X} controls the critical doping fraction $x = x_c = n_{\mathrm{ledge}}$ for the initial occupation of states around the X points. With the large DOS at $\omega \approx \omega_{\mathrm{X}}$ and the onset of (π, π) scattering processes connecting the four "hot spots" [131, 38, 9, 142, 103], this doping fraction should signify an instability, with superconductivity being a natural candidate.

4.2 Model and Methods

In standard notation, the Hamiltonian for the half-filled Hubbard model is

$$H = -t \sum_{\langle i,j \rangle, \sigma} \hat{c}^{\dagger}_{i,\sigma} \hat{c}_{j,\sigma} + U \sum_{i} (\hat{n}_{i,\uparrow} - \tfrac{1}{2})(\hat{n}_{i,\downarrow} - \tfrac{1}{2}), \tag{4.1}$$

here on the periodic square lattice with $N = L^2$ sites. We set $t = 1$ and use the AFQMC method to compute the imaginary-time dependent Green's function,

$$G_{\mathbf{k}}(\tau) = \langle T_{\tau} c_{\mathbf{k}}(\tau) c^{\dagger}_{\mathbf{k}}(0) \rangle, \tag{4.2}$$

at temperatures T low enough to converge to the ground state. We have confirmed that $\beta = 1/T = 2L$ is sufficient for all practical purposes for the repulsion strengths and system sizes $L \leq 20$ used here. In the AFQMC simulations, we use a time slice $\Delta\tau = 0.1$ and compute $G_{\mathbf{k}}(\tau)$ on this grid. The discretization error scales as Δ_{τ}^{2} and is insignificant, which we have confirmed using spot-checks of $\Delta\tau < 0.1$ for some of the smaller lattices. For the cases we study here (i.e., half-filling), AFQMC is sign-free [83] and numerically exact.

The single-particle spectral function $A_{\mathbf{k}}(\omega)$ is related to $G_{\mathbf{k}}(\tau)$ via the transform

$$G_{\mathbf{k}}(\tau) = \int_{-\infty}^{\infty} d\omega \, \frac{e^{-\tau\omega}}{1 + e^{-\beta\omega}} \, A_{\mathbf{k}}(\omega), \tag{4.3}$$

which we invert for $A_{\mathbf{k}}(\omega)$ using the SAC method. The DOS, $D(\omega)$, is the average of

$A_{\mathbf{k}}(\omega)$ over $\mathbf{k}$, which in imaginary time corresponds to

$$\frac{1}{N}\sum_{k} G_{\mathbf{k}}(\tau) = \frac{1}{N}\sum_{\mathbf{k},\mathbf{r}} G_{\mathbf{r}}(\tau)e^{i\mathbf{k}\cdot\mathbf{r}} \equiv G_{\text{loc}}(\tau), \tag{4.4}$$

where $G_{\text{loc}}(\tau) = G_{\mathbf{r}=\mathbf{0}}(\tau)$ is the local Green's function. Particle-hole symmetry at half-filling implies $D(-\omega) = D(\omega)$, which allows us to preform the analytic continuation of $G_{\text{loc}}(\tau)$ on only the positive frequency axis, by implementing a modified kernel in Eq. (4.3):

$$G_{\text{loc}}(\tau) = \int_{0}^{\infty} d\omega \left(\frac{e^{-\tau\omega}}{1+e^{-\beta\omega}} + \frac{e^{\tau\omega}}{1+e^{\beta\omega}} \right) D(\omega). \tag{4.5}$$

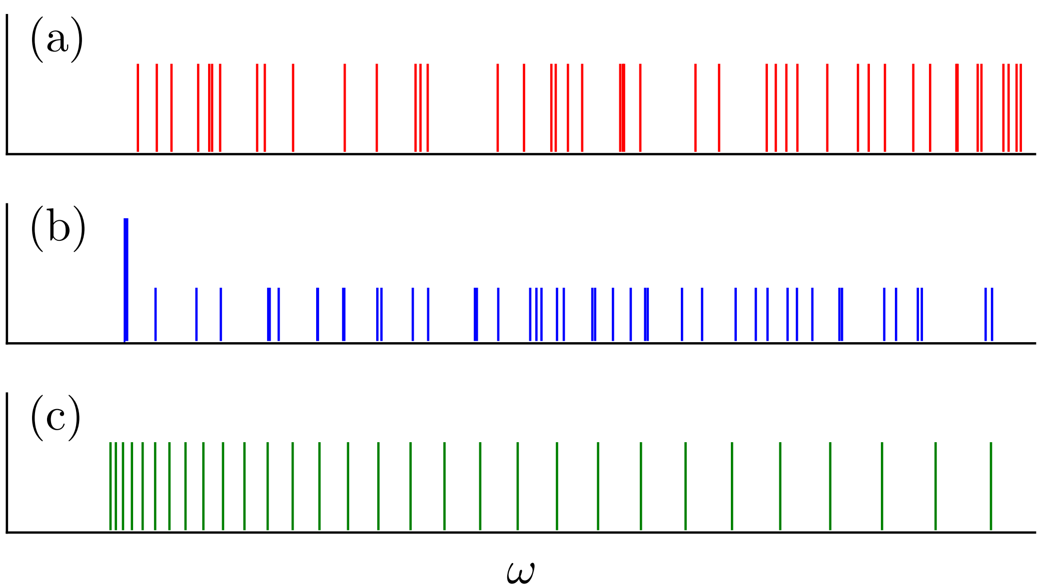

Figure 4·1: Schematic depictions of the $\omega > 0$ part of SAC sampling spaces, with the spectrum parametrized as a sum of δ-function with (a) unconstrained frequencies and fixed amplitudes, (b) a macroscopic δ-function (quasi particle) at ω_0, which also acts as a lower bound for the other (continuum) contributions, and (c) constrained so that the distance between adjacent δ-functions is monotonically increasing. This constraint produces an average spectrum with a sharp (divergent for $\omega \to \omega_0^+$ when $N_\omega \to \infty$) edge followed by an monotonically decaying continuum.

Fig. 4·1 depicts the three SAC parameterizations implemented to study this model, unconstrained (a), delta peak (b), and power-law edge (c). Details on these parameterizations can be found in Chapter 3. In the the particle-hole symmetric half-filled Hubbard model, $A_{\mathbf{k}}(\omega)$ will contain two quasiparticle peaks, at $\omega_0 = \pm|\omega_{\mathbf{k}}|$, which define the dispersion relations for injected holes and particles. In this case, we implement a modified delta peak parameterization, in which the optimized a_0 represents the sum of the two weights and the relative distribution between positive and negative part is sampled; see Sec. 3.4.4 and Sec. 4.3.1 for details on this optimization process. This constrained parametrization is suitable under the assumption of the true quasiparticle peak being very narrow, which can be expected here at least close to the minimum $|\omega_{\mathbf{k}}|$.

To resolve the DOS, we will use an extended form of the power-law edge parametrization for the DOS, touched on briefly in Sec. 3.4.3 and expanded upon in Sec. 4.4. In this modified paramterization, a continuum of unrestricted δ-functions is included to resolve the aforementioned ledge of the DOS.

4.3 Single-Particle Spectral Function

We first examine $A_{\mathbf{k}}(\omega)$ with $\mathbf{k}$ along lines of high symmetry in the Brillouin zone (BZ). Results for the $L = 16$ system at $U/t = 4$ obtained with both unconstrained sampling and the δ-edge constraint are presented in Fig. 4·2(a), with the optimal combined macroscopic δ-function weight a_0 displayed in each panel. The scanning procedure used to determine these optimal values in detailed below, in Sec. 4.3.1. The spectra from unconstrained SAC are qualitatively very similar to previous results obtained with the MEM [18, 102, 7, 8]. At the lowest energies, the peaks are sufficiently narrow for their centers to coincide with the location of the δ-edge of the constrained spectrum. Correspondingly, the continuum weights of these constrained spectra are relatively small. However, at the higher energies the peaks are too broad to provide a reliable dispersion relation, due to the dominant continuum beyond the quasiparticle peak.

The dispersion relation corresponding to the δ-edge locations is shown along three lines in the BZ in Fig. 4·2(b). Unlike the noninteracting system, the Σ and X states are no longer degenerate, with the former being the lowest in energy. We have carried out these calculations for many system sizes for both $U/t = 4$ and $U/t = 6$ and show the size dependence of ω_Σ and ω_X in Sec. 4.3.2. There is no sign of the energy difference vanishing as $L \to \infty$, which was suggested previously [8]. The extrapolated difference $\omega_\mathrm{X} - \omega_\Sigma$ is about 0.05 for $U/t = 4$ and 0.1 for $U/t = 6$.

We next consider the functional form of the dispersion about the Σ and X points. It was previously argued that the dispersion is quartic around X, supporting a metal-insulator transition with dynamic exponent $z = 4$ [6, 59, 7, 8, 87]. Our data can also be very well fitted to a quartic form, as shown in Fig. 4·2(b), except on the line connecting Σ and X (i.e. along the noninteracting Fermi surface). Examining lines extending from the X point at other angles, we find that the energy drops below ω_X only along the noninteracting Fermi surface, while elsewhere the dispersion is asymptotically quartic above ω_X. The density of states from the neighborhood of the X points is therefore divergent, of the form $D(\omega \to \omega_\mathrm{X}^+) \propto (\omega - \omega_\mathrm{X})^{-1/2}$. In the narrow range $\omega \in [\omega_\Sigma, \omega_\mathrm{X})$, the DOS should be roughly constant, on account of the quadratic dispersion around the Σ points; see Fig. 4·2(b) and Sec. 4.3.3 for results along different cuts in $\mathbf{k}$-space.

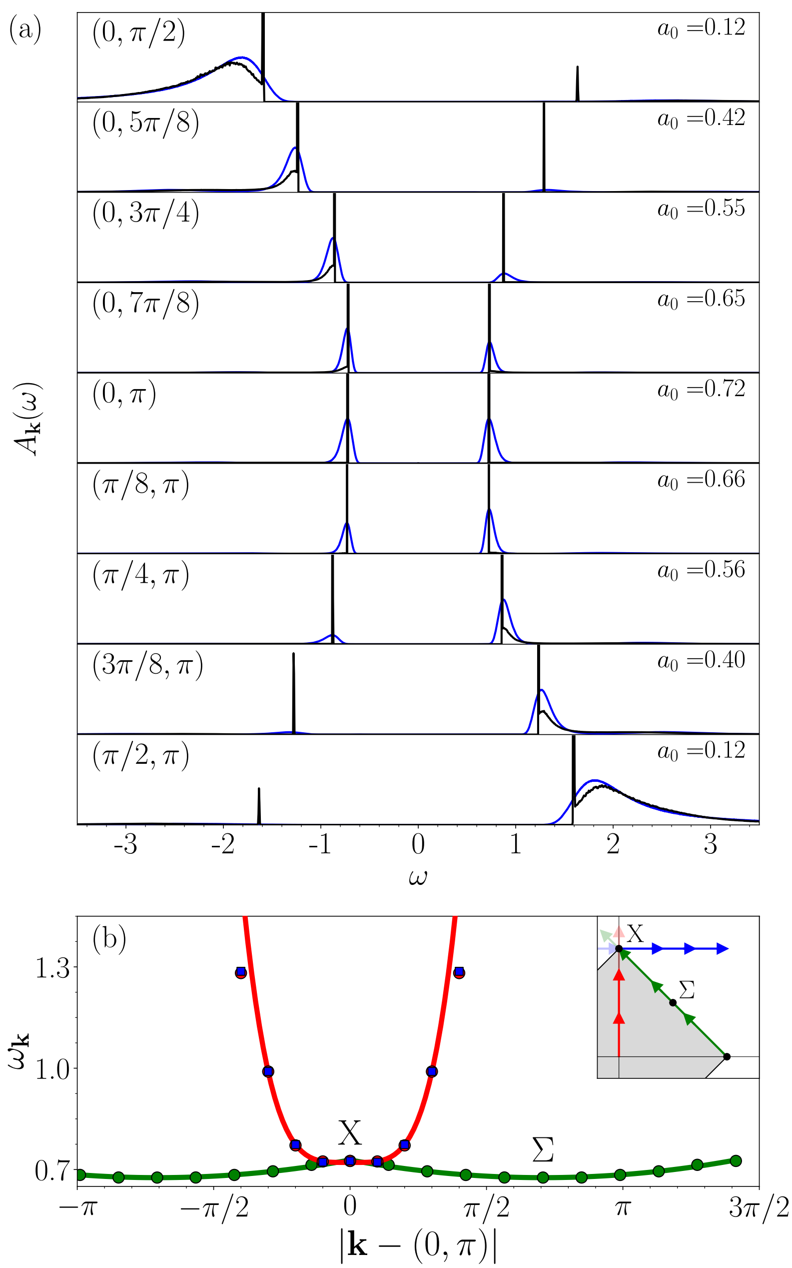

Figure 4·2: (a) $A_{\mathbf{k}}(\omega)$ for a series of $\mathbf{k}$ points in an $L = 16$ system with $U/t = 4$, comparing results of unconstrained (blue) and δ-edge constrained (black) parameterizations. The optimal macroscopic quasi-particle weights a_0 are indicated in each case. Both sets of spectra have been rescaled so that the maximum values of the unconstrained spectra are normalized to unity. (b) The dispersion relation obtained from the δ-edge locations, where the different colors correspond to the $\mathbf{k}$-space cuts depicted in the inset. The red and blue symbols coincide. The green and red curves are fits of the form $\omega_\Sigma + aq^2$ and $\omega_\mathrm{X} + bq^4$, respectively, with q being the distance to the respective reference points.

4.3.1 Optimization of a_0 for the Single-Particle Spectral Function

Here, we present exemplary results for the procedure used to determine the optimal value of the macroscopic δ-function weight a_0 using a scan, as described in detail in Sec. 3.4.1.

In Fig. 4·3 we show the scans used to determine a_0 for spectra between the Σ and X points in Fig. 4·2, In the left columns, $\langle \chi^2 \rangle$ is plotted versus a_0, and in the right column the corresponding location of this edge is shown. In the cases where the location of the $\langle \chi^2 \rangle$ minima drifts as the Θ is lowered, the edge location (colored points in the right panels) shifts only moderately. The minimum flattens as Θ is lowered and becomes hard to discern below the value corresponding to the optimal sampling temperature, indicated with dashed lines. We note that the minimum is the sharpest for the largest a_0 values; for the **k**-points closest to the the non-interacting Fermi surface. These BZ points are also the most important ones for the purposes of this work.

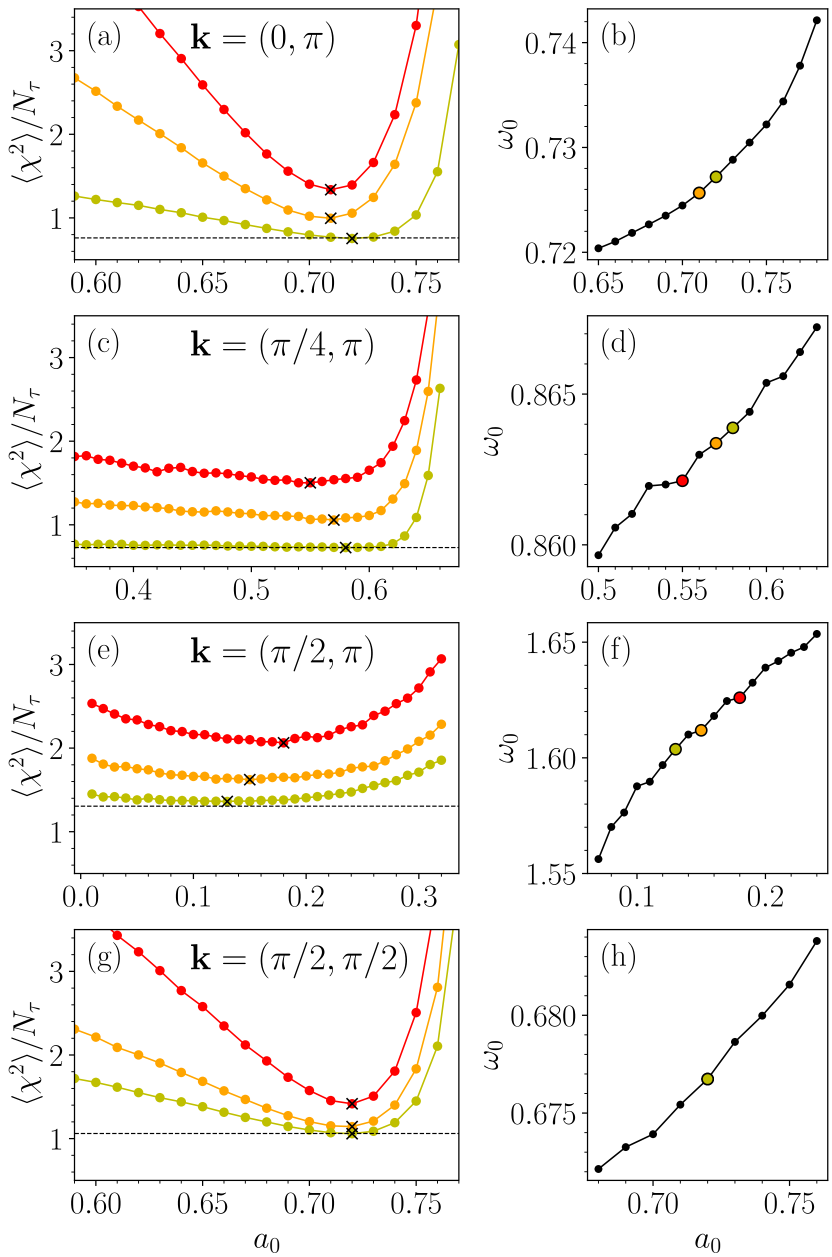

Figure 4·3: Scans over the δ weight a_0 in $A_{\mathbf{k}}(\omega)$ at four select momenta ($L = 16$, $U/t = 4$). The three colors correspond to three gradually lowered sampling temperatures (red to orange to yellow), where for the lowest Θ value, the χ^2 minimum coincides with the simple criterion typically used to fix Θ when performing SAC (black dashed lines).

4.3.2 Finite-Size Analysis of the Quasiparticle Energies

To monitor finite-size effects of the Σ and X excitation energies, we analyze the scaling of ω_{X} and ω_{Σ} as a function of system size. As shown in Fig. 4·4, both energies converge with L, with the difference between the $L = 16$ and $L = 20$ values being $<1\%$ for both U/t values. In light of the system being gapped with a finite correlation length, exponentially fast convergence can be expected, which is confirmed by the fits in Fig. 4·4.

For each system size, we estimated the uncertainty in the energy values by monitoring how the location of the macroscopic δ-function edge changes as a_0 is slightly increased and decreased from its optimal value, as depicted in the right column of Fig. 4·3. In all cases, the uncertainties are $<10\%$ of the magnitude of the energy difference, demonstrating that the splitting of these energy levels is neither a finite-size effect nor a consequence of SAC-associated resolution limitations.

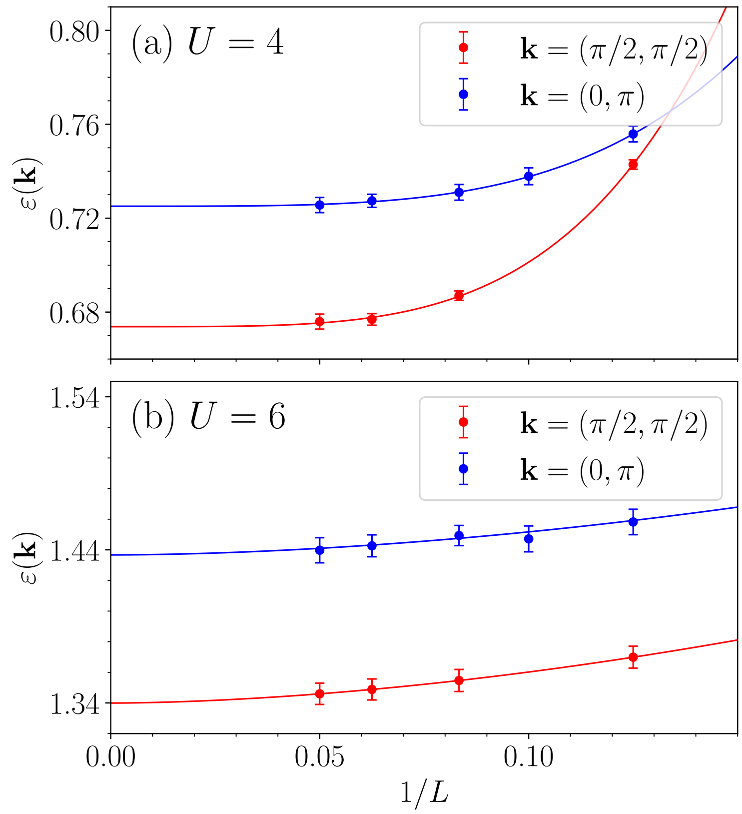

Figure 4·4: ω_X (blue) and ω_Σ (red) as a function of system size for $U/t = 4$, panel (a), and $U/t = 6$, panel (b). The solid lines are power law fits $\omega(L) = a + bL^{-c}$.

4.3.3 Anisotropic Quadratic Dispersion

The quartic dispersion about the X point is asymptotically highly isotropic, except along the noninteracting Fermi Surface where the energy increases on approach to X. In contrast, the quadratic dispersion about the Σ point, graphed in Fig. 4·5, is very anisotropic. The energy increases much more rapidly along the cut perpendicular to the noninteracting Fermi surface (purple points in Fig. 4·5), but are still well fit by a quadratic function. We note than an anisotropic quadratic dispersion still corresponds to a constant DOS in the narrow range $\omega \in [\omega_\Sigma, \omega_X)$.

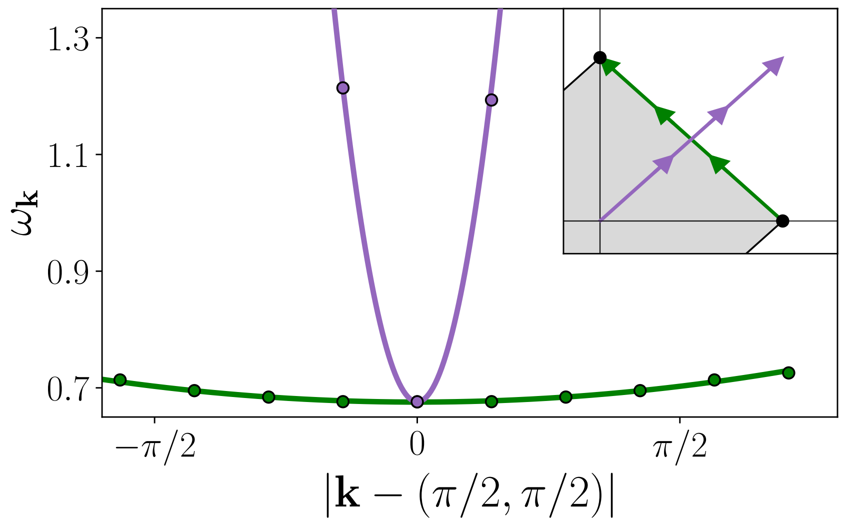

Figure 4·5: Quasiparticle dispersion along (green) and perpendicular to (purple) the non-interacting Fermi surface as a function of the distance in $\mathbf{k}$-space to Σ. The green points and fit are the same as those shown in Fig. 4·2(b), which are for $L = 20$ and $U/t = 4$.

4.4 Full Density of States

Our results for the dispersion relation can now be fed in as prior information in SAC to extract the full DOS from $G_{\text{loc}}(\tau)$. We know the singular points ω_Σ and ω_{X} and the smooth behavior that is expected between them, so we can use the parametrization in Fig. 4·1(a) to resolve the DOS in this region. For $\omega \geq \omega_{\text{X}}$, we use the constrained parametrization in Fig. 4·1(c), with the lowest of the δ-functions locked at ω_{X}. To determine the fraction ρ_{ledge} of states (i.e. the spectral weight) in the ledge of the DOS in $[\omega_\Sigma, \omega_{\text{X}})$, a scan is performed over this parameter to locate a goodness-of-fit minimum, in analogy to the scan over a_0 of the δ-edge discussed in Sec. 3.4.1. Results are shown in Fig. 4·6 along with $\langle\chi^2\rangle$ versus ρ_{ledge}. For reference, we also include results obtained by unconstrained sampling of the entire DOS, which cannot resolve the singularities and look very similar to previous results generating using MEM [18, 7].

Based on the rapid size convergence of ω_Σ, ω_X, and the DOS (using the L-specific values of ω_Σ and ω_X), we believe that the results in Fig. 4·6 well represent the thermodynamic limit. For both U/t values, the DOS between the two singular points is indeed very flat, motivating our designation of this part as the ledge. The singular peak above ω_X is followed by a thick tail with significant weight all the way up to $\omega = 6 \sim 7$. Given that the χ^2 values at the minimum of the scans are statistically acceptable (with the sampling temperature at its optimal value [124, 119]), there is no statistical evidence for any additional peaks beyond the edge at ω_X. Such additional peaks, often referred to as "ringing", are common in MEM results and, as seen in Fig. 4·6, are also produced by unconstrained SAC. The ringing behavior has been explained as a compensating behavior stemming from the presence of spectral weight inside the true gap.

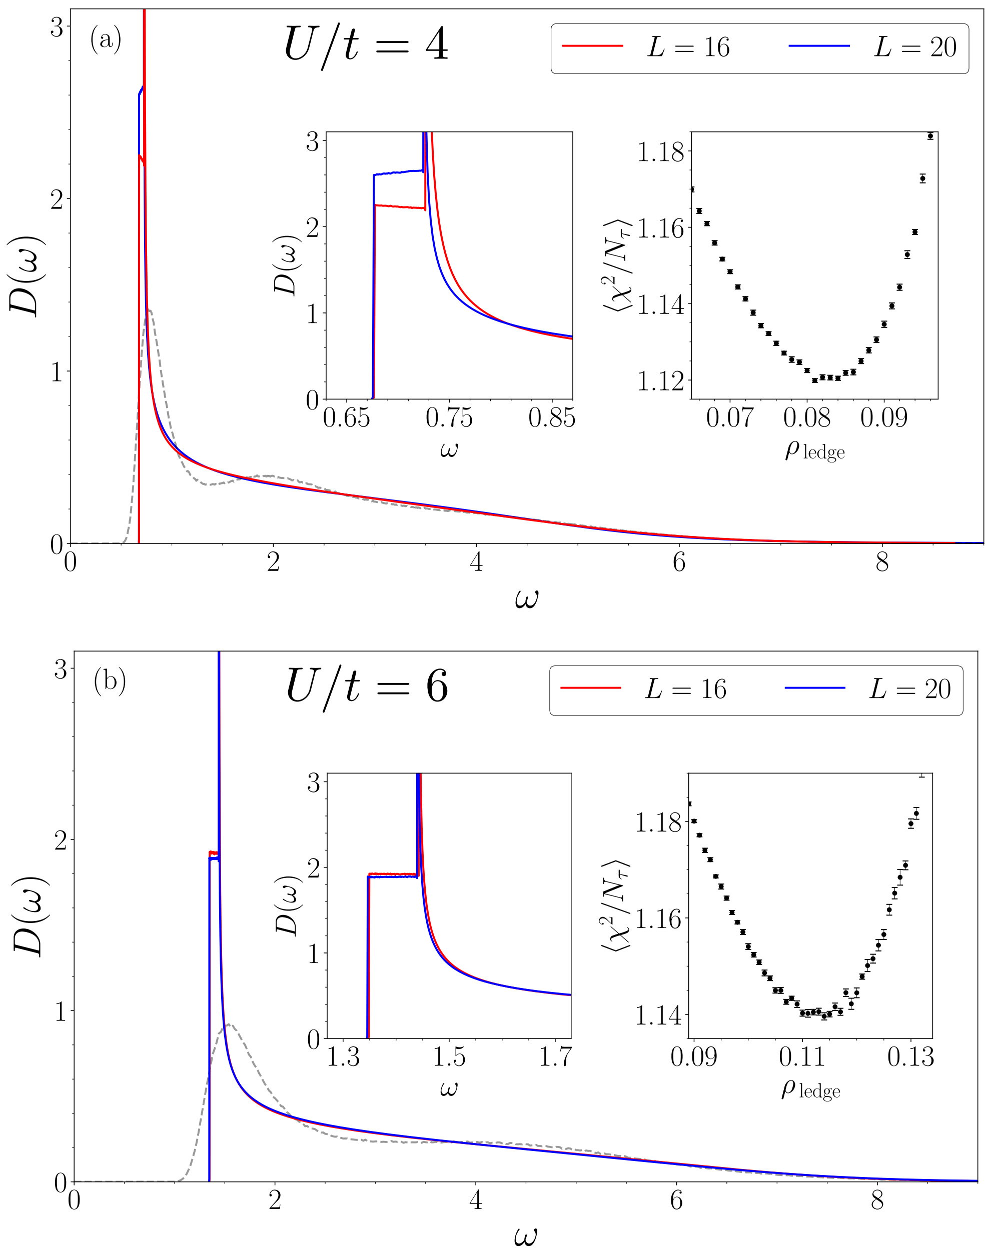

Figure 4·6: DOS obtained with the double-edge SAC parameterization for both $L = 16$ (blue) and $L = 20$ (red), with $U/t = 4$ in (a) and $U/t = 6$ in (b). The dashed grey lines are the results from unconstrained sampling for $L = 20$. The left insets show zoomed in views of the ledge region between ω_Σ and ω_X. The right insets show the goodness-of-fit versus the ledge fraction ρ_{ledge} for $L = 20$. The spectra shown were obtained using ρ_{ledge} of the minimum; in (a) $\rho_{\text{ledge}} = 0.071$ for $L = 16$ and $\rho_{\text{ledge}} = 0.082$ for $L = 20$, while in (b) $\rho_{\text{ledge}} = 0.115$ for $L = 16$ and $\rho_{\text{ledge}} = 0.113$ for $L = 20$.

4.5 Summary and Discussion

Despite the fundamental importance of the 2D Hubbard, crucial details of the quasiparticle dispersion relation and the density of states have been lacking. It has for long been established that the related t-J model hosts its lowest $x = 0$ quasiparticle for small J/t at $\mathbf{k} = \Sigma$, and a rather flat band around $\mathbf{k} = \mathrm{X}$ has been observed (whereas a local maximum appears at X for larger J/t [17, 88, 78, 89]). The energy splitting of the Σ and X quasiparticles is also a well documented feature of the underdoped cuprates [98, 79, 132, 135, 86, 3, 57]. However, this behavior has not been established in the case of the Hubbard model at moderate U values of relevance to the cuprates [7, 8, 59].

Our results demonstrate that the smallest gap in the Hubbard model is clearly at Σ, and we confirm a dispersion close to quartic around X. There is a barely resolvable local X maximum, seen in Fig. 4·2(b), which implies a minor rounding of the singular DOS at ω_{X} that cannot be resolved with our methods. Having established these facts, we determined the fraction ρ_{ledge} of states below ω_{X}. Within the RBA, four hole pockets would form upon light doping $x > 0$ (as observed in the form of "Fermi arcs" in underdoped cuprates [132, 135]), and merge into a contiguous Fermi sea as the energy approaches ω_{X}.

The doping at which ledge states will be exhausted and the Fermi surface will reconnect, x_c, would at first sight be be equal to ρ_{ledge}, ~ 0.1 for the two U/t values considered here. However, this estimate of x_c neglects the fact that there is a significant continuum extending rather far above ω_X, corresponding to a collection of excited states that dress the $\omega_{\mathbf{k}}$ quasiparticles. A simple way to correct for the fact that only the peak contribution to $A_{\mathbf{k}}(\omega)$ is accounted for in ρ_{ledge} is to divide by the quasiparticle weight a_0 for the states with $\omega_{\mathbf{k}}$ within the ledge, ~ 0.7 for $U/t = 4$ and ~ 0.6 for $U/t = 6$. A more accurate approach to calculating x_c in the RBA is by direct counting of the fraction n_{ledge} of quasiparticle energies below ω_{X}.

We performed a series of systematic quadratic fits to the dispersion across the noninteracting Fermi surface, in order to map out how the Fermi sea evolves upon doping within the RBA [128, 33, 32]. The perpendicular quadratic dispersion gradually and continuously flattens as $\mathbf{k}$ approaches X. The quadratic coefficient can be well fit to a cosine function, though this form is likely only approximate but sufficiently precise for our purposes. Using this function, along with the quadratic fit along the noninteracting Fermi surface (green in Fig. 4·5), we can model the quasiparticle dispersion in the vicinity of Σ by a continuous function $\omega(\mathbf{k})$. This allows us to map out the Fermi surface at any level of doping in the RBA by identifying equal energy contours, as shown in Fig. 4·7 for both $U/t = 4$ and $U/t = 6$. We note that as the energy approaches ω_{X} (the blue contour) , the Fermi surface flattens near the X point, reflecting that ω_{X} actually corresponds to a slight local maximum in the quasiparticle energy. This is also found in the closely related t-J model [17, 88, 78, 89].

The $\mathbf{k}$-space area of the contour with energy ω_{X} provides an alternative estimate for the critical level of doping x_c. In the rigid band picture, the fraction of the total number of $\mathbf{k}$-points with quasiparticle energies lying within the ω_{X} (blue) contour, n_{ledge}, is exactly equal to x_c. With these ledge stages exhausted, doped holes would begin to occupy the higher energy states around X, merging the four hole pockets that form in the underdoped system. This gives a critical doping level of $x_c = n_{\mathrm{ledge}} \approx 0.13$ and 0.20 for $U/t = 4$ and 6, respectively; both larger than the values of ρ_{ledge} extracted from the DOS. The discrepancy between n_{ledge} and ρ_{ledge} can be attributed to the fact $A_{\mathbf{k}}(\omega)$ contains both a quasiparticle peak and an incoherent continuum—a doped hole with momentum $\mathbf{k}$ will thus be in a superposition of states with energies determined by the spectral weight distribution of these two features. This implies that ρ_{ledge} should underestimate x_c exactly by a factor of a_0, which is verified by the observed relation $n_{\text{ledge}} \approx \rho_{\text{ledge}}/a_0$, with $a_0 \sim 0.7$ and 0.6 for $U/t = 4$ and 6, respectively.

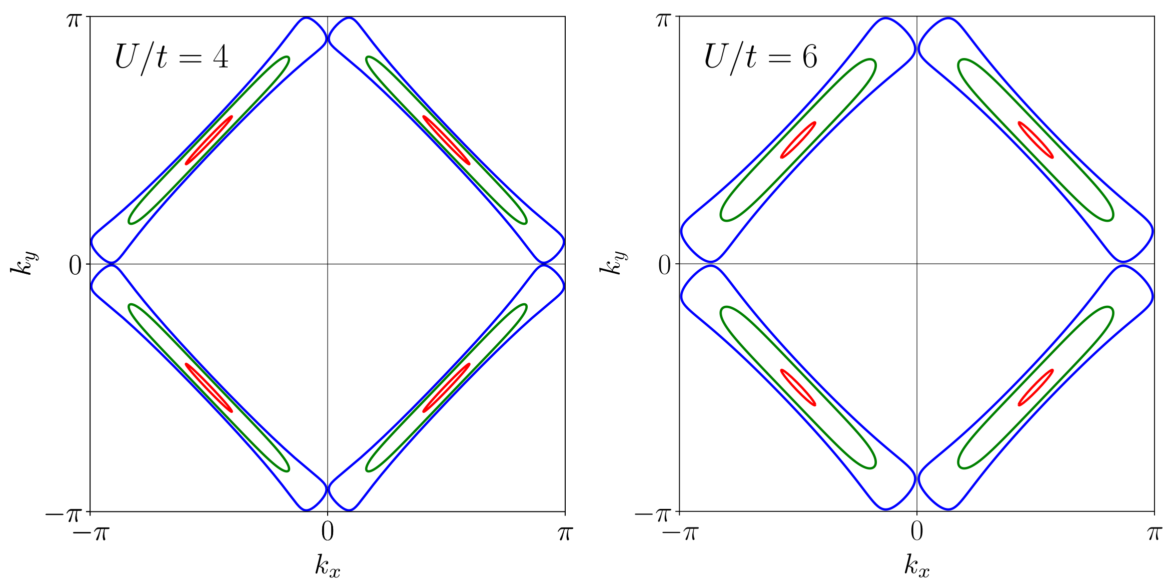

Figure 4·7: Fermi surfaces at doping levels corresponding to three energy levels: just above ω_Σ (red), just below ω_{X} (blue), and an intermediate energy (green).

The exact values of n_{ledge} obtained here should not be taken as specific predictions for the cuprates, as there will clearly be significant effects of interactions beyond the Hubbard model. Indeed, the rather large dependence on U can be taken as a general high sensitivity of n_{ledge} to model parameters. Our main point here is the presence of the second singularity at ω_X, which should imply a drastic change in the doped state at $x_c = n_{\text{ledge}}$, a change from a "plain" doped Mott insulator [80, 103], likely with strong spin and charge density correlations [150], to something else. While it appears plausible that the RBA applies at low doping in the Hubbard model, in some calculations, a charge-density-wave (CDW) or stripe instability takes place that would likely have to involve breakdown of the rigid band; a Fermi surface reconstruction [131, 38, 9, 103]. Furthermore, it is now believed that Luttinger's theorem is violated in the underdoped cuprates [155, 121, 85] due to the emergence of topologically ordered states, as investigated within a fermion-doped dimer-singlet model [112, 111]. If this pictures applies immediately upon doping, it would invalidate the application of the RBA to describe the low doping behavior of these systems. But in the absence of such complications, a compelling scenario emerging from our study is the rigid band breaking down only when the doping exceeds the ledge fraction at $x_c = n_{\mathrm{ledge}}$. A reason for this doping induced Fermi surface reconstruction could be the onset of (π, π) scattering between the X "hot spots", facilitated by exchange of magnons and amplified by the sharp DOS peak at ω_{X} [142].

Though the standard Hubbard model may not itself have a superconducting phase [104, 105], the second singularity in the DOS could still correspond to a critical point. The extended superconducting phase would then be induced only in the presence of additional interactions, with t' hopping the most promising candidate so far [52, 150]. While the mechanism of high-temperature superconductivity is still an open question, the Hubbard model remains a key piece of this puzzle, and our presented results provide further insight into their deep connection.

Chapter 5

Cross Validation in Stochastic Analytic Continuation

5.1 Introduction

In both the unconstrained and constrained SAC sampling schemes, the average χ^2 of the sampled spectra is used to optimize various parameters and gauge the ability of the spectral representation to describe the data. While these methods that rely on the average χ^2 alone have been shown to produce reliable and consistent results, in this chapter, we investigate the predictive power of these simple schemes by taking advantage of concepts from the machine learning and statistics literature. We do this by implementing a cross validation procedure [13, 84].

Cross validation is a tool used to compare how well different models describe a set of data. This is generally achieved by splitting the available data into mutually exclusive sets; a training, or sampling, set and a series of testing, or validation, sets. A model is trained on the sampling set, typically by optimizing some type of loss function. After training, the ability of the optimized model to describe the validation data sets is determined using this same loss function. This can provide information about the suitability of the model to describe the data, as well as the level of over-fitting to noise in the sampling data set.

Cross validation has a natural extension to SAC. The imaginary-time correlation function data, calculated using QMC simulations, are typically stored as individual bins, which can be split up into sampling and validation sets. Here, the goodness-of-fit χ^2 is used as the loss function to "train" and to "validate" the resulting spectral function. The concept of using cross validation within SAC was first explored by Efremkin et al. [34], whose work we draw on and significantly expand upon. Here, we perform an exhaustive study using both synthetic and real QMC-generated data.

In Sec. 5.2 we explain our cross validation scheme in detail, including the necessary modifications to the SAC procedure detailed in Chapter 3. In Sec. 5.3.1 and 5.3.2 we implement cross validation using the unconstrained SAC parameterizations and demonstrate how it can be used to determine the optimal sampling temperature, which we compare to the simple criterion proposed previously [124], Eq. (3.23). In Sec. 5.4.1 and 5.4.2, we show how cross validation can be used in practice to determine which *unconstrained* SAC sampling parameterization, using equal amplitude δ-functions or allowing for the amplitudes to vary and be sampled, produces the most statistically likely spectrum. We extend this model selection procedure to the class of *constrained* SAC parameterization in Sec. 5.5. In Sec. 5.5.1 we perform a test using QMC-generated data for the $S = 1/2$ antiferromagnetic (AFM) Heisenberg chain, where it is known that the spectral function for the operator $O = S^z_q$ contains a power-law divergent edge. In Sec. 5.5.2 we perform a preliminary test on the $S = 1/2$ AFM Heisenberg chain with long-range interactions [151, 152], a model where the exact features of the spectral function are unknown. We summarize our results and discuss the remaining open questions and future directions in Sec. 5.6.

5.2 Cross Validation Procedure

We begin by describing the implementation of cross validation within the SAC framework. As with the standard SAC procedure, we consider N_B bins of imaginary-time correlation function data evaluated at a set of N_τ points τ_i, $\{G^b(\tau_i)\}$, where

$b = 1, 2, \ldots, N_B$ and $i = 1, 2, \ldots, N_\tau$. The N_B bins are then split into $K+1$ mutually exclusive sets $\{G^b(\tau_i)\}_k$ of equal length $N_k = N_B/(K+1)$. For a single cross validation run, one of the $K+1$ groups will act as the sampling data set, while the other K groups will be used for cross validation. For each group, we compute the average

$$\bar{G}_k(\tau_i) = \frac{1}{N_k} \sum_{b \in k} G^b(\tau_i), \tag{5.1}$$

where the sum is over all bins b in the kth data set. We define the "error level" for these sets of $G(\tau)$ data as the magnitude of the error bar on the averaged correlation function, $\bar{G}_k(\tau_i)$, for the largest τ point included, when normalized so that $\bar{G}_k(0) = G^b(0) = 1$. For larger β, where the relative error can become substantial at large τ, a cutoff is typically introduced at the τ value where this relative error on $\bar{G}_k(\tau)$ is greater than 10%. In this case, we instead use this data point to quantify the error level.

The error level provides a rough gauge of the data quality, but, just as in the standard SAC prescription, the covariance matrix is required for the full characterization of the error and the calculation of χ^2 [62]. We still use the bootstrap method to calculate the covariance matrix, but we now do so for each of the $K+1$ data sets:

$$C_{k,ij} = \frac{1}{M} \sum_{m=1}^{M} (G_k^m(\tau_i) - \bar{G}_k(\tau_i))(G_k^m(\tau_j) - \bar{G}_k(\tau_j)), \tag{5.2}$$

where G_k^m represents the average of a bootstrap sample of N_k randomly chosen bins within the kth data set.

Using the $k = 0$th set of bins as the sampling data set, we carry out a simulated annealing run described in Sec. 3.3. Along with the sampling χ^2,

$$\chi_0^2 = \sum_{i=1}^{N_\tau} \left(\frac{G(\tau_i) - \bar{G}_0(\tau_i)}{\sigma_0(\tau_i)} \right)^2, \tag{5.3}$$

at each temperature, we use the sets $k = 1, \ldots, K$ to compute the validation χ^2 value:

$$\chi^2_{\text{val}} = \frac{1}{K} \sum_{k=1}^{K} \sum_{i=1}^{N_\tau} \left(\frac{G(\tau_i) - \bar{G}_k(\tau_i)}{\sigma_k(\tau_i)} \right)^2 . \tag{5.4}$$

As always, we use the full covariance matrix to calculate χ^2, but transform to its eigenbasis so the calculation can be done using the computationally simpler diagonal form.

5.2.1 Connection to the Bias-Variance Trade Off in Machine Learning

The inquisitive reader may be curious why the set of sampling $G(\tau)$ bins is so much smaller than the set of validation $G(\tau)$ bins in our cross validation procedure, in contrast to the standard used in traditional machine learning [84]. This choice can be motivated by performing a decomposition of the χ^2_{val} formula in Eq. (5.4) in a manner similar to the bias-variance decomposition of the out-of-sample error [84] in machine learning.

If we denote the exact correlation function by $G_{\text{ex}}(\tau)$, Eq. (5.4) can be rewritten as

$$\chi^2_{\text{val}} = \frac{1}{K} \sum_{k=1}^{K} \sum_{i=1}^{N_\tau} \left(\frac{G(\tau_i) - G_{\text{ex}}(\tau_i) + G_{\text{ex}}(\tau_i) - \bar{G}_k(\tau_i)}{\sigma_k(\tau_i)} \right)^2 , \tag{5.5}$$

which can be split into three terms:

$$\chi^2_{\text{val}} = X_1 + X_2 + X_3, \tag{5.6}$$

where

$$X_1 = \frac{1}{K}\sum_{k=1}^{K}\sum_{i=1}^{N_\tau}\left(\frac{G(\tau_i) - G_{\rm ex}(\tau_i)}{\sigma_k(\tau_i)}\right)^2, \tag{5.7a}$$

$$X_2 = \frac{1}{K}\sum_{k=1}^{K}\sum_{i=1}^{N_\tau}\left(\frac{\bar{G}_k(\tau_i) - G_{\rm ex}(\tau_i)}{\sigma_k(\tau_i)}\right)^2, \tag{5.7b}$$

$$X_3 = \frac{2}{K}\sum_{k=1}^{K}\sum_{i=1}^{N_\tau}\frac{\left(\bar{G}_k(\tau_i) - G_{\rm ex}(\tau_i)\right)\left(G_{\rm ex}(\tau_i) - G(\tau_i)\right)}{\sigma_k^2(\tau_i)}. \tag{5.7c}$$

Here X_1 contains purely information about how well the estimated $G(\tau)$, calculated from the sampled $S(\omega)$, describes the exact correlation function. This will vary as we change Θ, or any other parameter used in the SAC program, and will be minimized by the most statistically likely spectral function, given input data $\bar{G}_0(\tau)$ and the corresponding covariance matrix C_0. In the standard bias-variance decomposition, X_1 is analogous to the the Bias^2 + Variance term [84].

X_2 is the the χ^2 value for the validation data with respect to the underlying data $G_{\rm ex}$, averaged over each of the K validation data sets. The value of X_2 will follow the χ^2 distribution with $N_{\rm dof} = N_\tau$, and will have a mean of $E\left[X_2\right] = N_{\rm dof} = N_\tau$ and a variance of $V\left[X_2\right] = 2N_{\rm dof} = 2N_\tau$. This term is independent of any SAC parameter and will simply contribute a background value of N_τ, in the limit of large K. In the standard bias-variance decomposition, X_2 is analogous to the the Noise term [84].

X_3 contains both information about the accuracy of $G(\tau)$ with respect to $G_{\rm ex}(\tau)$ and the deviation of the $\bar{G}_k(\tau)$ with respect to $G_{\rm ex}(\tau)$. However, these two factors are uncorrelated, and since we expect that the fluctuations of $\bar{G}_k(\tau)$ will follow a Gaussian distribution with a mean of zero, this term will also have a mean of zero.

Using the above properties, the *reduced* validation χ^2 becomes simply

$$\frac{\chi^2_{\rm val}}{N_\tau} := \frac{1}{N_\tau K}\sum_{k=1}^{K}\sum_{i=1}^{N_\tau}\left(\frac{G(\tau_i) - G_{\rm ex}(\tau_i)}{\sigma_k(\tau_i)}\right)^2 + 1. \tag{5.8}$$

It is not straightforward to separately define the bias and variance of a SAC model, as one can do in a standard machine-learning model. Each SAC parameterization does make its own simplifying assumptions about the underlying spectral function (i.e. contain its own biases) and does fluctuate according to the input data to its own degree (i.e. contains its own model variance), but distilling these two effects as a function of parameterization complexity is unfeasible. Consequently, in SAC the bias-variance tradeoff does not play quite as large of a role in model selection as it does in typical machine learning.

For our goal of cross validation as a tool for model selection, we simply aim to minimize the combined Bias^2 and Variance term X_1 within as a class of suitable SAC parameterization, while minimizing the fluctuations in X_2 that may mask the signal.

5.2.2 Choosing the Optimal Number K of Validation.

To determine what K, we analyze the individual K scaling of the X_1:

$$X_1 = \sum_{i=1}^{N_\tau} \left(G(\tau_i) - G_{\text{ex}}(\tau_i)\right)^2 \times \frac{1}{K} \sum_{k=1}^{K} \frac{1}{\left(\sigma_k(\tau_i)\right)^2} \tag{5.9}$$

$$\propto \sum_{i=1}^{N_\tau} \left(G(\tau_i) - G_{\text{ex}}(\tau_i)\right)^2 \times \frac{1}{K} \sum_{k=1}^{K} \frac{N_k}{\left(\sigma_{\text{bin}}(\tau_i)\right)^2} \propto \frac{1}{K}, \tag{5.10}$$

where we have used that variance of the the the τ_i data point in validation group k is equal to $\sigma^2_{\text{bin}}/N_k$, where σ^2_{bin} is the variance of the individual bin measurements. As K decreases, and the number of data points in each validation set increases, the signal in X_1 will be amplified. However, the variance in X_2, which contributes a constant background value of N_τ, also scales as $1/K$. So the optimal K will balance these two effects–maximizing the signal in X_1 while minimizing the noise contributed by X_2.

One option would be to optimize K using cross validation itself. We can perform the cross validation procedure for a series of K values, and use the model selection results from the K value that produces the lowest χ^2_{val} amongst all parameterizations tested. However, when preforming cross validation, it is also important to consider the error level of the sampling data set. A smaller K lowers the error level on the sampling set, which will produce an overall better result for $S(\omega)$. As demonstrated in the many examples investigated in Ref. [124], an error level of $\sigma \sim 10^{-5}$ is typically sufficient to accurately reproduce the underlying spectra. When the spectral function contains many fine features, such as collections of small peaks or sharp edges, a smaller error level may be necessary to accurately resolve said features. However, for the spectra we consider here, this is not the case, so we have used data with $\sigma \sim 10^{-5}$ in all of the examples shown in the following sections. For the tests using QMC-generated data, this restricted us to using $K = 20$ validation data sets. For consistency, we have also used this number of validation sets for our tests using synthetic data, where we have direct control of both the error level and number of $G(\tau)$ bins.

We have tested an alternative approach that maintains as high a data quality as possible in the sampling/validation sets and maximizes the signal in X_1 by splitting the $G(\tau)$ data into just two mutually exclusive sets. To reduce fluctuations in the X_2, we repeat this process many times, with new random partitions for the sampling and validation sets, and average the results. This method produced results that agreed with those generated using the procedure outlined above ($K = 20$), which is consistent with our experience using SAC. Increasing the error level by a factor of $\sqrt{21} \approx 4.5$ (which is the effect of using only 1/21st of the total number of $G(\tau)$ bins) will rarely be detrimental to the accuracy of SAC, assuming the data quality is high to begin with and the spectral function does not contain many fine features. This approach does offer the advantage of being able to generate additional validation data sets as needed, until the results converge, and is ultimately the method we suggest using moving forward.

5.3 Optimal Sampling Temperature

5.3.1 Single Gaussian Peak

Our first implementation of cross validation will be to provide support for the optimal Θ criterion, Eq. (3.23). The motivation behind this criterion is to achieve a statistically good data fit while placing the simulation in a temperature range safely above the regime of over-fitting, which we will be able to identify by tracking the behavior of the validation χ^2 as Θ is lowered.

For our tests, we will use synthetic $G(\tau)$ data aimed to mimic the data generated in a QMC simulation. Given an artificial bosonic spectral function $S(\omega)$, we use Eq. (3.7) to calculate the exact imaginary-time correlation function. To introduce the statistical errors necessarily present in QMC-generated data, we generate many $G(\tau)$ bins with noise that is correlated in imaginary-time. This is done by first generating normally distributed noise for each τ point $\varepsilon_0(\tau_i)$, and then taking the weighted averaging over all τ points with an exponentially decaying weight function:

$$\varepsilon(\tau_i) = \sum_{j=1}^{N_\tau} \varepsilon_0(\tau_j) e^{-|\tau_i - \tau_j|}. \tag{5.11}$$

It was shown that the presence of covariance in the data actually improves the results of the SAC method [124], so including this in our synthetic data is critical when comparing to QMC-generated data.

In our first test, we consider a spectrum consisting of a single Gaussian peak, centered at $\omega = 3$. This spectrum can be essentially perfectly reproduced using the unconstrained SAC procedure, Fig. 3·1(a) and (b). In this example, we use $N_\omega = 5,000$ total δ-functions, which is large enough to no longer detect the dependence of $S(\omega)$ on N_ω.

We generated $K = 20$ independent validation data sets, each containing 1,000 bins of synthetic $G(\tau)$ data with an error level $\sigma = 10^{-5}$. While it is certainly reasonable to generate this volume of data using modern QMC methods, this would be an inefficient use of computational resources given that the majority of the data is being used for validation. However, we used this large quantity so that the simplifications made in the large K limit are more accurately realized. As we show in Sec. 5.5, one can still use cross validation with far fewer $G(\tau)$ bins, even when this limit is not fully reached.

When converting from $S(\omega)$ to $G(\tau)$, we set the inverse temperature $\beta = 2$ and use a τ spacing of $\Delta_\tau = 0.0625$, giving $N_\tau = 16$ imaginary time points. The relatively low value of β used here was chosen in light of the fact that spectra consisting of broad and smooth features, such as the simple Gaussian used here, would usually be found at higher temperatures.

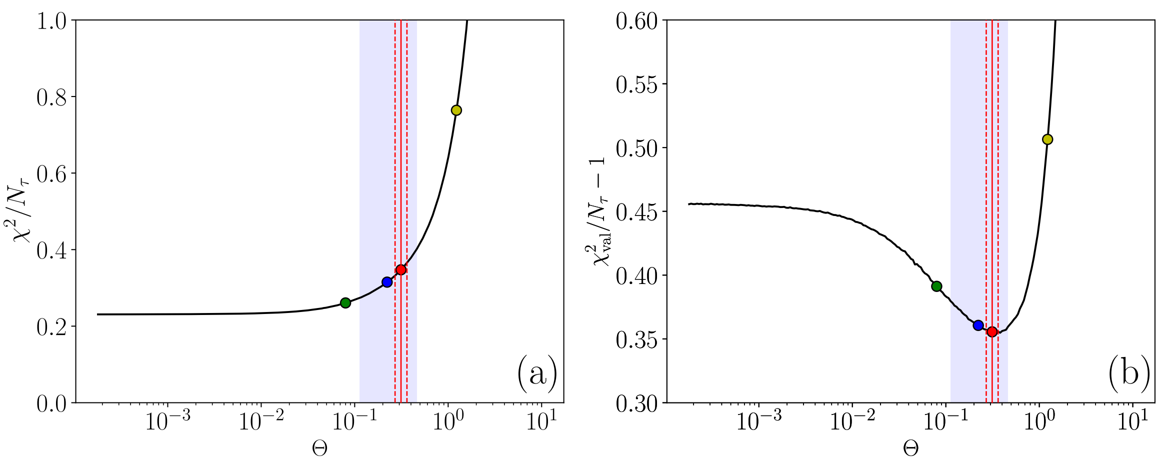

Figure 5·1: For the spectral function shown in Fig. 5·2, the average of $K = 20$ cross validation annealing runs using synthetic data with a noise level of $\sigma = 10^{-5}$. Panels (a) and (b) show the behavior χ^2 and χ^2_{val} versus Θ. The temperature range corresponding to Eq. (3.23) with $a = 0.25 - 1.0$ is shaded in blue. The temperature where χ^2_{val} is at its minimum is marked by a red vertical line. The spread (one standard deviation) in the location of the χ^2_{val} minimum is denoted by the red dashed lines. The χ^2 and χ^2_{val} values have been normalized by the number of τ points, and the background value of 1, corresponding to X_2 in Eq. (5.6) has been subtracted from χ^2_{val}. The colored points along the curves mark the sampling temperatures of the corresponding spectra in Fig. 5·2

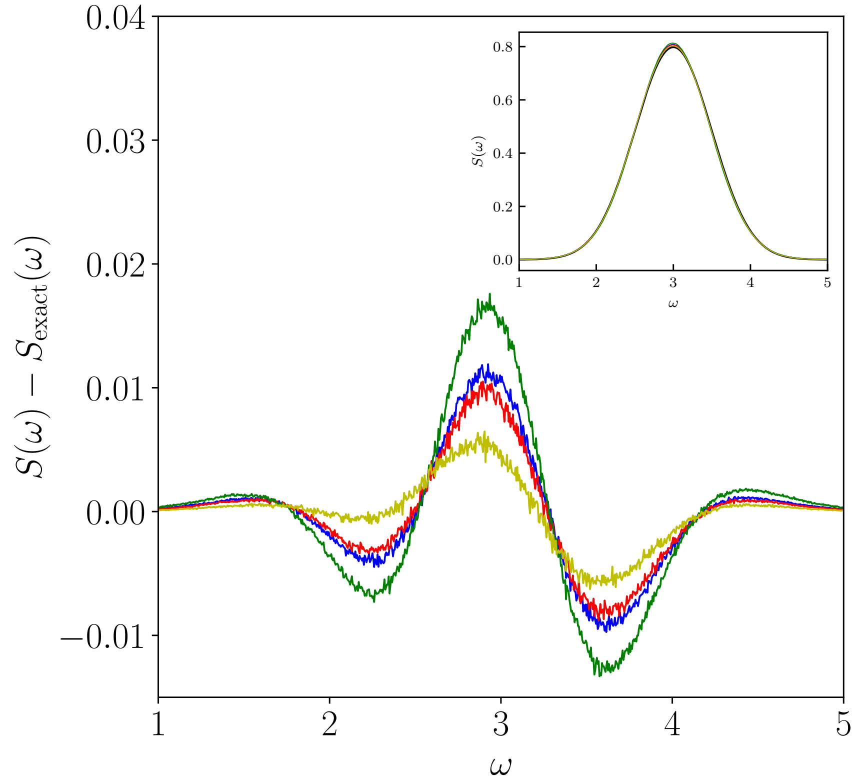

Figure 5·2: Deviation between the exact spectrum and those sampled at four different temperatures Θ: that where $a = 0.5$ (blue), where χ^2_{val} is minimized (red), and Θ values above/below (yellow/green) this minimum. Inset: the spectral function themselves. In this case, it is difficult to distinguish the different spectra, as they nearly overlap.

The results of the cross validation procedure are presented in Fig. 5·1 and Fig. 5·2. Fig. 5·1 show the behavior of the sampling χ^2 (panel (a)) and the validation χ^2 (panel (b)) versus Θ, averaged over all $K = 20$ validation data sets and normalized by the number of τ points. The temperature range corresponding to Eq. (3.23) with $a = 0.25 - 1.0$ is shaded in blue. The temperature where the validation χ^2 is at its minimum is marked by a red vertical line, with the red dashed lines denoting the standard deviation of the location of the minimum. In practice, it is useful to repeat this process, rotating which of the $K + 1$ data sets is used for sampling and which are used for validation, and averaging the results in the end. However, for this test we found this this was not necessary given the large number of synthetic data bins that can be easily generated.

In Fig. 5·2, we show the difference between the the exact spectrum and the SAC spectra sampled at the temperature where $a = 0.5$ (blue), at the validation χ^2 minimum (red), and at temperatures above (yellow) and below (green) the validation minimum, for comparison. In the inset of Fig. 5·2, we show the spectral functions themselves. To generate these spectra we ran the SAC procedure using all bins, recombined from the $K + 1 = 21$ sets, thus taking full advantage of the data at hand.

In this case, the minimum of the validation χ^2 is quite sharp and agrees very well with the temperature from Eq. (3.23). Consequently, the red and blue spectra in Fig. 5·1 are nearly identical, and, for all intents and purposes, are equally valid representations of the true spectrum, shown in black. The green and yellow spectra are included to show the relatively small Θ dependence for this simple spectrum.

5.3.2 Two Gaussian Peaks

We now turn to an example where the unconstrained sampling parameterization cannot reproduce the artificial spectrum quite as closely. The spectrum we consider is composed of a pair of Gaussian peaks, one broader than the other, that intersect slightly. The area near the sharp dip, where the two peaks overlap, poses difficulties and would require a lower error level than one could reasonably achieve using real QMC-generated data to resolve fully [124]. As such, the criterion in Eq. (3.23) provides an estimate for the temperature corresponding to the best possible spectrum, given these limitations.

We performed the same test as in Sec. 5.3.1, and present the results in Fig. 5·3. In this case, both the validation minimum and the temperature range corresponding to the criterion with a ranging from 0.25 and 1 are much broader, but still align quite well with one another. The higher level of uncertainty is a reflection of the inability of the unconstrained sampling scheme to reproduce this spectrum nearly as well as in the previous case, which can be seen clearly in Fig. 5·4. Here we again plot the spectra at the validation χ^2 minimum (red) and at the temperature where $a = 0.5$ (blue), as well as at a slightly lower/higher temperature, corresponding to one standard deviation below/above from the validation minimum (green/yellow). It is interesting to note that, at the lower sampling temperature (green), SAC is able to reproduce certain features of the exact spectrum better than at the temperatures corresponding to the validation minimum or $a = 0.5$ (such as the position of the first peak and area where the two peaks intersect), while reproducing others not as well (such as the position of the second peak). This is consistent with the broadness of the validation χ^2 minimum, as it is not necessarily clear which is the better choice of spectra. Furthermore, all of the spectra shown here are acceptable reproductions of the underlying artificial spectrum; the most prominent features of the spectrum are captured well in all four cases and the deviations plotted in Fig. 5·4 reflect only small shifts in the locations and amplitudes of the two Gaussian peaks.

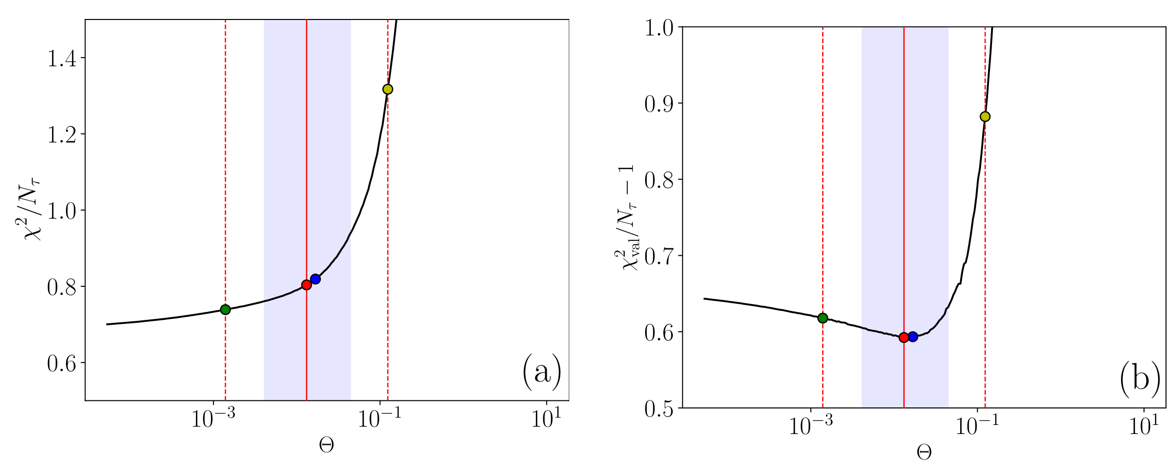

Figure 5·3: For the spectral function shown in Fig. 5·4, the average of $K = 20$ cross validation annealing runs using synthetic data with a noise level of $\sigma = 10^{-5}$. Panels, symbols, and set up are identical to those in Fig. 5·1.

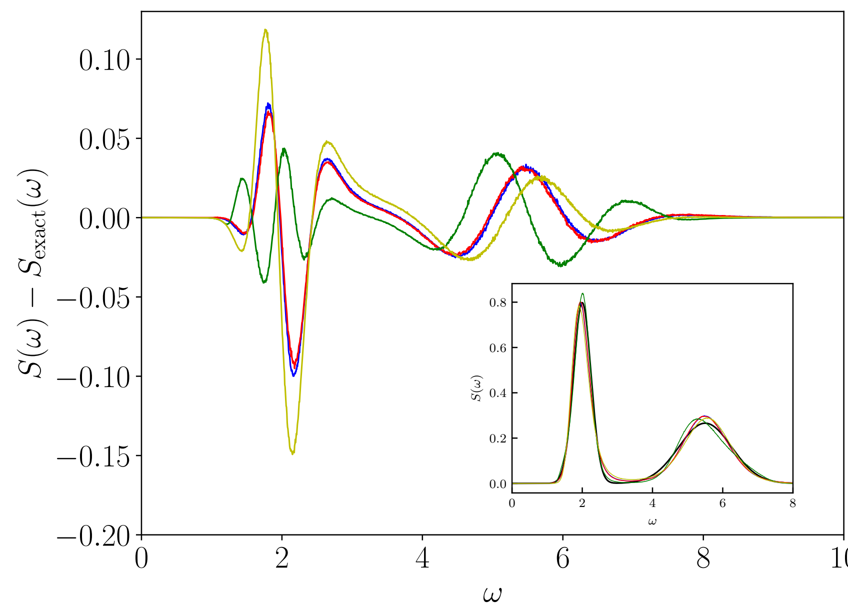

Figure 5·4: Deviation between the exact spectrum and those sampled at three different temperatures: the temperature where $a = 0.5$ (blue), the temperature where χ^2_{val} is minimized (red), at a slightly lower temperature, corresponding to one standard deviation below this minimum (green), and at a slightly higher temperature, corresponding to one standard deviation above the minimum (yellow). We note that the magnitude of the deviation is both a reflection of the error in the amplitudes and the locations of the two Gaussian peaks. Inset: the spectral function themselves. Just as in Fig. 5·2, the red and blue spectra are nearly indistinguishable.

We emphasize that using cross validation to determine the optimal sampling temperature is not necessary when running SAC with real QMC-generated data. Rather, these test-cases show that the criterion in Eq. (3.23) indeed samples safely above the regime of over-fitting, while still keeping χ^2 at an acceptable level, and *should* be used as the method of fixing Θ in real use-cases. In the following sections we will explore a more practical and decisive use for our cross validation procedure, as a tool for selecting the most applicable SAC parameterization.

5.4 Optimal Spectral Parameterization: Unconstrained Sampling

Our second application of cross validation will be to help determine which SAC parameterization to use when the exact spectral features are unknown, i.e. model selection [13, 84]. In the unconstrained parameterization, we can either use fixed ampltiude δ-functions and only sample their positions (Fig. 3·1(a)), or additionally, sample their amplitudes (Fig. 3·1(b)). These two parameterizations, which we'll refer to as equal and variable amplitudes (EA and VA, respectively), have different configurational entropies as well as different sampling efficiencies, depending on the exact form of the underlying spectrum. For the EA parameterization, the entropy is the conventional Shannon information entropy [10, 124], while for the VA parameterization, a generalized Rényi entropy is instead appropriate [42].

In many cases, such as the two considered in the previous section, the output spectra will depend very little, if at all, on whether amplitudes are sampled along with frequencies. More specifically, if N_ω is large enough, the spectral features are not very sharp, and the error level is low, the different entropy contents will only minimally affect the sampling efficiency and the underlying spectrum can be reproduced with high fidelity using either parameterization [124]. But this is not always the case, as we will explore in the following sections.

5.4.1 Two Equal Amplitude Gaussian Peaks

In the first case we consider, we used an artificial spectrum composed of two equal amplitude Gaussian peaks centered at $\omega = 2$ and $\omega = 4$. While this spectrum is relatively simple, the proximity of the two peaks poses similar difficulties to those encountered in Sec. 5.3.2; the low amplitude region in between the two peaks cannot be resolved fully at this error level, which causes a ripple effect of distortions throughout the spectrum. The two unconstrained sampling schemes, EA and VA, are able to overcome these challenges to different extents, which can be seen clearly when performing tests using artificial spectra. However, when performing analytic continuation using QMC-generated data, one must instead rely on cross validation to asses the abilities of these two parameterizations.

Just as in Sec. 5.3, we generated $K+1 = 21$ sets of synthetic data with error level $\sigma = 10^{-5}$. In this case, we used a slightly higher inverse temperature, $\beta = 16$, and accordingly a larger τ spacing, $\Delta\tau = 0.1$, when converting the artificial spectrum to $G(\tau)$. When defining the error level, we introduce a cutoff of $\tau_{\rm max} \approx 4$, chosen so that the relative error on $\bar{G}(\tau)$ does not exceed 10%. We ran the cross validation procedure using two different SAC parameterizations, the unconstrained sampling method with equal amplitude and variable, sampled amplitude δ-functions, and compare their performance.

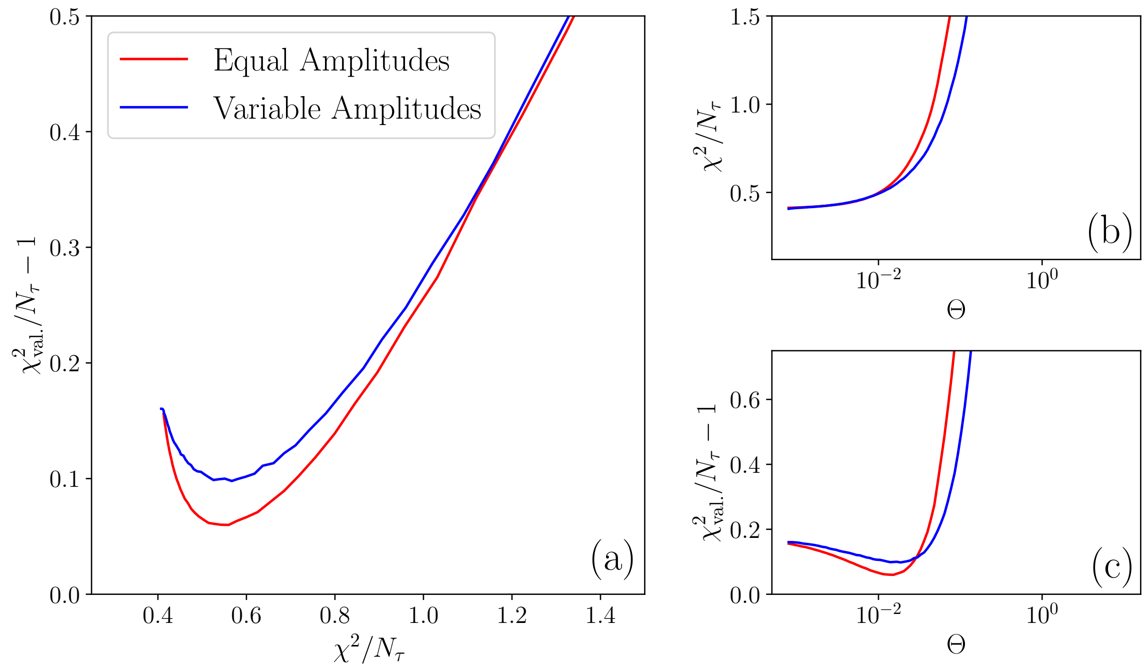

Figure 5·5: Cross validation results for the spectral function shown in Fig. 5·6. In panel (b), the sampling χ^2 values, normalized by the number of τ points, is plotted versus the sampling temperature Θ. In panel (c), the validation χ^2 values, also normalized by the number of τ points, but now also averaged over all cross validations data sets, is plotted versus the sampling temperature Θ. We have subtracted the background constant of one from the validation χ^2 to account for the term X_2 in Eq. (5.6). In panel (a), the validation χ^2 as a function of the sampling χ^2 is plotted to allow for direct comparison of the two parameterizations at a fixed value of the sampling χ^2.

Figure 5·5 shows the results of this cross validation test. The normalized sampling χ^2 during the annealing run is shown in panel (b) for both updating methods, equal (red) and variable (blue) amplitudes. Because the entropy contents of the spectra defined with each parameterization are different, the annealing paths they take differ, so it is also informative to compare the validation χ^2 values at the same value of the sampling χ^2. This is shown in Fig. 5·5(a) where we plot the validation χ^2 versus the sampling χ^2 throughout the annealing run.

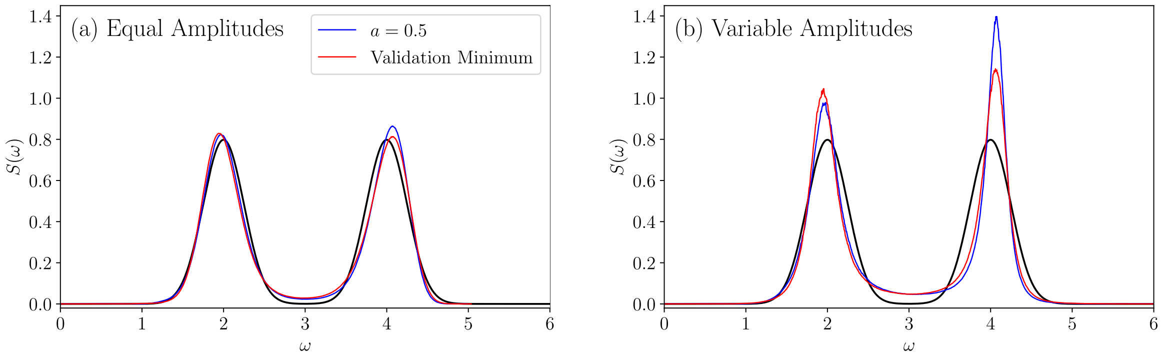

Figure 5·6: Spectra corresponding to the two parameterizations shown in Fig. 5·5, equal amplitudes (panel (a)) and variable amplitudes (panel (b)). The exact spectrum is plotted in black, while the two colored spectra correspond to two different sampling temperatures: the validation χ^2 minimum (red) and the criterion in Eq. (3.23) with $a = 0.5$ (blue). In both cases, $N_\omega = 4{,}000$ δ-functions were used, which was chosen to ensure that there was no dependance of $S(\omega)$ on N_ω.

For this spectrum, the validation χ^2 reaches the lowest minimum value for the EA parameterization, which can be clearly seen in Fig. 5·5 panel (c), and even more clearly in panel (a). We can see why this is the case by comparing the spectra produced using both parameterizations, shown in Fig. 5·6. Here, we plot the spectra at two sampling temperatures, corresponding to the validation χ^2 minima (red, with $\chi/N_\tau \sim 0.56$) and the criterion in Eq. 3.23 with $a = 0.5$ (blue, with $\chi/N_\tau \sim 0.54$).

The EA parameterization (panel (a)) produces a spectrum where both the shapes and heights of the two Gaussian peaks are very close to those in the exact spectrum, shown in black. On the other hand, the VA parameterization (panel (b)) produces peaks that are far too sharp with amplitudes that are unequal and too large. While both parameterizations are unable to correctly reproduce the region in between the two peaks, the VA parameterization performs far worse. The excessively tall peaks produced by the variable amplitudes parameterization may be compensation for the missing weight in this region.

Each SAC parameterization comes with its own set of entropic pressures associated with the degrees of freedom that are being sampled. In the case of the VA parameterization, it is known that these pressures overly favor sharp peaks [124]. In this example, that comes as a detriment, but in others, this can be advantageous, as we will explore in the following section.

5.4.2 Sharp Gaussian Peak with a Shoulder

For our second test, we consider an artificial spectrum composed of a tall and sharp Gaussian peak centered at $\omega = 2$, followed by a short and broad Gaussian at $\omega = 3$, which forms a shoulder-like feature for the first peak. A similar artificial spectrum was used in Ref. [124] to study this exact issue, whether or not to sample amplitudes along with frequencies. It was found that including amplitude updates was necessary to resolve both the dominant peak and the shoulder for data with error levels attainable using QMC [124]. Being able to determine this using cross validation would be very valuable when the exact underlying spectrum is unknown.

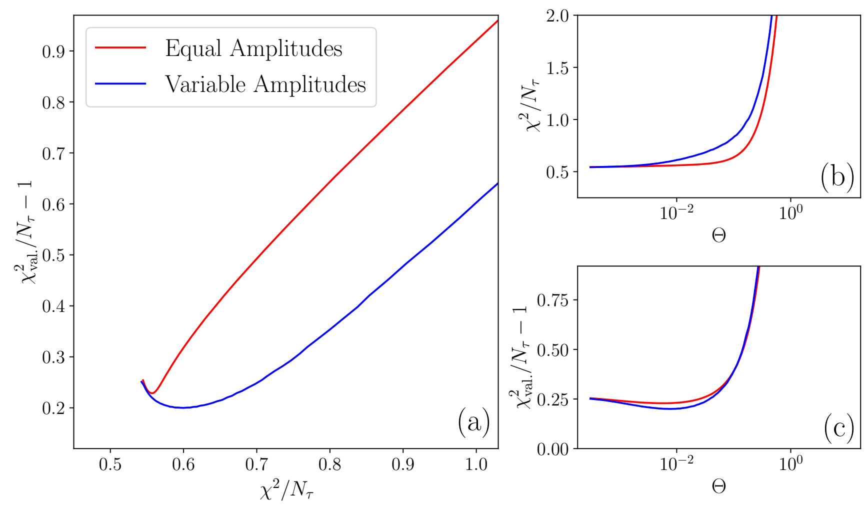

Figure 5·7: Cross validation results for the spectral function shown in Fig. 5·8. Panels and set up are identical to those in Fig. 5·5

Again, we generated $K + 1 = 21$ sets of synthetic data with error level $\sigma = 10^{-5}$

and use $\beta = 16$ with $\Delta\tau = 0.1$ (again, with $\tau_{\max} \approx 4$), which is the same inverse temperature used in Ref. [124] for their tests on this type of spectral function. The results of the cross validation procedure are shown in Fig. 5·7 and the corresponding spectral functions are shown in Fig. 5·8.

In this case, cross validation suggests that the VA parameterization performs better, in contrast to the previous example. Here, the validation χ^2 reaches a lower minimum value (panel (c)), but is also lower for the majority of the annealing process, when compared at fixed values of the sampling χ^2 (panel (a)).

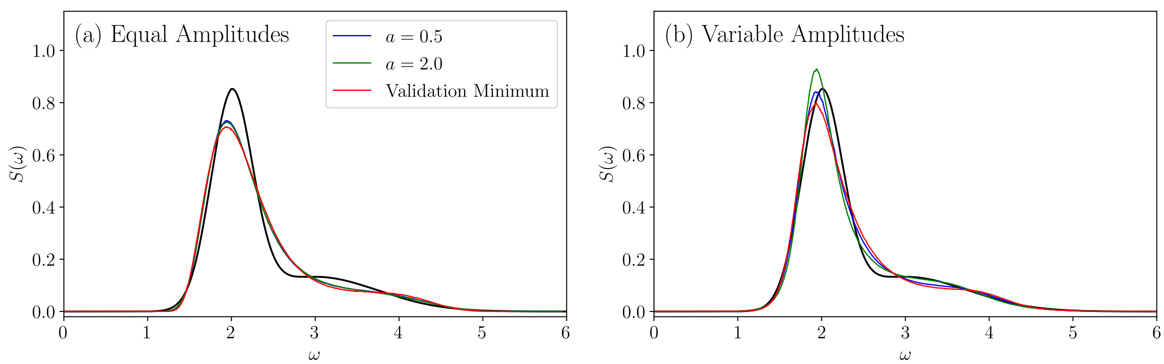

Figure 5·8: Spectra corresponding to the two parameterizations shown in Fig. 5·6. Panels and set up are identical to those in Fig. 5·6, except an additional spectrum is plotted, corresponding to the criterion in Eq. (3.23) with $a = 2.0$.

In Fig. 5·8 we plot the spectra at three different sampling temperatures for both parameterizations. The red spectra were sampled at the temperature corresponding to the validation minima, and the blue and green spectra were sampled at temperatures corresponding criterion in Eq. (3.23) with $a = 0.5$ and $a = 2.0$, respectively.

At all three sampling temperatures, the VA parameterization is able to resolve the first peak with higher fidelity. In both cases, the spectra sampled at the validation χ^2 minima (red, with $\chi/N_\tau \sim 0.55$) do not reproduce the shoulder feature very well, instead displaying a ringing pattern around this structure. But as the temperature, and consequently the χ^2 value, are increased (blue, with $\chi/N_\tau \sim 0.66$ and green, with $\chi/N_\tau \sim 1.0$), the shoulder is reproduced with much higher accuracy by the VA parameterization. This is, however, not the case for the EA parameterization, as suggested by the results shown in Fig. 5·7(c).

In the previous examples, there was much better agreement between the validation χ^2 minima and the criterion in Eq. (3.23) than there is for this case. We explain this discrepancy by noting that at the error levels considered here, $\sigma = 10^{-5}$, this spectrum is unable to be produced with nearly as high a fidelity. If the SAC results will always deviate from the exact solution in some considerable way, then the validation χ^2 may be minimized by a spectrum that does not visually appear to match the exact spectrum quite as well as one that is produced at a slightly higher or lower sampling temperature. This is exemplified by the $a = 2.0$ spectrum in Fig. 5·8(b); while it may visually appear to be the most accurate, the results of the cross validation procedure suggests that this is not the case. Cross validation identifies the spectrum with the corresponding $G(\tau)$ that is closest to the exact solution, and in some cases this may be a spectrum that deviates from the exact $S(\omega)$ in multiple different ways, i.e. a compensating effect that produces multiple different distortions. Even considering these limitations, it is clear that cross validation can accurately determine which unconstrained sampling scheme produces the most accurate spectrum.

At this point it should be pointed out that in the four unconstrained sampling tests we have presented, the fidelity of the SAC spectrum varies depending on the exact form of the underlying spectrum. In light of this, the cross validation χ^2 value should not be taken as an absolute measure, but rather a relative statistic used to compare two (potentially imperfect) solutions. The ability of cross validation to identify which imperfect solution is more accurate, as gauged by direct comparison to the exact spectrum, is a testament to the robustness of this procedure.

5.5 Optimal Spectral Parameterization: Constrained Sampling

We now extend the concept of cross validation as a tool for model selection to the constrained SAC sampling schemes (Fig. 3·1(c)-(e)). In many cases, it is known that the spectral function contains a sharp feature, but the exact form of this feature, and thus which constrained sampling scheme to use, is unknown. Relying on the value of χ^2 alone to compare the resulting spectra is not always reliable, since the SAC method can produce spectra with acceptable χ^2 values even when the improper parameterization is being used (a consequence of the "ill-posed" problem). Our cross validation procedure can be used as an unbiased method not only to determine whether an unconstrained or constrained sampling scheme should be used, but also to determine which of the constrained sampling scheme produces the most statistically likely spectrum.

5.5.1 Heisenberg Chain

We first demonstrate the ability to discriminate between models using real QMC data, generated using the stochastic series expansion (SSE) method [115], for the $S = 1/2$ AFM Heisenberg chain. We followed the procedure outlined in Sec. 5.2, again with $K = 20$ validation sets with an error level of $\sigma = 10^{-5}$. The system size used here is $L = 512$, and the QMC simulation was run using the inverse temperature $\beta = 1024$, which is large enough to access ground state properties. We considered the spectral function for $O = S^z_q$, $S_q(\omega)$. at two momenta, $q = \pi/2$ and $3\pi/4$, as test cases. For both q values, the optimal weights of the macroscopic δ-peak were determined using the method detailed in Sec. 3.4.1 ($a_0 = 0.40$ and 0.35 for $q = \pi/2$ and $3\pi/4$, respectively, see Fig. 3·4). Due to the logarithmic correction to the square-root divergent edge form of the spectral function [24, 25, 99], the optimal power p

for each q value considered may not necessarily be equal to -0.50. There are some subtleties in determining the optimal p value, as discussed in Sec. 3.4.2 and expanded upon below.

p Scan

To produce a spectrum with an edge that diverges asymptotically with an arbitrary exponent $p \neq -0.5$, the δ-functions amplitudes can be varied according to the form,

$$a_i \propto i^c, \quad c = 2p + 1 \tag{5.12}$$

Here, the index $i \in \{1, N_\omega\}$ denotes the index of the sampled δ-functions (note that $p = -0.5$ corresponds to $c = 0$, i.e. equal amplitudes). In both the $p = 0.5$ and $p \neq -0.5$ cases, the spectrum will only exhibit the power-law divergent form very close to the edge; away from this edge, the QMC-generated data instead dictates the exact shape of the monotonically decaying tail. Therefore, it is not necessary to varying the amplitudes according to Eq. (5.12) for all N_ω of the δ-functions, and the instead a cross-over to equal amplitudes is implemented at index $i = n_0$. The exact location of n_0 is determined by the QMC-generated data, and is thus sampled along with the other updates of the spectrum within the SAC program. Figure 5·9 shows an example of how the amplitudes may vary with ω_i for an exponent $p = -3/4$, which is what was used for the for the diagram in Fig. 3·1(e). The cross-over point for this example was $n_0/N_\omega = 0.50$, meaning that half of the amplitudes vary according to Eq. (5.12).

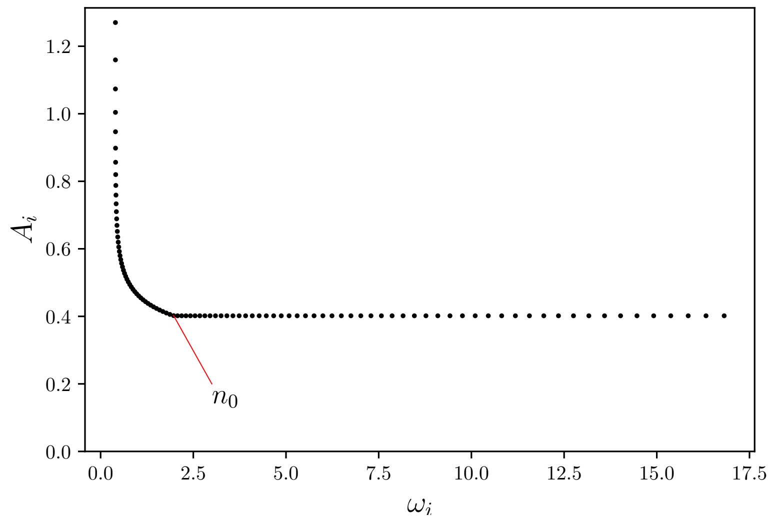

Figure 5·9: Variable amplitudes used to produce the δ-function diagram in Fig. 3·1(d). For an exponent $p = -0.75$, the amplitudes increase as the edge is approached. The location of the equal amplitude cross-over point, n_0, is marked.

Because of the sampling of the crossover point, one has to be slightly more careful when analyzing the results of a scan over p. If the p value that minimizes the χ^2 value is accompanied by a very small value of n_0, this implies that this optimal exponent is "fictitious." In this case, the δ-functions transition very quickly to the equal amplitude form and the asymptotic divergence $(\omega - \omega_0)^p$ is only realized is a very small region of the spectrum. In contrast, if the optimal exponent has a n_0 value that is close to one, the amplitude form is applied to the majority of the δ-functions and is thus realized for a large portion of the spectrum. This suggests that the accompanying low χ^2 value indeed indicates a good statistical fit for a spectrum with this specific divergent edge form.

There are no definite rules for determining the optimal exponent under these considerations, but Fig. 5·10 provides two characteristic examples of how one may use the χ^2 and n_0 values from a scan over p to determine the optimal exponent. Panels (a) and (b) show the χ^2 and n_0 values versus $-p$ for $S(q = \pi/2, \omega)$. In this case, the χ^2 value is nearly constant for $-p < 0.50$ and rapidly increases when $-p$ exceeds 0.50. The value of n_0 is small (less than ~ 0.30 and as small as 0.02 for $-p = 0.40$) during this constant χ^2 platform, which explains this behavior. The low values of n_0 imply that the edge is, for the most part, of the form $(\omega - \omega_0)^{-1/2}$, which is corroborated by the fact that the χ^2 value when p is exactly -0.50 is roughly equal to the χ^2 values for $-p < 0.50$.

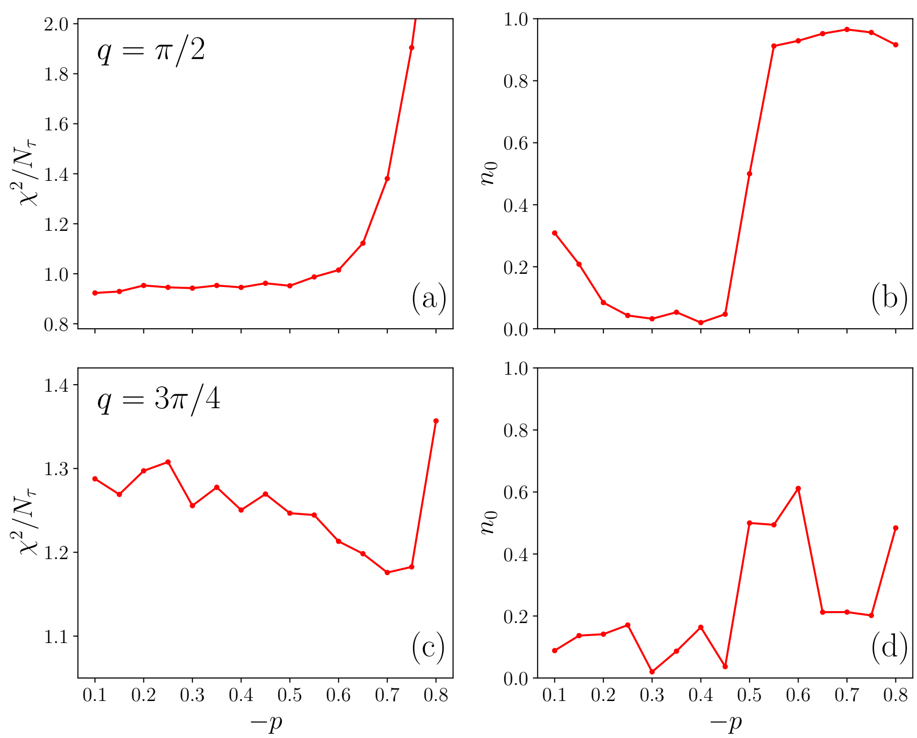

Figure 5·10: Results of p scan for the $S = 1/2$ Heisenberg chain. Panels (a) and (c) show the χ^2 value versus p ($\Delta p = 0.05$) for $q = \pi/2$ and $q = 3\pi/4$, respectively and panels (b) and (d) show the corresponding values of n_0.

When $-p$ exceeds 0.50, the value of n_0 jumps close to one. Since the χ^2 value also begins to increase in this region, $-p > 0.50$ can be ruled out as optimal exponent values. The combination of these two trends strongly imply that $p = -0.50$ is the true exponent, which is to be fully expected for dynamic spin structure factor of the Heisenberg chain, according to the BA solution [24, 25, 99].

Panels (c) and (d) show the results of the p scan for $S(q = 3\pi/4, \omega)$. Just as for $q = \pi/2$, n_0 is small for $-p < 0.50$ (less than ~ 0.20 and as small as 0.02 for $-p = 0.30$). However, in contrast to $q = \pi/2$, this range of p values does not coincide with the χ^2 minimum. Therefore, we can say with certainty that the optimal exponent does not lie within this range. In this case, there is a clear χ^2 minimum at $-p = 0.70$, but the accompanying value of n_0 is small, ~ 0.20. The n_0 value for $-p = 0.60$ is much larger, around 0.60, and the χ^2 value is within 4 % of the overall minimum value. This difference in the χ^2 is likely too small to say with certainty that $p = -0.70$ is really preferable over $p = -0.60$, so considering the much larger n_0 value, we conclude that the true optimal exponent for $q = 3\pi/4$ is $p = -0.60$. This is further supported by the very good agreement between the SAC spectra with $p = -0.60$ and the BA solution, shown in Fig. 3·5(f).

Cross Validation Results

For this test, we performed the cross validation process outlined above for three parameterizations (unconstrained, single-peak, and edge). We also rotated which of the $K + 1 = 21$ sets of QMC-generated data was used for sampling, and which ones were used for validation, and averaged the results over all of the rotations. We found that this extra step was necessary given the quality and quantity of the QMC-generated data at hand. In this case, the alternative cross validation method (50/50 splits, repeated many times in a bootstrap fashion) may actually be preferable, for this very reason, which we will touch on at the end of this section.

Fig. 5·11 shows the results of this test. The the normalized sampling (panels (b) and (e)) and validation (panels (c) and (f)) χ^2 values are plotted versus Θ for each parameterization: unconstrained (red), single-peak (blue), and edge (green). Just as with the two unconstrained sampling parameterizations, the entropy contents of the spectra defined with each constrained sampling parameterization are different, so to better compare the three parameterizations, we plot the validation χ^2 versus the sampling χ^2 as well, shown in panels (a) and (d).

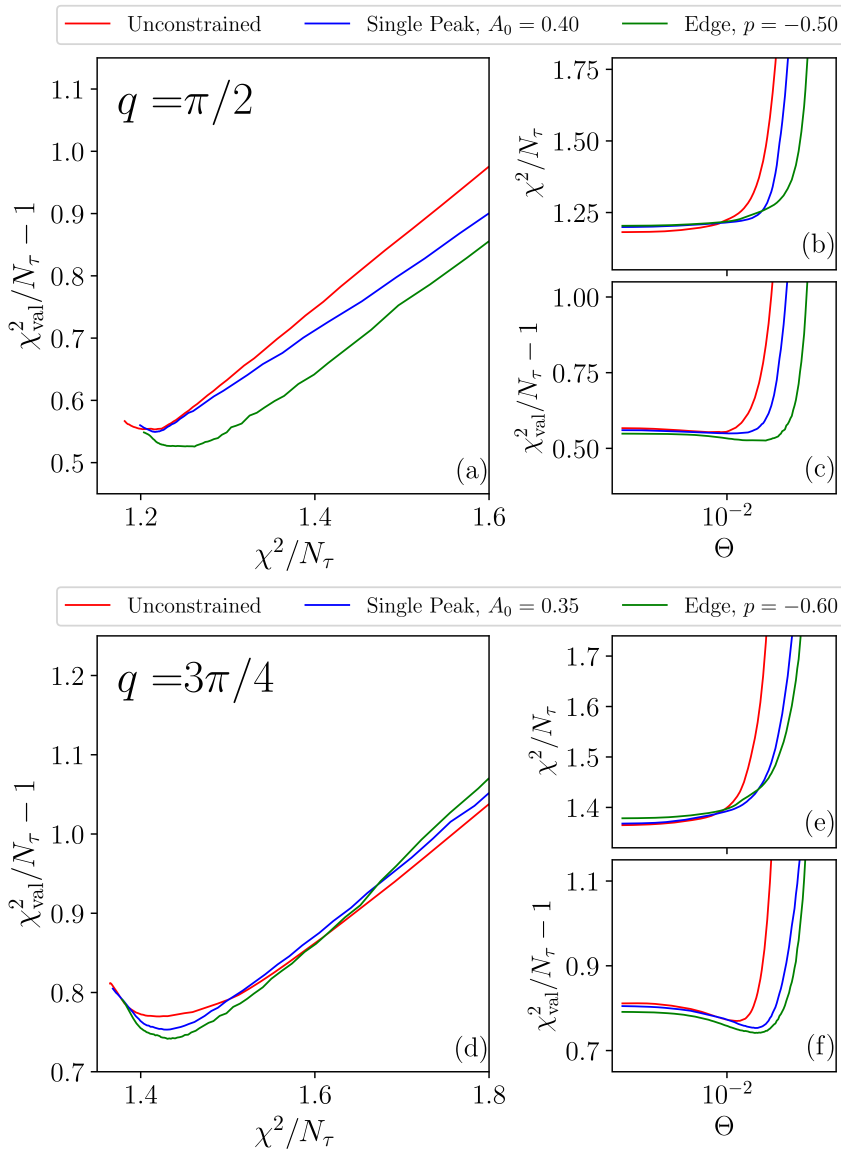

Figure 5·11: Cross validation results for $S(q = \pi/2, \omega)$ and $S(q = 3\pi/4, \omega)$ of the $S = 1/2$ Heisenberg chain ($L = 512$, $\beta = 1024$). In panels (b) and (e), the sampling χ^2 values, normalized by the number of τ points and averaged over all cross validations runs, is plotted versus the sampling temperature Θ. In panels (c) and (f), the validation χ^2 values, also normalized by the number of τ points (50 in this case) and averaged over all cross validations runs, is plotted versus the sampling temperature Θ. We have subtracted the background constant of one from the validation χ^2 to account for the term X_2 in Eq.(5.6). In panels (a) and (d), we plot the validation χ^2 as a function of the sampling χ^2, to allow for direct comparison of the three parameterizations at a fixed χ^2 value.

While it is difficult to glean any information about the relative performance of the three parameterizations from the sampling χ^2, the validation χ^2 paints a much different picture. At each temperature, the the edge parameterization has the lowest validation χ^2, followed by the single-peak case, and then unconstrained sampling. This would suggest that the edge is indeed the most appropriate parameterization, in agreement with the BA solution for this model. This conclusion can also be made when comparing the values of the validation χ^2 at fixed χ^2. For $q = \pi/2$, the edge clearly performs the best of the three parameterizations during the entire annealing run. For $q = 3\pi/4$, the validation χ^2 for the edge parametrization only dips below that of the other two at the end of the annealing run, but clearly reaches the lowest minimum value. We thus find that it is important take into consideration both the minimum value of each validation χ^2 curve and the path the curves take when comparing the various parameterizations.

We now compare the corresponding spectra, shown originally in Chapter 3, but reproduced here in Fig. 5·12, to see if the hierarchy suggested by the validation χ^2 is borne out in the results of the analytic continuation process. As reflected by the relative differences between the validation χ^2 curves, the single-peak and edge spectra share similar features, such as lower bounds near $\omega = 1.5$ ($q = \pi/2$) and $\omega = 1.0$ ($q = 3\pi/4$), while the unconstrained sampling produces spectra with considerable spectral weight below the edge. The BA results [1], shown in black, agree very well with the edge spectra (with edges that nearly perfectly align), moderately well for the peak spectra (but with some large deviations), and very poorly with the unconstrained spectra, supporting the results of our cross validation test. Here, it should be noted that the BA spectra have been subject to broadening, and furthermore, the calculation of these spectra only take into consideration two-spinon and four-spinon contributions [26]. As such, it is known that a few percent of the total spectral weight is missing, which could at least partially explain the minor deviations between the two spectra in Fig. 5·12(c) and (f).

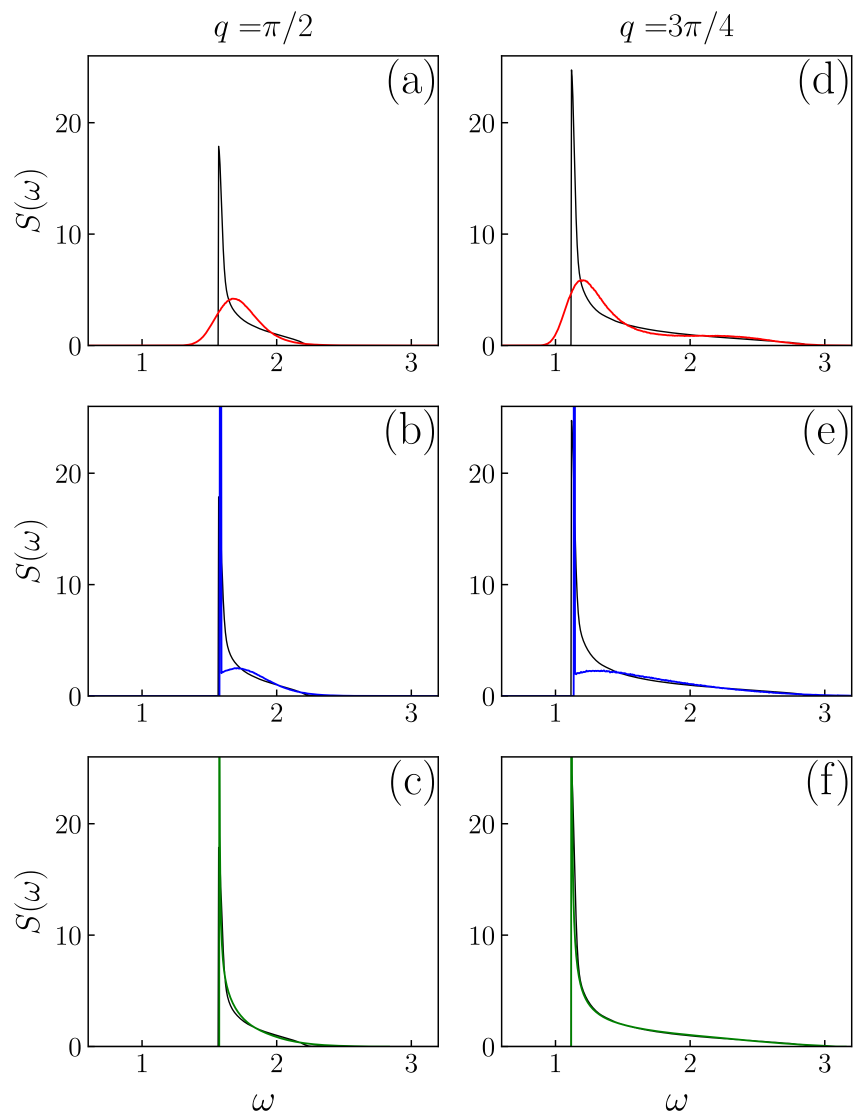

Figure 5·12: Reprint of Fig. 3·5 for reader's convenience.

The fact that the edge outperforms the unconstrained and single-peak parameterization strongly supports the viability of cross validation as a tool for model selection. Considering the fact that for a chain of length 512, the exact $T = 0$ spectrum likely only contains ∼100 delta functions with significant weight, deduced by extrapolating from the results for smaller system sizes [141, 149], it is remarkable that our method is able correctly identify the true form of the spectral function in the thermodynamic limit. It is quite possible that if smaller system size was considered instead, a spectrum composed of a dominate peak plus a smooth continuation would actually be the best representation and the single-peak parameterization would perform the best. But this is not the case here, implying that finite size effects are sufficiently mitigated.

Cross Validation Results using Alternative Procedure

We tested the alternative procedure introduced at the end of Sec. 5.2 and present the results in Fig. 5·13. The same conclusions are drawn from both cross validation methods (cf. Fig. 5·11); the edge parameterization has the lowest validation χ^2 when compared to the single-peak and unconstrained methods, at a fixed value of the χ^2. The validation χ^2 curves for the peak and edge parameterizations ($q = \pi/2$) and unconstrained sampling ($q = 3\pi/4$) do not have well defined minima, which differs from the results in Fig. 5·11. Nevertheless, the results of this test are clearly inline with those presented in Sec. 5.5.1 and emphasize that the value of the validation χ^2 throughout the entire anneal, as well as its minimum value, are important to consider when preforming cross validation.

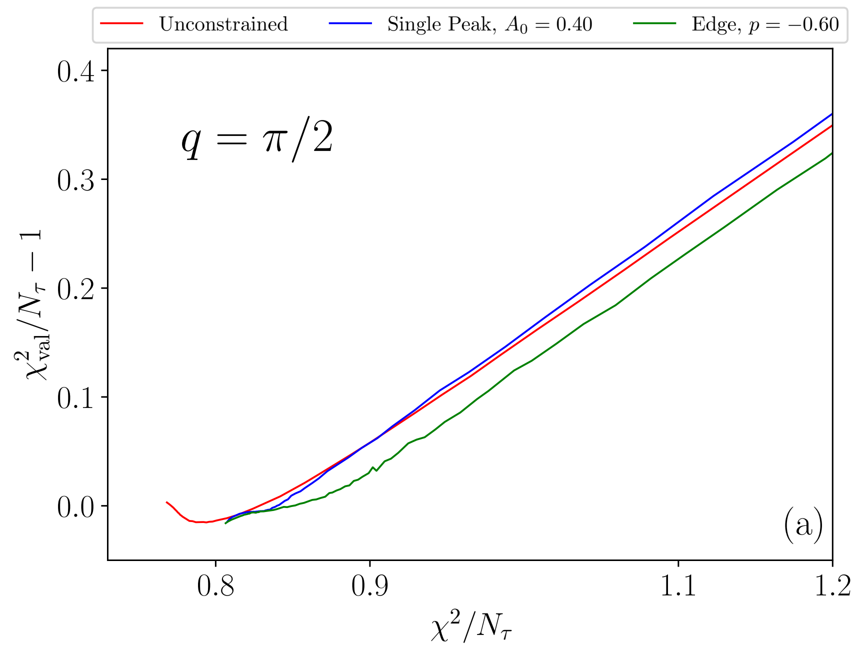

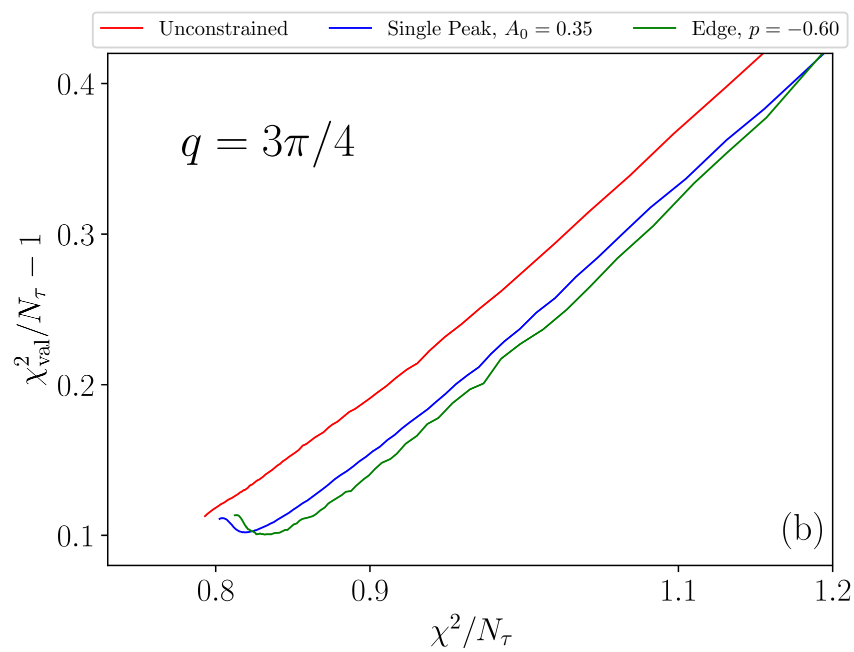

Figure 5·13: Cross validation results for $S(q = \pi/2, \omega)$ and $S(q = 3\pi/4, \omega)$ of the $S = 1/2$ Heisenberg chain ($L = 512$, $\beta = 1024$) using the alternative cross validation procedure. Panels (a) and (b) show validation χ^2 as a function of the sampling χ^2 for $q = \pi/2$ and $q = 3\pi/4$, respectively.

5.5.2 Heisenberg Chain with Long-Range Interactions

We now turn to a case where the exact features contained in the spectral function are unknown, the unfrustrated AFM Heisenberg chain with power-law decaying interactions [156, 75]:

$$H = \sum_{r=1}^{L/2} J_r \sum_{i=1}^{L} \boldsymbol{S}_i \cdot \boldsymbol{S}_{\boldsymbol{i+r}}, \tag{5.13}$$

where

$$J_r = G\frac{(-1)^{r-1}}{r^\alpha}, \quad G = \left(1 + \sum_{r=2}^{L/2} \frac{1}{r^\alpha}\right)^{-1}. \tag{5.14}$$

We note that the normalization of the coupling, G, chosen so that $\sum_r |J_r| = 1$, differs slightly from that used in previous related studies. Interest in this model is rooted in its possible connection to the 2D square lattice AFM Heisenberg model, as the long-range, staggered interactions allow for symmetry breaking and true long-range order to form in this 1D quantum magnet at $T = 0$, effectively increasing its dimensionality [151]. The excitation spectra of the long-range Heisenberg chain is a key ingredient in understanding the nature of its ground state, and the ability to resolve sharp features with the aforementioned constrained SAC sampling schemes may help clear up the long-standing uncertainties surrounding this model. An extensive study of this system will be published in a separate paper, where spectral functions are calculated in both the Néel ordered phase ($\alpha < \alpha_c = 2.22(1)$), as well as in the quasi long-range ordered (QLRO) phase ($\alpha > \alpha_c$), on the other side of the quantum phase transition (QPT) [152].

The transition between the Néel and QLRO phases provides an excellent application for cross validation, because of the uncertainty surrounding the nature of the ground state excitations. Deep in the QLRO regime, it is expected that the spectral functions contains a power-law edge, in light of the BA results discussed in Sec. 5.5.2. While deep in the Néel phase, the spectral function should contain a dominant magnon-peak, in analogy to the 2D AFM Heisenberg model [58, 22]. A previous study on a Heisenberg chain with similar long-range interactions only resolved spectral functions with sharp magnon peaks, followed by weak features at higher energy, when the system is in Néel phase [151]. This study used the time-dependent Density Matrix Renormalization Group method, which is known to produce artificially broadened peaks due to limits on the width of the real-time window used in the calculations. Due to the uncertainties introduced by this method, the nature of the higher-energy continuum was not examined in detail. The δ-function peak SAC parameterization could potentially provide a significant improvement to these results, and the application of the edge parameterization could provide completely new insights into this model.

Here, we use cross validation to test three SAC parameterizations (unconstrained, single-peak, and edge) in the Néel ordered regime ($\alpha = 1.5$), in the Néel, but close to the QPT ($\alpha = 2.0$), and on the QLRO side of the QPT ($\alpha = 3.0$) for the spectral function of the operator $O = S^z_{q=\pi/2}$. In Ref. [152], a broader set of SAC parameterizations are tested and all momenta are examined, but here, we will only focus on this single momentum and the three parameterizations discussed thus far.

Figure 5·14 shows the results of the cross validation procedure for this model. Here, we again used QMC data generated using the SSE method, for a system of length $L = 256$ at inverse temperature $\beta = 256$. The number of validation data sets used in this run was $K = 20$, which was chosen to limit the error level to $\sigma = 10^{-5}$. Just as before, the values of a_0 and p used in these tests were identified using scans [152].

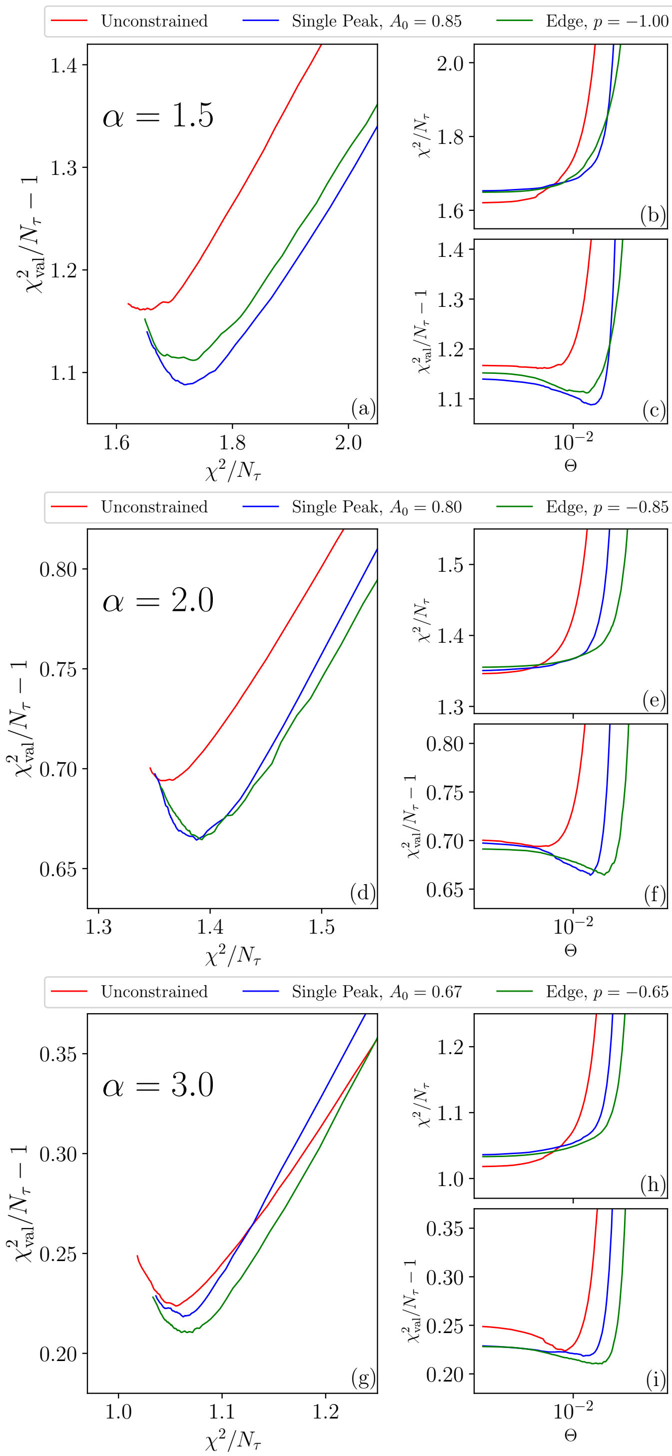

Figure 5·14: Cross validation results for the $S = 1/2$ Heisenberg chain with long-range interactions ($L = 256$, $\beta = 256$), for $\alpha = 1.5$ (panels (a)-(c)), 2.0 (panels (d)-(f)), and 3.0 (panels (g)-(i)). The spectral function considered here is for the operator $O = S^z_{q=\pi/2}$. Panels and set up are the same as in Fig. 5·11.

For all values of α considered here, the unconstrained sampling expectedly performs the worst of the three parameterizations. It is known that in both the large and small α limits, the spectral function contains a sharp edge feature, so it is likely that for intermediate values of α, the spectral function also contains a low energy edge of some form. It is most instructive, however, to look at the behavior of the single-peak and edge curves as α is increased. For $\alpha = 1.5$, the long-range Heisenberg chain exhibits Néel order and will thus host coherent, spin-wave excitations, albeit with anomalous dispersion [156, 76, 152]. We thus expect the spectral function to contain a dominant magnon peak, as was resolved using the single-peak parameterization in the 2D square lattice Heisenberg model [123]. Accordingly, the single-peak curve in Fig. 5·14(a) (blue) reaches the lowest validation χ^2 minimum. The spectral functions produced using these two parameterizations for this value of α are shown in Fig. 5·15. While the two spectral edges align nearly perfectly, the distributions of the spectral weight above this lower bound differ. We attribute the observed difference in the validation χ^2 values to this difference and conclude that the spectral function for $\alpha = 1.5$ does not contain a monotonically decaying edge. We contrast that with the curves in Fig. 5·14(g). For $\alpha = 3.0$, the ground state of the long-range Heisenberg chain is much closer to that of the standard Heisenberg chain discussed in Sec. 5.5.1. This is reflected in the cross validation χ^2, where the edge parameterization instead reaches the lowest minimum value.

Finally, for $\alpha = 2.0$ the validation χ^2 does not appear to favor either of the single-peak or edge parameterizations due to the proximity to the QPT at α_c. As α is increased from, the SAC parameterization that has the lowest validation χ^2 switches over from the single-peak to the edge parameterization. This implies that there exists a crossover regime around the vicinity of α_c where cross validation is unable to identify a clear preference between the two spectral forms. The notion of a crossover between these two spectral forms is further supported by the smoothly increasing value of the optimal exponent p as α is reduced. The edge parameterization spectra diverging more sharply as the systems enters deeper into the Néel ordered phase can be considered a natural transition to the expected dominant δ-peak form. However, in the intermediate regime near α_c, it is unlikely that the spectral function can be described purely by one of these two sharp features.

Different values of both α and momentum q will be studied in Ref. [152], but these preliminary tests provide further support for the viability of cross validation as a tool for model selection in the analytic continuation of QMC-generated data.

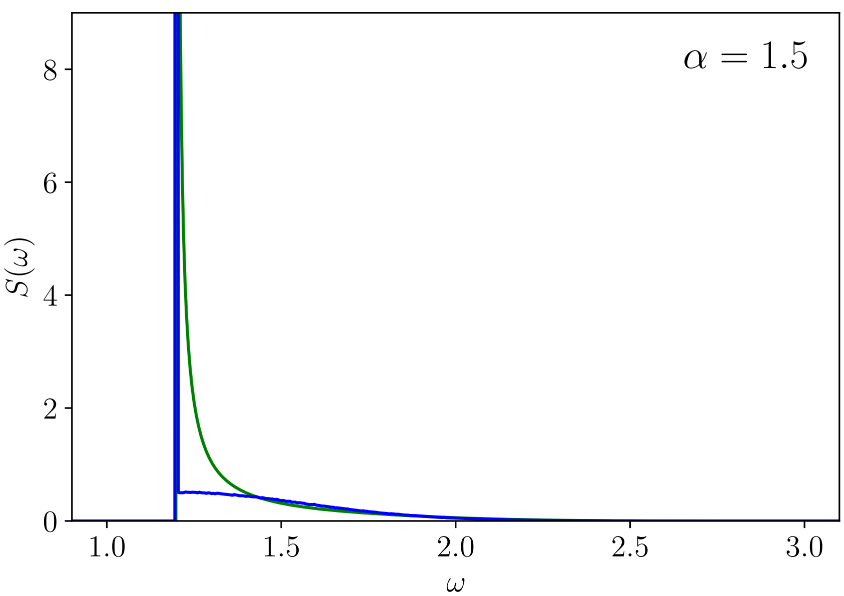

Figure 5·15: Comparison of the single-peak and edge parameterization spectral functions for the for the $S = 1/2$ Heisenberg chain with long-range interactions, corresponding to the cross validation results shown in Fig. 5·11 panels (a)-(c). For the single-peak parameterization (blue) we use the leading δ-function weight $a_0 = 0.85$ and $N_c = 2,000$ continuum δ-functions. For the edge parameterization (green), we use $N_\omega = 640$ and the exponent $p = -1.00$.

5.6 Summary and Discussion

We have explored the use of cross validation in the SAC method, as both a confirmation of the optimal Θ criterion, as well as a tool to select the most likely spectral parameterization for a given model. In our test cases using synthetic data, we found excellent agreement between the optimal Θ criterion Eq. (3.23) and the location of the cross validation minimum. For the more complex of the two test cases, the double Gaussian peak spectrum, the broad validation χ^2 minimum reflected the difficulties posed by the relatively sharp features contained in this spectrum. In principle this spectrum can be resolved using SAC, but the unconstrained sampling introduces entropic distortions that produce noticeable deviations from the exact result. As a result, all sampling temperatures within the broad minimum produce acceptable representations of the exact spectrum, which we demonstrated by using a Θ value one standard deviation below the average minimum value. Interestingly, at this lower sampling temperature some features of the spectrum were actually reproduced better. Since using cross validation to fix Θ would be a very poor use of computational resources, it is very encouraging that the optimal Θ criterion agrees so well with the the results of the validation procedure.

A promising, and more decisive, use of cross validation is in model selection. The different entropy contents of the two unconstrained SAC parameterizations result in differing abilities to reproduce a given spectrum depending on its exact features. In some cases, it it advantageous to use the equal amplitudes parameterization, and in others the variable amplitudes parameterization performs substantially better. In previous studies, this was only able to be shown retrospectively, when comparing the SAC results to an artificial spectrum used to generate synthetic data. Here, we showed that cross validation can indeed act as a tool for model selection when comparing the two different unconstrained SAC parameterizations. This will be extremely valuable in practice, as often times these two parameterizations produce considerably different results, which can be even more dramatic when other parameterizations are considered, such as the fixed grid method [124] or the orthogonal polynomial representation [108]. When there is no artificial spectrum to compare to, it is impossible to determine which of these spectra is best.

Cross validation as a tool for model selection was also applied to the constrained SAC parameterizations. While the recent developments in the constrained sampling schemes have allowed for the resolution of spectra with sharp features, the ill-posed nature of the analytic continuation problem still presents obstacles for the implementation of these new methods. When many different spectra give acceptable χ^2 values, cross validation is the least biased way to determine which parameterization is most likely, and our results for the $S = 1/2$ AFM Heisenberg chain suggest that cross validation indeed can accomplish this task. Cross validation should be an extremely valuable tool in the study of highly correlated, many-body systems if it can be used to differentiate quantum phases of matter via their spectral features. Further tests on other systems with known spectral features will be very instructive.

The test on the Heisenberg chain with long-range interactions also produced promising results, correctly identify which side of the QPT the system was on based on the spectral features alone. While the exact features are only known with certainty in the limiting cases of α, $\alpha \to 0$ and $\alpha \to \infty$, it is likely that these features persists throughout either side of the QPT. Further tests at more values of α and q are still needed, and some will be presented in Ref. [152], but as a general diagnostic tool, cross validation shows much promise in the practice of numerical analytic continuation of QMC-generated data.

Our cross validation scheme is not only applicable to SAC, but can also be applied other numerical analytic continuation methods, such as MEM. Based on our tests, a promising application of cross validation would be in the selection of the default model used in MEM. While not exactly analogous to the SAC parameterization, there are similarities between these two user-controlled inputs, such as their effects on the shape of the output spectra. Often times, perturbation theory is used to select the MEM default model, where some parameter in the perturbative solution is optimized to select the "best" spectra [62]. Cross validation could be used to aide in this optimizations process or to compare entirely different default models, perhaps from different perturbative solutions of the model under consideration.

Chapter 6

Conclusions

In this thesis, we explored the numerical study of quantum many-body systems, focusing two main classes of observables: statics and dynamics.

In the discussion of static observables, our most notable results were the extraction of scaling dimensions, Δ. Defined as the decay exponent of the two-point correlator in a conformal field theory, in the context of lattice models, scaling dimensions help us characterize critical points and identify the universality class of their classical or quantum phase transitions. As explored in Chapter 2, scaling dimensions can be used as keys to connect a critical lattice model to its corresponding continuum field theory, and this connection can in turn enable a deeper understanding of the lattice model itself. At first glance, the dimer order parameter of the FFTFIM seems to be the natural tool to characterize the ground state, as it corresponds to a visual and tangible description of the ordered phase. However, the quantum critical scaling revealed its secondary nature—its scaling dimension was that of a secondary order parameter in the field-theoretical description. This discovery is what enabled us to define a proper primary order parameter, fully understand the 8-fold degeneracy of the ground state, and thus map out the model's complex finite temperature phase diagram.

Scaling dimensions also have practical use in the study of out-of-equilibrium dynamics, as explored in Sec. 2.7. These are key pieces of the puzzle when analyzing the results of annealing protocols, and are crucial in extracting the dynamic exponent z.

As the field of quantum computing continues to develop, the importance of numerical simulation as a benchmarking tool will continue to grow. Comprehensive studies of building block models like the FFTFIM will thus be pivotal in the verification of this new and exctiting technology.

When discussing dynamics, one more typically is referring to the frequency dependent information contained in the dynamic response function. This class of observables is the primary focus of Chapters 3-5. After laying out of theory, framework, and my Julia implementation of the SAC method in Chapter 3, in Chapter 4 we explored its application to one of the quintessential models of condensed matter physics, the 2D Hubbard model. By leveraging the new capabilities of the SAC method, we were able to unveil fine details about the DOS of the 2D Hubbard model, that could not have been resolved using any other numerical analytical continuation method. This exciting advancement provided more pieces to the superconducting puzzle that is the cuprates; the reason for the finite doping level needed to observe superconducting may be contained in this minimal model. Novel results like these highlight the ongoing need for the advancement and study of computational physics methodology.

This is the motivation behind the development of the cross validation method presented in Chapter 5. SAC has enabled the ability to resolve spectral functions with fine and sharp structure, but without prior knowledge to motivate their existence, it is challenging to determine with certainty if they are genuine features of the underlying spectrum. Our cross validation methodology will thus be important for SAC to have a broad and significant impact in the field. This has already been demonstrated with the example discussed in Sec. 5.5.2, the long-range Heisenberg chain, and has led to exciting new insights for this yet to be fully understood model [152].

List of Journal Abbreviations

Acta Phys. Polon. B	....	Acta Physica Polonica B
J. Phys. France	....	Journal de Physique
Phys. Rep.	....	Physics Reports
Phys. Rev.	....	Physical Review
Phys. Rev. B	....	Physical Review B
Phys. Rev. E	....	Physical Review E
Phys. Rev. Lett.	....	Physical Review Letters
Phys. Rev. X	....	Physical Review X
Rev. Mod. Phys	....	Reviews of Modern Physics
SciPost Phys.	....	SciPost Physics

Bibliography

[1] BA results from Ref. [25] provided by J.-S. Caux (private communication).

[2] Amnon Aharony, R. J. Birgeneau, C. W. Garland, Y.-J. Kim, V. V. Lebedev, R. R. Netz, and M. J. Young. Critical behavior of the structure factor for higher harmonics in density wave systems. *Phys. Rev. Lett.*, 74:5064–5067, Jun 1995.

[3] M. Aichhorn, E. Arrigoni, Z. B. Huang, and W. Hanke. Superconducting gap in the hubbard model and the two-gap energy scales of high-T_c cuprate superconductors. *Phys. Rev. Lett.*, 99:257002, Dec 2007.

[4] Fabien Alet, Jesper Lykke Jacobsen, Grégoire Misguich, Vincent Pasquier, Frédéric Mila, and Matthias Troyer. Interacting classical dimers on the square lattice. *Phys. Rev. Lett.*, 94:235702, Jun 2005.

[5] Ammar Ali, Hanjing Xu, William Bernoudy, Alberto Nocera, Andrew D. King, and Arnab Banerjee. Quantum quench dynamics of geometrically frustrated ising models. *Nature Communications*, 15(1):10756, Dec 2024.

[6] F. F. Assaad and M. Imada. Insulator-metal transition in the one- and two-dimensional hubbard models. *Phys. Rev. Lett.*, 76:3176–3179, Apr 1996.

[7] F. F. Assaad and M. Imada. Doping-induced metal-insulator transition in two-dimensional hubbard $t-u$ and extended hubbard $t-u-w$ models. *Phys. Rev. B*, 58:1845–1852, Jul 1998.

[8] F. F. Assaad and M. Imada. Unusually flat hole dispersion relation in the two-dimensional hubbard model and restoration of coherence by addition of pair-hopping processes. *The European Physical Journal B - Condensed Matter and Complex Systems*, 10(4):595–598, Jun 1999.

[9] S. Badoux, S. A. A. Afshar, B. Michon, A. Ouellet, S. Fortier, D. LeBoeuf, T. P. Croft, C. Lester, S. M. Hayden, H. Takagi, K. Yamada, D. Graf, N. Doiron-Leyraud, and Louis Taillefer. Critical doping for the onset of fermi-surface reconstruction by charge-density-wave order in the cuprate superconductor $\mathrm{la}_{2-x}\mathrm{sr}_x\mathrm{cuo}_4$. *Phys. Rev. X*, 6:021004, Apr 2016.

[10] K. S. D. Beach. Identifying the maximum entropy method as a special limit of stochastic analytic continuation. 2004. arXiv:cond-mat/0403055 [cond-mat.str-el].

[11] Dominic Bergeron and A.-M. S. Tremblay. Algorithms for optimized maximum entropy and diagnostic tools for analytic continuation. *Phys. Rev. E*, 94:023303, Aug 2016.

[12] K. Binder. Critical properties from monte carlo coarse graining and renormalization. *Phys. Rev. Lett.*, 47:693–696, Aug 1981.

[13] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[14] Daniel Blankschtein, M. Ma, and A. Nihat Berker. Fully and partially frustrated simple-cubic ising models: Landau-ginzburg-wilson theory. *Phys. Rev. B*, 30:1362–1365, Aug 1984.

[15] Daniel Blankschtein, M. Ma, A. Nihat Berker, Gary S. Grest, and C. M. Soukoulis. Orderings of a stacked frustrated triangular system in three dimensions. *Phys. Rev. B*, 29:5250–5252, May 1984.

[16] Henk W. J. Blöte and M. Peter Nightingale. Antiferromagnetic triangular ising model: Critical behavior of the ground state. *Phys. Rev. B*, 47:15046–15059, Jun 1993.

[17] M. Brunner, F. F. Assaad, and A. Muramatsu. Single hole dynamics in the one-dimensional - model. *The European Physical Journal B - Condensed Matter and Complex Systems*, 16(2):209–212, Jul 2000.

[18] N. Bulut, D. J. Scalapino, and S. R. White. Electronic properties of the insulating half-filled hubbard model. *Phys. Rev. Lett.*, 73:748–751, Aug 1994.

[19] J. K. Burton, Jr., and Christopher Lee Henley. A constrained potts antiferromagnet model with an interface representation. *Journal of Physics A*, 30:8385–8413, 1997.

[20] Pasquale Calabrese and Pietro Parruccini. Harmonic crossover exponents in $o(n)$ models with the pseudo-ϵ expansion approach. *Phys. Rev. B*, 71:064416, Feb 2005.

[21] Massimo Campostrini, Martin Hasenbusch, Andrea Pelissetto, and Ettore Vicari. Theoretical estimates of the critical exponents of the superfluid transition in ^{4}He by lattice methods. *Phys. Rev. B*, 74:144506, Oct 2006.

[22] C. M. Canali and Mats Wallin. Spin-spin correlation functions for the square-lattice heisenberg antiferromagnet at zero temperature. *Phys. Rev. B*, 48:3264–3280, Aug 1993.

[23] John Cardy. *Scaling and Renormalization in Statistical Physics*. Cambridge Lecture Notes in Physics. Cambridge University Press, 1996.

[24] Jean-Sébastien Caux and Jean Michel Maillet. Computation of dynamical correlation functions of heisenberg chains in a magnetic field. *Phys. Rev. Lett.*, 95:077201, Aug 2005.

[25] Jean-Sébastien Caux, Rob Hagemans, and Jean Michel Maillet. Computation of dynamical correlation functions of heisenberg chains: the gapless anisotropic regime. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09):P09003, sep 2005.

[26] Jean-Sébastien Caux, Rob Hagemans, and Jean Michel Maillet. Computation of dynamical correlation functions of heisenberg chains: the gapless anisotropic regime. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09):P09003, sep 2005.

[27] Anushya Chandran, Amir Erez, Steven S. Gubser, and S. L. Sondhi. Kibble-zurek problem: Universality and the scaling limit. *Phys. Rev. B*, 86:064304, Aug 2012.

[28] Shai M. Chester, Walter Landry, Junyu Liu, David Poland, David Simmons-Duffin, Ning Su, and Alessandro Vichi. Carving out ope space and precise o(2) model critical exponents. *Journal of High Energy Physics*, 2020(6):142, Jun 2020.

[29] I Chuang, R Durrer, N Turok, and B Yurke. Cosmology in the laboratory: defect dynamics in liquid crystals. *Science*, 251(4999):1336–1342, March 1991.

[30] Tommaso Coletta, Sergey E. Korshunov, and Frédéric Mila. Semiclassical evidence of columnar order in the fully frustrated transverse-field ising model on the square lattice. *Phys. Rev. B*, 90:205109, Nov 2014.

[31] Jacek Dziarmaga. Dynamics of a quantum phase transition and relaxation to a steady state. *Advances in Physics*, 59(6):1063–1189, 2010.

[32] R. Eder and Y. Ohta. Doping dependence of the spectral function in the t-j model. *Physica C: Superconductivity*, 263(1):104–106, 1996. Proceedings of the International Symposium on Frontiers of High - Tc Superconductivity.

[33] R. Eder, Y. Ohta, and T. Shimozato. Validity of the rigid-band picture for the t-j model. *Phys. Rev. B*, 50:3350–3355, Aug 1994.

[34] Vladislav Efremkin, Jean-Louis Barrat, Stefano Mossa, and Markus Holzmann. Time correlation functions for quantum systems: Validating Bayesian approaches for harmonic oscillators and beyond. *The Journal of Chemical Physics*, 155(13):134108, 10 2021.

[35] Michael E. Fisher. Statistical mechanics of dimers on a plane lattice. *Phys. Rev.*, 124:1664–1672, Dec 1961.

[36] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum monte carlo simulations of solids. *Rev. Mod. Phys.*, 73:33–83, Jan 2001.

[37] Eduardo Fradkin, Steven A. Kivelson, and John M. Tranquada. Colloquium: Theory of intertwined orders in high temperature superconductors. *Rev. Mod. Phys.*, 87:457–482, May 2015.

[38] Hermann Freire, Vanuildo S. de Carvalho, and Catherine Pépin. Renormalization group analysis of the pair-density-wave and charge order within the fermionic hot-spot model for cuprate superconductors. *Phys. Rev. B*, 92:045132, Jul 2015.

[39] Sebastian Fuchs, Thomas Pruschke, and Mark Jarrell. Analytic continuation of quantum monte carlo data by stochastic analytical inference. *Phys. Rev. E*, 81:056701, May 2010.

[40] Khaldoon Ghanem and Erik Koch. Average spectrum method for analytic continuation: Efficient blocked-mode sampling and dependence on the discretization grid. *Phys. Rev. B*, 101:085111, Feb 2020.

[41] Khaldoon Ghanem and Erik Koch. Extending the average spectrum method: Grid point sampling and density averaging. *Phys. Rev. B*, 102:035114, Jul 2020.

[42] Khaldoon Ghanem and Erik Koch. Generalized maximum entropy methods as limits of the average spectrum method. *Phys. Rev. B*, 108:L201107, Nov 2023.

[43] J. E. Gubernatis, Mark Jarrell, R. N. Silver, and D. S. Sivia. Quantum monte carlo simulations and maximum entropy: Dynamics from imaginary-time data. *Phys. Rev. B*, 44:6011–6029, Sep 1991.

[44] Martin C. Gutzwiller. Effect of correlation on the ferromagnetism of transition metals. *Phys. Rev. Lett.*, 10:159–162, Mar 1963.

[45] Martin Hasenbusch. Monte carlo study of an improved clock model in three dimensions. *Phys. Rev. B*, 100:224517, Dec 2019.

[46] Martin Hasenbusch and Ettore Vicari. Anisotropic perturbations in three-dimensional o(n)-symmetric vector models. *Phys. Rev. B*, 84:125136, Sep 2011.

[47] Kazumasa Hattori and Hirokazu Tsunetsugu. Antiferro quadrupole orders in non-kramers doublet systems. *Journal of the Physical Society of Japan*, 83(3):034709, 2014.

[48] Kazumasa Hattori and Hirokazu Tsunetsugu. Classical monte carlo study for antiferro quadrupole orders in a diamond lattice. *Journal of the Physical Society of Japan*, 85(9):094001, 2016.

[49] Sumner N. Hearth, Siddhardh C. Morampudi, and Chris R. Laumann. Quantum orders in the frustrated ising model on the bathroom tile lattice. *Phys. Rev. B*, 105:195101, May 2022.

[50] Christopher L. Henley. Ordering due to disorder in a frustrated vector antiferromagnet. *Phys. Rev. Lett.*, 62:2056–2059, Apr 1989.

[51] Christopher L. Henley. Relaxation time for a dimer covering with height representation. *Journal of Statistical Physics*, 89(3):483–507, Nov 1997.

[52] Motoaki Hirayama, Youhei Yamaji, Takahiro Misawa, and Masatoshi Imada. Ab initio effective hamiltonians for cuprate superconductors. *Phys. Rev. B*, 98:134501, Oct 2018.

[53] J. E. Hirsch. Two-dimensional hubbard model: Numerical simulation study. *Phys. Rev. B*, 31:4403–4419, Apr 1985.

[54] J. E. Hirsch and S. Tang. Antiferromagnetism in the two-dimensional hubbard model. *Phys. Rev. Lett.*, 62:591–594, Jan 1989.

[55] Yun-Da Hsieh, Ying-Jer Kao, and Anders W Sandvik. Finite-size scaling method for the berezinskii–kosterlitz–thouless transition. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(09):P09001, sep 2013.

[56] J. Hubbard and Brian Hilton Flowers. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 276(1365):238–257, 1963.

[57] S Hüfner, MA Hossain, A Damascelli, and GA Sawatzky. Two gaps make a high-temperature superconductor? *Reports on Progress in Physics*, 71(6):062501, 2008.

[58] Jun-ichi Igarashi. 1/s expansion for thermodynamic quantities in a two-dimensional heisenberg antiferromagnet at zero temperature. *Phys. Rev. B*, 46:10763–10771, Nov 1992.

[59] Masatoshi Imada, Atsushi Fujimori, and Yoshinori Tokura. Metal-insulator transitions. *Reviews of modern physics*, 70(4):1039, 1998.

[60] S. V. Isakov and R. Moessner. Interplay of quantum and thermal fluctuations in a frustrated magnet. *Phys. Rev. B*, 68:104409, Sep 2003.

[61] Rodolfo A. Jalabert and Subir Sachdev. Spontaneous alignment of frustrated bonds in an anisotropic, three-dimensional ising model. *Phys. Rev. B*, 44:686–690, Jul 1991.

[62] Mark Jarrell and J.E. Gubernatis. Bayesian inference and the analytic continuation of imaginary-time quantum monte carlo data. *Physics Reports*, 269(3):133–195, 1996.

[63] Ying Jiang and Thorsten Emig. String picture for a model of frustrated quantum magnets and dimers. *Phys. Rev. Lett.*, 94:110604, Mar 2005.

[64] Ying Jiang and Thorsten Emig. Ordering of geometrically frustrated classical and quantum triangular ising magnets. *Phys. Rev. B*, 73:104452, Mar 2006.

[65] Jorge V. José, Leo P. Kadanoff, Scott Kirkpatrick, and David R. Nelson. Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model. *Phys. Rev. B*, 16:1217–1241, Aug 1977.

[66] Junjiro Kanamori. Electron correlation and ferromagnetism of transition metals. *Progress of Theoretical Physics*, 30(3):275–289, 09 1963.

[67] T W B Kibble. Topology of cosmic domains and strings. *Journal of Physics A: Mathematical and General*, 9(8):1387, aug 1976.

[68] Andrew D. King et al. Observation of topological phenomena in a programmable lattice of 1,800 qubits. *Nature*, 560(7719):456–460, Aug 2018.

[69] Andrew D. King et al. Scaling advantage over path-integral monte carlo in quantum simulation of geometrically frustrated magnets. *Nature Communications*, 12(1):1113, Feb 2021.

[70] Andrew D King, Alberto Nocera, Marek M Rams, Jacek Dziarmaga, Roeland Wiersema, William Bernoudy, Jack Raymond, Nitin Kaushal, Niclas Heinsdorf, Richard Harris, et al. Beyond-classical computation in quantum simulation. *Science*, 388(6743):199–204, 2025.

[71] Andrew D King, Jack Raymond, Trevor Lanting, Richard Harris, Alex Zucca, Fabio Altomare, Andrew J Berkley, Kelly Boothby, Sara Ejtemaee, Colin Enderud, et al. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, 617(7959):61–66, 2023.

[72] Andrew D King, Sei Suzuki, Jack Raymond, Alex Zucca, Trevor Lanting, Fabio Altomare, Andrew J Berkley, Sara Ejtemaee, Emile Hoskinson, Shuiyuan Huang, et al. Coherent quantum annealing in a programmable 2,000 qubit ising chain. *Nature Physics*, 18(11):1324–1328, 2022.

[73] S. E. Korshunov. Finite-temperature phase transitions in the quantum fully frustrated transverse-field ising models. *Phys. Rev. B*, 86:014429, Jul 2012.

[74] J M Kosterlitz. The critical properties of the two-dimensional xy model. *Journal of Physics C: Solid State Physics*, 7(6):1046, mar 1974.

[75] Nicolas Laflorencie, Ian Affleck, and Mona Berciu. Critical phenomena and quantum phase transition in long range heisenberg antiferromagnetic chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(12):P12001, dec 2005.

[76] Nicolas Laflorencie, Ian Affleck, and Mona Berciu. Critical phenomena and quantum phase transition in long range heisenberg antiferromagnetic chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(12):P12001, dec 2005.

[77] Pablo Laguna and Wojciech Hubert Zurek. Density of kinks after a quench: When symmetry breaks, how big are the pieces? *Phys. Rev. Lett.*, 78:2519–2522, Mar 1997.

[78] Catia Lavalle, Michael Brunner, Fakher F. Assaad, and Alejandro Muramatsu. Dynamical properties of the t-j model. In Egon Krause and Willi Jäger, editors, *High Performance Computing in Science and Engineering 2000*, pages 143–154, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[79] M. Le Tacon, A. Sacuto, A. Georges, G. Kotliar, Y. Gallais, D. Colson, and A. Forget. Two energy scales and two distinct quasiparticle dynamics in the superconducting state of underdoped cuprates. *Nature Physics*, 2(8):537–543, Aug 2006.

[80] Patrick A. Lee, Naoto Nagaosa, and Xiao-Gang Wen. Doping a mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.*, 78:17–85, Jan 2006.

[81] Cheng-Wei Liu, Anatoli Polkovnikov, and Anders W. Sandvik. Dynamic scaling at classical phase transitions approached through nonequilibrium quenching. *Phys. Rev. B*, 89:054307, Feb 2014.

[82] Cheng-Wei Liu, Anatoli Polkovnikov, and Anders W. Sandvik. Quantum versus classical annealing: Insights from scaling theory and results for spin glasses on 3-regular graphs. *Phys. Rev. Lett.*, 114:147203, Apr 2015.

[83] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B*, 41:9301–9307, May 1990.

[84] Pankaj Mehta, Marin Bukov, Ching-Hao Wang, Alexandre G.R. Day, Clint Richardson, Charles K. Fisher, and David J. Schwab. A high-bias, low-variance introduction to machine learning for physicists. *Physics Reports*, 810:1–124, 2019. A high-bias, low-variance introduction to Machine Learning for physicists.

[85] Jia-Wei Mei, Shinji Kawasaki, Guo-Qing Zheng, Zheng-Yu Weng, and Xiao-Gang Wen. Luttinger-volume violating fermi liquid in the pseudogap phase of the cuprate superconductors. *Phys. Rev. B*, 85:134519, Apr 2012.

[86] A. J. Millis. Gaps and our understanding. *Science*, 314(5807):1888–1889, 2006.

[87] Takahiro Misawa and Masatoshi Imada. Quantum criticality around metal-insulator transitions of strongly correlated electron systems. *Physical Review B—Condensed Matter and Materials Physics*, 75(11):115121, 2007.

[88] A. S. Mishchenko, N. V. Prokof'ev, and B. V. Svistunov. Single-hole spectral function and spin-charge separation in the $t-j$ model. *Phys. Rev. B*, 64:033101, Jun 2001.

[89] A.S. Mishchenko and N. Nagaosa. Arpes spectra of polaron in the t–j model. *Journal of Physics and Chemistry of Solids*, 67(1):259–261, 2006. Spectroscopies in Novel Superconductors 2004.

[90] R Moessner. Magnets with strong geometric frustration. *Canadian Journal of Physics*, 79(11-12):1283–1294, 2001.

[91] R. Moessner and S. L. Sondhi. Ising models of quantum frustration. *Phys. Rev. B*, 63:224401, May 2001.

[92] R. Monaco, J. Mygind, and R. J. Rivers. Zurek-kibble domain structures: The dynamics of spontaneous vortex formation in annular josephson tunnel junctions. *Phys. Rev. Lett.*, 89:080603, Aug 2002.

[93] A. Moreo, D. Scalapino, and E. Dagotto. Phase separation in the hubbard model. *Phys. Rev. B*, 43:11442–11444, May 1991.

[94] A. Moreo, D. J. Scalapino, R. L. Sugar, S. R. White, and N. E. Bickers. Numerical study of the two-dimensional hubbard model for various band fillings. *Phys. Rev. B*, 41:2313–2320, Feb 1990.

[95] Ojiro Nagai, Seiji Miyashita, and Tsuyoshi Horiguchi. Ground state of the antiferromagnetic ising model of general spin s on a triangular lattice. *Phys. Rev. B*, 47:202–205, Jan 1993.

[96] Roland R. Netz and Amnon Aharony. Critical behavior of energy-energy, strain-strain, higher-harmonics,and similar correlation functions. *Phys. Rev. E*, 55:2267–2278, Mar 1997.

[97] B Nienhuis, H J Hilhorst, and H W J Blote. Triangular sos models and cubic-crystal shapes. *Journal of Physics A: Mathematical and General*, 17(18):3559, dec 1984.

[98] M. Opel, R. Nemetschek, C. Hoffmann, R. Philipp, P. F. Müller, R. Hackl, I. Tüttő, A. Erb, B. Revaz, E. Walker, H. Berger, and L. Forró. Carrier relaxation, pseudogap, and superconducting gap in high-T_c cuprates: A raman scattering study. *Phys. Rev. B*, 61:9752–9774, Apr 2000.

[99] R. G. Pereira, J. Sirker, J.-S. Caux, R. Hagemans, J. M. Maillet, S. R. White, and I. Affleck. Dynamical spin structure factor for the anisotropic spin-1/2 heisenberg chain. *Phys. Rev. Lett.*, 96:257202, Jun 2006.

[100] Anatoli Polkovnikov. Universal adiabatic dynamics in the vicinity of a quantum critical point. *Phys. Rev. B*, 72:161201, Oct 2005.

[101] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore. Colloquium: Nonequilibrium dynamics of closed interacting quantum systems. *Rev. Mod. Phys.*, 83:863–883, Aug 2011.

[102] R. Preuss, W. Hanke, and W. von der Linden. Quasiparticle dispersion of the 2d hubbard model: From an insulator to a metal. *Phys. Rev. Lett.*, 75:1344–1347, Aug 1995.

[103] Cyril Proust and Louis Taillefer. The remarkable underlying ground states of cuprate superconductors. *Annual Review of Condensed Matter Physics*, 10(Volume 10, 2019):409–429, 2019.

[104] Mingpu Qin, Chia-Min Chung, Hao Shi, Ettore Vitali, Claudius Hubig, Ulrich Schollwöck, Steven R. White, and Shiwei Zhang. Absence of superconductivity in the pure two-dimensional hubbard model. *Phys. Rev. X*, 10:031016, Jul 2020.

[105] Mingpu Qin, Thomas Schäfer, Sabine Andergassen, Philippe Corboz, and Emanuel Gull. The hubbard model: A computational perspective. *Annual Review of Condensed Matter Physics*, 13(Volume 13, 2022):275–302, 2022.

[106] Mingpu Qin, Hao Shi, and Shiwei Zhang. Benchmark study of the two-dimensional hubbard model with auxiliary-field quantum monte carlo method. *Phys. Rev. B*, 94:085103, Aug 2016.

[107] Yan Qi Qin, B. Normand, Anders W. Sandvik, and Zi Yang Meng. Amplitude mode in three-dimensional dimerized antiferromagnets. *Phys. Rev. Lett.*, 118:147207, Apr 2017.

[108] Wu Quan-Sheng, Wang Yi-Lin, Fang Zhong, and Dai Xi. Acceleration of the stochastic analytic continuation method via an orthogonal polynomial representation of the spectral function. *Chinese Physics Letters*, 30(9):090201, sep 2013.

[109] David R. Reichman and Eran Rabani. Analytic continuation average spectrum method for quantum liquids. *The Journal of Chemical Physics*, 131(5):054502, 08 2009.

[110] Daniel S. Rokhsar and Steven A. Kivelson. Superconductivity and the quantum hard-core dimer gas. *Phys. Rev. Lett.*, 61:2376–2379, Nov 1988.

[111] Subir Sachdev, Erez Berg, Shubhayu Chatterjee, and Yoni Schattner. Spin density wave order, topological order, and fermi surface reconstruction. *Phys. Rev. B*, 94:115147, Sep 2016.

[112] Subir Sachdev and Debanjan Chowdhury. The novel metallic states of the cuprates: Topological fermi liquids and strange metals. *Progress of Theoretical and Experimental Physics*, 2016(12):12C102, 11 2016.

[113] Anders W. Sandvik. Stochastic method for analytic continuation of quantum monte carlo data. *Phys. Rev. B*, 57:10287–10290, May 1998.

[114] Anders W. Sandvik. Stochastic series expansion method for quantum ising models with arbitrary interactions. *Phys. Rev. E*, 68:056701, Nov 2003.

[115] Anders W. Sandvik. Computational Studies of Quantum Spin Systems. *AIP Conference Proceedings*, 1297(1):135–338, 11 2010.

[116] Anders W. Sandvik. Constrained sampling method for analytic continuation. *Phys. Rev. E*, 94:063308, Dec 2016.

[117] Douglas J. Scalapino. Numerical simulations of the two-dimensional hubbard model: Dynamic properties. *Physica C: Superconductivity*, 185-189:104–113, 1991.

[118] Gabe Schumm, Hui Shao, Wenan Guo, Frédéric Mila, and Anders W. Sandvik. Primary and secondary order parameters in the fully frustrated transverse-field ising model on the square lattice. *Phys. Rev. B*, 109:L140408, Apr 2024.

[119] Gabe Schumm, Sibin Yang, and Anders W. Sandvik. Cross validation in stochastic analytic continuation. *Phys. Rev. E*, 110:055307, Nov 2024.

[120] Gabe Schumm, Shiwei Zhang, and Anders W. Sandvik. Single-particle dispersion and density of states of the half-filled 2d hubbard model, 2025. arXiv:2504.02614 [cond-mat.str-el].

[121] T. Senthil, Subir Sachdev, and Matthias Vojta. Fractionalized fermi liquids. *Phys. Rev. Lett.*, 90:216403, May 2003.

[122] Hui Shao, Wenan Guo, and Anders W. Sandvik. Monte carlo renormalization flows in the space of relevant and irrelevant operators: Application to three-dimensional clock models. *Phys. Rev. Lett.*, 124:080602, Feb 2020.

[123] Hui Shao, Yan Qi Qin, Sylvain Capponi, Stefano Chesi, Zi Yang Meng, and Anders W. Sandvik. Nearly deconfined spinon excitations in the square-lattice spin-1/2 heisenberg antiferromagnet. *Phys. Rev. X*, 7:041072, Dec 2017.

[124] Hui Shao and Anders W. Sandvik. Progress on stochastic analytic continuation of quantum monte carlo data. *Phys. Rep.*, 1003:1–88, 2023. Progress on stochastic analytic continuation of quantum Monte Carlo data.

[125] R. N. Silver, J. E. Gubernatis, D. S. Sivia, and M. Jarrell. Spectral densities of the symmetric anderson model. *Phys. Rev. Lett.*, 65:496–499, Jul 1990.

[126] R. N. Silver, D. S. Sivia, and J. E. Gubernatis. Maximum-entropy method for analytic continuation of quantum monte carlo data. *Phys. Rev. B*, 41:2380–2389, Feb 1990.

[127] John Stephenson. Ising-Model Spin Correlations on the Triangular Lattice. III. Isotropic Antiferromagnetic Lattice. *Journal of Mathematical Physics*, 11(2):413–419, 10 2003.

[128] Edward A. Stern. Rigid-band model of alloys. *Phys. Rev.*, 157:544–551, May 1967.

[129] R. B. Stinchcombe. *Phase Transitions*, pages 295–340. Springer US, Boston, MA, 1988.

[130] Olav F. Syljuåsen. Using the average spectrum method to extract dynamics from quantum monte carlo simulations. *Phys. Rev. B*, 78:174429, Nov 2008.

[131] Louis Taillefer. Scattering and pairing in cuprate superconductors. *Annual Review of Condensed Matter Physics*, 1(Volume 1, 2010):51–70, 2010.

[132] Kiyohisa Tanaka, W. S. Lee, D. H. Lu, A. Fujimori, T. Fujii, null, I. Terasaki, D. J. Scalapino, T. P. Devereaux, Z. Hussain, and Z.-X. Shen. Distinct fermi-momentum-dependent energy gaps in deeply underdoped bi2212. *Science*, 314(5807):1910–1913, 2006.

[133] Hirokazu Tsunetsugu, Takayuki Ishitobi, and Kazumasa Hattori. Quadrupole orders on the fcc lattice. *Journal of the Physical Society of Japan*, 90(4):043701, 2021.

[134] K. Vafayi and O. Gunnarsson. Analytical continuation of spectral data from imaginary time axis to real frequency axis using statistical sampling. *Phys. Rev. B*, 76:035115, Jul 2007.

[135] T. Valla, A. V. Fedorov, Jinho Lee, J. C. Davis, and G. D. Gu. The ground state of the pseudogap in cuprate superconductors. *Science*, 314(5807):1914–1916, 2006.

[136] C. N. Varney, C.-R. Lee, Z. J. Bai, S. Chiesa, M. Jarrell, and R. T. Scalettar. Quantum monte carlo study of the two-dimensional fermion hubbard model. *Phys. Rev. B*, 80:075116, Aug 2009.

[137] J Villain. Spin glass with non-random interactions. *Journal of Physics C: Solid State Physics*, 10(10):1717, may 1977.

[138] J Villain. Spin glass with non-random interactions. *Journal of Physics C: Solid State Physics*, 10(10):1717, may 1977.

[139] Villain, J., Bidaux, R., Carton, J.-P., and Conte, R. Order as an effect of disorder. *J. Phys. France*, 41(11):1263–1272, 1980.

[140] Ettore Vitali, Hao Shi, Mingpu Qin, and Shiwei Zhang. Computation of dynamical correlation functions for many-fermion systems with auxiliary-field quantum monte carlo. *Phys. Rev. B*, 94:085140, Aug 2016.

[141] Ling Wang and Hai-Qing Lin. Dynamic structure factor from real time evolution and exact correction vectors with matrix product states, 2019. arXiv:1901.07751 [cond-mat.str-el].

[142] Tatiana A. Webb, Michael C. Boyer, Yi Yin, Debanjan Chowdhury, Yang He, Takeshi Kondo, T. Takeuchi, H. Ikuta, Eric W. Hudson, Jennifer E. Hoffman, and Mohammad H. Hamidian. Density wave probes cuprate quantum phase transition. *Phys. Rev. X*, 9:021021, May 2019.

[143] Sandro Wenzel, Tommaso Coletta, Sergey E. Korshunov, and Frédéric Mila. Evidence for columnar order in the fully frustrated transverse field ising model on the square lattice. *Phys. Rev. Lett.*, 109:187202, Nov 2012.

[144] S. R. White. The average spectrum method for the analytic continuation of imaginary-time data. In David P. Landau, K. K. Mon, and Heinz-Bernd Schüttler, editors, *Computer Simulation Studies in Condensed Matter Physics III*, pages 145–153, Berlin, Heidelberg, 1991. Springer Berlin Heidelberg.

[145] S. R. White, D. J. Scalapino, R. L. Sugar, and N. E. Bickers. Monte carlo calculation of dynamical properties of the two-dimensional hubbard model. *Phys. Rev. Lett.*, 63:1523–1526, Oct 1989.

[146] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar. Numerical study of the two-dimensional hubbard model. *Phys. Rev. B*, 40:506–516, Jul 1989.

[147] Steven R. White. Spectral weight function for the two-dimensional hubbard model. *Phys. Rev. B*, 44:4670–4673, Sep 1991.

[148] Lei Wu, M. J. Young, Y. Shao, C. W. Garland, R. J. Birgeneau, and G. Heppke. Critical behavior of the second harmonic in a density wave system. *Phys. Rev. Lett.*, 72:376–379, Jan 1994.

[149] H. D. Xie, R. Z. Huang, X. J. Han, X. Yan, H. H. Zhao, Z. Y. Xie, H. J. Liao, and T. Xiang. Reorthonormalization of chebyshev matrix product states for dynamical correlation functions. *Phys. Rev. B*, 97:075111, Feb 2018.

[150] Hao Xu, Chia-Min Chung, Mingpu Qin, Ulrich Schollwöck, Steven R. White, and Shiwei Zhang. Coexistence of superconductivity with partially filled stripes in the hubbard model. *Science*, 384(6696):eadh7691, 2024.

[151] Luhang Yang and Adrian E. Feiguin. From deconfined spinons to coherent magnons in an antiferromagnetic Heisenberg chain with long range interactions. *SciPost Phys.*, 10:110, 2021.

[152] Sibin Yang, Gabe Schumm, and Anders W. Sandvik. Dynamic structure factor of a spin-1/2 heisenberg chain with long-range interactions. *Phys. Rev. B*, 111:224404, Jun 2025.

[153] Sibin Yang, Gabe Schumm, and Anders W. Sandvik. Single-hole dynamics in the 1d t-j model, 2025. Manuscript in preparation.

[154] Andrew Yates and Wojciech H. Zurek. Vortex formation in two dimensions: When symmetry breaks, how big are the pieces? *Phys. Rev. Lett.*, 80:5477–5480, Jun 1998.

[155] T. Yoshida, X. J. Zhou, T. Sasagawa, W. L. Yang, P. V. Bogdanov, A. Lanzara, Z. Hussain, T. Mizokawa, A. Fujimori, H. Eisaki, Z.-X. Shen, T. Kakeshita, and S. Uchida. Metallic behavior of lightly doped $\mathrm{la}_{2-x}\mathrm{sr}_x\mathrm{cuo}_4$ with a fermi surface forming an arc. *Phys. Rev. Lett.*, 91:027001, Jul 2003.

[156] Eddy Yusuf, Anuvrat Joshi, and Kun Yang. Spin waves in antiferromagnetic spin chains with long-range interactions. *Phys. Rev. B*, 69:144412, Apr 2004.

[157] Shiwei Zhang. 15 auxiliary-field quantum monte carlo for correlated electron systems. *Emergent Phenomena in Correlated Matter*, 2013.

[158] Xingjiang Zhou, Wei-Sheng Lee, Masatoshi Imada, Nandini Trivedi, Philip Phillips, Hae-Young Kee, Päivi Törmä, and Mikhail Eremets. High-temperature superconductivity. *Nature Reviews Physics*, 3(7):462–465, Jul 2021.

[159] W. H. Zurek. Cosmological experiments in superfluid helium? *Nature*, 317(6037):505–508, Oct 1985.

[160] W. H. Zurek. Cosmic strings in laboratory superfluids and the topological remnants of other phase transitions. *Acta Phys. Polon. B*, 24:1301–1311, 1993.

[161] W.H. Zurek. Cosmological experiments in condensed matter systems. *Physics Reports*, 276(4):177–221, 1996.

[162] Wojciech H. Zurek, Uwe Dorner, and Peter Zoller. Dynamics of a quantum phase transition. *Phys. Rev. Lett.*, 95:105701, Sep 2005.

CURRICULUM VITAE

[redacted]

[redacted]