Numerical approaches to isolated many-body quantum systems

Michael H. Kolodrubetz

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Physics
Adviser: Professor David Huse

September 2012

Abstract

Ultracold atoms have revolutionized atomic and condensed matter physics. In addition to having clean, controllable Hamiltonians, ultracold atoms are near-perfect realizations of isolated quantum systems, in which weak environmental coupling can be neglected on experimental time scales. This opens new opportunities to explore these systems not just in thermal equilibrium, but out of equilibrium as well.

In this dissertation, we investigate some properties of closed quantum systems, utilizing a combination of numerical and analytical techniques. We begin by applying full configuration-interaction quantum Monte Carlo (FCIQMC) to the Fermi polaron [1], which we use as a test bed to improve the algorithm. In addition to adapting standard QMC techniques, we introduce novel controlled approximations that allow mitigation of the sign problem and simulation directly in the thermodynamic limit. We also contrast the sign problem of FCIQMC with that of more standard techniques, focusing on FCIQMC's capacity to work in a second quantized determinant space.

Next, we discuss nonequilibrium dynamics near a quantum critical point, focusing on the one-dimensional transverse-field Ising (TFI) chain. We show that the TFI dynamics exhibit critical scaling, within which the spin correlations exhibit qualitatively athermal behavior. We provide strong numerical evidence for the universality of dynamic scaling by utilizing time-dependent matrix product states to simulate a non-integrable model in the same equilibrium universality class. As this non-integrable model has been realized experimentally [2], we investigate the robustness of our predictions against the presence of open boundary conditions and disorder. We find that the qualitatively athermal correlations remain visible, although other phenomena such as even/odd effects become relevant within the finite size scaling theory.

Finally, we investigate the properties of the integrable TFI model upon varying the strength of a non-integrable perturbation, attempting to understand the finite-size scaling behavior of the eigenstate thermalization crossover. We numerically estimate

the crossover scale of this perturbation strength using exact diagonalization, finding that the crossover scale decreases strongly with system size, and analytically provide a lower bound on the finite-size scaling of this crossover. However, we are unable to solve for the thermodynamic limit of this crossover; doing so will require larger systems and more comprehensive theory.

Acknowledgements

I could not have gotten to this point without the help of many people. Having switched topics fairly late in the game, I have had to rely on the help of others to bring me up to speed, and they have not let me down. To be able to walk down the hall and receive a master's course in the fractional quantum Hall effect or topological insulators, or just to be able to sit around for two hours discussing the ramifications of Majoranas in a nanowire, I could not imagine a better environment for developing as a physicist. While I will never be able to thank everyone who has helped me, it is the support of the entire Princeton physics community that has enabled any of my success.

The two most important individuals in helping rejuvenate my graduate career are David and Bryan. Two years ago, I happened into what has been a nearly ideal situation for me. Between Bryan's categorical knowledge of all things numeric and David's categorical knowledge of all things physic, I could not imagine a better team to study under. Working with them has reminded me of why I got into physics in the first place, and I'm looking forward to continuing my collaboration with Prof. Huse and, I both hope and expect, Prof. Clark.

Bryan and David have also been very valuable in helping me get a postdoc, providing useful advice despite my silly, neurotic questions. The other person who has been helpful in this respect is Shivaji, who not only wrote letters but provided a valuable sounding board as I weighed postdoc-related decisions. I'd also like to thank Shivaji for allowing me to tag along to many of his group functions; these not only helped me broaden my physics horizons, but generated the ideas that became Chp. 5.

I have greatly enjoyed learning physics from the other condensed matter theorists here in Jadwin, three of whom I would like to specifically thank: Sid, Amir, and Anushya. Having met Sid as a first year when he was kicking my ass in the condensed matter course, I was lucky enough to have him as a guide into the realm of condensed

matter theory, providing a template for me to aspire to. Meanwhile, listening to Amir and Anushya discuss Kibble-Zurek (before I even know what that meant) taught me much of what I know about dynamics. I'd like to particularly thank Anushya who, as my office mate, has had to listen to many of my stupid ideas before they made any sense, allowing me to run really dumb ideas past her until she helped make them sensible. It is nice to be able to walk less than the distance of a Jadwin hallway and be able run your ideas past a brilliant physicist; in Sid, Amir, and Anushya, I have been blessed with this opportunity.

I would like to thank Jason and all of his lab members for getting me started along my grad school path. I wish I had been able to hack it as an experimentalist, but, despite my shortcomings, working with you all left me much better prepared for my work with David and Bryan. Between journal clubs, group meeting, and just talking physics with y'all, I received the experience necessary to hit the ground running with David and Bryan much quicker than I could have done when I first came to Princeton. I'd like to thank all my out-of-town collaborators, including Prof. Krishnendu Sengupta, Dr. David Pekker, Dr. James Spencer, and Prof. Matthew Foulkes. Finally, I'd like to thank my parents for their support and for not complaining (too much) about their current lack of grandchildren.

This acknowledgments section was supported in part by ARO Award W911NF-07-1-0464 with funds from the DARPA OLE Program. Some of the computation was performed using the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. Additional computation was done on the Feynman and Della clusters at Princeton.

Relation to Previously Published Work

Chp. 4 is divided into two sections, the first of which (Sec. 4.1.3) comes from Ref. 3, which we are currently in the late stages of editing. Ref. 3 will also include the proof found in Appendix A. The second part, Sec. 4.2, is based on Ref. 4. Similarly, Chp. 5 involves the work from two papers, Refs. 5 and 6. In addition, work from the supplement to Ref. 6 on dynamics in open boundary conditions appears in Sec. 6.2. All other chapters are either introductions involving others' work (Chps. 2 and 3) or involve unpublished work (Chps. 6 and 7).

Contents

Chapter 1

Introduction

Ultracold atoms are an exciting field sitting at the boundary between condensed matter and atomic physics. With cold atoms one has the ability to realize and manipulate a set of designer Hamiltonians, with an amazing range of tools for control and measurement. Certain models that were invented as low-energy theories to model traditional condensed matter systems, such as the Bose-Hubbard model, can be realized more or less exactly using cold atoms. This provides an exciting avenue into exploring condensed matter physics in a new setting, with an unprecedented interplay between experiment and theory.

However, cold atoms also offer exciting new challenges distinct from those of traditional condensed systems. For one, cooling these systems is quite a different beast than cooling a sample in a cryostat. It may seem as if the micro- and nanokelvin temperatures routinely achievable in modern cold atom experiments would be sufficient to cool any system to its ground state, but in fact the temperature needed to see quantum effects (the so-called degeneracy temperature) is much lower in ultracold atoms. While the nanokelvin temperatures achievable in ultracold atoms are much colder than the millikelvin temperatures found in a dilution fridge, one must compare this temperature to the other energy scales in the problem. Since the characteristic

density is much lower in the the cold atom clouds – no more than one atom per $(1000\text{Å})^3$ for cold atoms versus one atom per $(1\text{Å})^3$ in a room-temperature crystal – this corresponds to a much lower degeneracy temperature for the cold atoms. So while nanokelvin temperatures are impressive, rendering the atoms ultracold on *our* energy scale, it is not immediately apparent whether they are cold on *their* energy scale.

More fundamentally, though, the concept of temperature and cooling is not as clearly defined in the context of cold atoms. This comes from the fact that cold atoms are well isolated from their environment. So whereas in a crystal of bismuth, the electrons quickly thermalize to the temperature of their environment due to interactions with lattice phonons, it is much less clear what mechanisms allow cold atoms to thermalize in the absence of such a heat bath. This is a major open question in the context of cold atom physics, and understanding the mechanism for presence and/or absence of thermalization is a step toward suggesting techniques to reach ever lower temperatures.

From a fundamental physics standpoint, the existence of such well-isolated systems allows us to turn the question on its head and ask what interesting properties the isolated quantum states have when out of thermal equilibrium. While aspects of these questions have been mulled for the past century by a number of theoretical physicists, their recent realization in ultracold atoms has brought non-equilibrium dynamics to the forefront of modern theoretical and experimental research. While dynamics are understood, and even exactly solvable, for some specific cases, the general rules for understanding non-equilibrium dynamics are still being written.

In order to study these questions for all but the simplest systems, one of two approaches is generally utilized. The more traditional approach is to make some approximation for which the system is solvable, e.g., a mean field theory treatment. This is a good way to get at the basic physics of many problems, and has been

very successful particularly in weakly-interacting cases such as the BCS theory of superconductivity.

However, to push past mean field or other analytical approximations for solving strongly interacting many-body problems, it is often necessary to utilize numerical methods. As a rule of thumb, these approaches should complement each other; mean field or variational solutions provide initial intuition about the physical system, then numerics help confirm or refine this understanding through methods that are exact, or at least involve weaker approximations than purely analytic methods.

1.1 Thesis outline

In this thesis, we discuss the application of numerical and analytical methods to a few problems in isolated quantum systems, with an eye on experimental applications in cold atoms. We begin by introducing cold atoms in Chp. 2, followed by some essential numerical methods in Chp. 3. As a first application of these numerical techniques, in Chp. 4 we develop and use full configuration-interaction quantum Monte Carlo (FCIQMC) for solving the ground state properties of the Fermi polaron. We then discuss dynamics near an isolated quantum critical point in Chp. 5, specifically solving the critical scaling of non-equilibrium dynamics for the quintessential quantum phase transition of the transverse-field Ising model. In the same chapter, we numerically compare the dynamics of the Ising model to those of an experimentally-realizable non-integrable model in the same universality class, providing evidence that the dynamics is universal. We further detail experimental considerations for this non-integrable model in Chp. 6. To better understand the onset of statistical mechanics in closed quantum systems, in Chp. 7 we investigate the finite size scaling of thermalization in a non-integrable extension of the Ising model. Finally, we draw some overall conclusions and discuss open questions in Chp. 8.

1.2 A brief note on notation

In keeping with the quantum statistical mechanical spirit of this work, we use the convention $\hbar = k_B = 1$ throughout, where $\hbar$ is Planck's constant divided by 2π and k_B is Boltzmann's constant. In addition, when discussing matrix product states we use the non-relativistic version of the Einstein summation, meaning that all repeated indices are summed over regardless of whether the indices are raised or lowered.

Chapter 2

Cold atoms

The cold atom revolution first caught fire in 1995 with the realization of a Bose-Einstein condensate (BEC) in a gas of cold atoms by three separate experimental groups [7–9]. Since that time, the field has taken off, demonstrating cooling of fermions [10], artificial lattices [11], artificial electric [11] and magnetic [12] fields, and much more. Ultracold atoms have matured to the point where certain models that have been used for years as theoretically elegant descriptions of basic phenomena, such as the transverse-field Ising model, have recently been implemented in a cold atom setting [2, 13, 14].

In this chapter, we will discuss the basic properties of cold atoms. We start by reviewing some procedures used for working with cold atoms experimentally in Sec. 2.1. We focus on more traditional techniques, such as those predominately implemented by the MIT groups [1, 10, 15], noting that similar methods underly more complicated setups such as the Greiner group's cold atom microscope [2]. Having introduced the basic properties of cold atoms, in Sec. 2.2 we address to what degree they can be treated as isolated quantum systems, a topic that we return to in later chapters.

2.1 Basic tools of cold atom experiments

It has been roughly one century since the initial realization of an atomic beam [16], yet only recently have the atoms in such a beam been slowed to the point where they could be trapped, cooled, and studied. The largest step toward this development was the invention of the laser, which has become the single most important tool in modern atomic physics. In addition to helping trap and cool the atoms, lasers enable a wide range of techniques for controlling them, including the creation of artificial lattices that act as a bridge between cold atom research and condensed matter physics. Magnetic and electric fields provide additional tools for controlling these systems; of particular note is ability to modify interactions through the use of a static magnetic field (see Sec. 2.1.2). In this section, we review the physics behind cooling, trapping, and manipulating atoms, closely following two reviews by Ketterle et al. [10, 15].

2.1.1 Trapping and cooling atoms

Starting from a hot atomic beam (at $\sim$500 K), a series of steps are used to cool the atoms down to quantum degeneracy and below. These are, in broad strokes,

1. Zeeman slowing, which reduces the energy of the gas to the point where it can be trapped.

2. Optical molasses cooling, to remove the majority of the remaining thermal motion, pre-conditioning the gas for evaporation.

3. Evaporative cooling, where much lower temperatures are achieved by removing the most energetic atoms from the trap.

Both Zeeman slowing and optical molasses cooling depend on Doppler cooling, which works as follows. Consider a particular atomic transition at frequency ω_0. If the atom is moving to the right at some velocity v, one irradiates it with left-moving

photons at frequency $\omega = \omega_0 - \Delta$, with small detuning Δ off of the atomic resonance. In the rest frame of the atom, the light is blue-shifted by the Doppler effect, leading to absorption if Δ is tuned to match the (v-dependent) Doppler shift. Assuming a fairly quick atomic decay rate, the atom spontaneously re-emits the photon in a random direction, bringing the atom back to its internal ground state. However, to conserve momentum, the atom must on the average receive a momentum kick to the left during this process, which in the original frame corresponds to slowing down the velocity v. This process can also be thought of as drag due to the radiation pressure of the incoming photons, but only acting on atoms with velocity near Doppler resonance.

To continually reduce the velocity, one must either adjust the frequency of the laser or of the atomic transition to keep the velocity in resonance. In Zeeman slowing, the atomic transition frequency is modified by passing the atoms through a magnetic field gradient. By using two hyperfine levels of the atom, the external magnetic field B allows control of the transition frequency ω_0 via the Zeeman effect $\Delta E_z \propto \mu B$, where μ is the magnetic dipole moment (see Fig. 2.1). The magnetic field gradient reduces the splitting as the atoms slow down, allowing the fixed frequency laser to remain in Doppler resonance with the atoms. The velocities achieved via Zeeman slowing, generally of order 40 m/s [17], are sufficiently small to allow trapping of the atoms.

After Zeeman slowing, there are two basic tools for trapping atoms: magnetic fields and lasers. Magnetic trapping has been around for longer, and was the manner in which the earliest experiments, such as the first BECs, were trapped and manipulated [7–9]. However, contemporary experiments often use hybrid magnetic-optical or even all-optical trapping, since the additional expense and/or complications are balanced by an improved ability to trap and control atoms (such as the ability to use strong field seeking states, see below). In general, any such trap is called a magneto-optical trap (MOT).

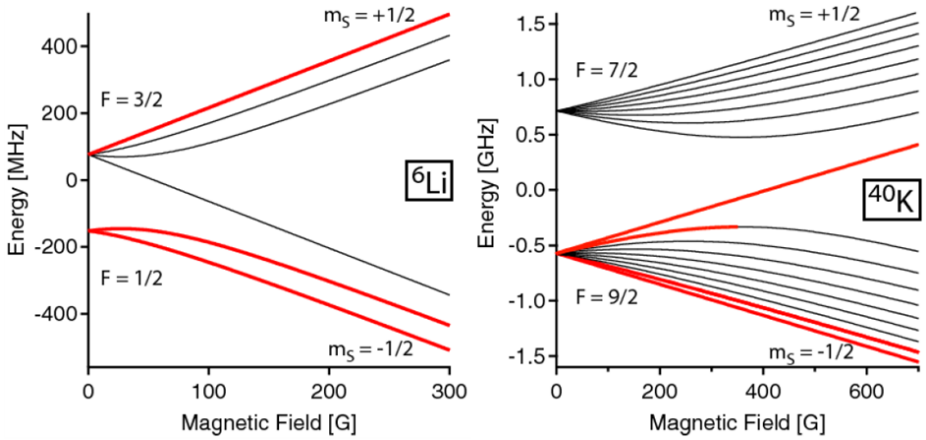

Figure 2.1: Hyperfine spectra of two fermionic species (figures reproduced from Ref. 10). Red lines indicate the most commonly used hyperfine states in ultracold fermion experiments.

Magnetic trapping involves an atomic gas polarized in a single hyperfine state; for reference, the hyperfine spectra of ^{6}Li and ^{49}K are shown in Fig. 2.1. These spectra split into weak field seeking states ($\mu > 0$), whose energy is lowest in a small magnetic field, and strong field seeking states ($\mu < 0$), whose energy is lowest in a large magnetic field. By surrounding the sample with coils and magnets, magnetic field configurations can be created so that $|B|$ is minimum (but not a maximum [15]) at the center of the trap. Therefore, only the weak field seeking states can be trapped by such a static magnetic field configuration, limiting the states that can be used in cold atom experiments using magnetic traps. Optical traps, however, can be made to work for any hyperfine state in the manifold, with the additional benefit of allowing an arbitrary external magnetic field to be added, which can be very useful (see Sec. 2.1.2).

The essence of optical trapping and manipulation is the *AC Stark effect* [18], which describes the interaction of light at frequency ω with an atomic transition $|g\rangle \leftrightarrow |e\rangle$ at frequency ω_0 (see Fig. 2.2). For a weak light field of magnitude $\mathcal{E}$ and intensity $I = 2\epsilon_0 c|\mathcal{E}|^2$, this interaction induces level repulsion between the (dressed) atomic states at second order in perturbation theory. If $\delta = \omega - \omega_0$ is the detuning between

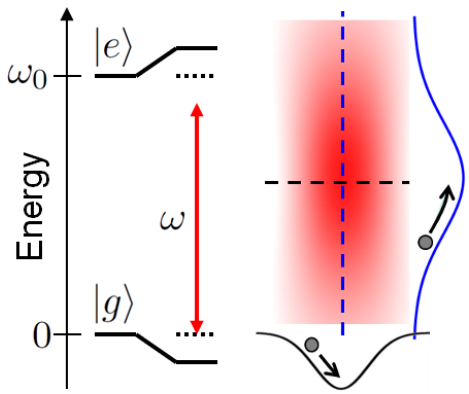

Figure 2.2: The AC Stark effect for a red-detuned laser, which can be used for optical trapping in the vicinity of its focal plane.

the laser and the atomic frequencies, then the (dressed) level $|g\rangle$ undergoes an energy shift [18]

$$\Delta E = \frac{|\langle e|\hat{p}|g\rangle|^2}{\delta}|\mathcal{E}|^2 \propto \frac{I}{\delta}\,, \tag{2.1}$$

known as the AC Stark shift or light shift, where $\hat{p} = -e\hat{x}$ is the electric dipole operator. This formula contains two important pieces of information about the AC Stark shift: first, the shift is proportional to the intensity of light, and second, the sign of the shift depends on the sign of the detuning δ.

The AC Stark shift is used for trapping the atoms by varying the intensity of the laser spatially, creating local minima or maxima in the effective potential seen by the atom. The dependence of the AC Stark shift on the sign of δ allows for two distinct modes of operation: red-detuned lasers ($\delta < 0$), for which the atoms are drawn to maxima of intensity, and blue-detuned lasers ($\delta > 0$), for which the atoms seek intensity minima. In practice, the easiest way to vary the intensity is to focus the laser using lenses, in which case there exists an intensity maximum in the focal plane. Therefore, the dominant mode of operation for optical dipole traps is to use one or more focused lasers in the red-detuned regime, such that the atoms are trapped in the vicinity of the focal plane (see Ref. 10, Fig. 2.2).

After the atoms have been trapped, optical molasses pushes their temperature down further through a similar Doppler cooling process as in Zeeman slowing. Now the atom cloud is surrounded by six counter-propagating beams (one shining "forward" and one "backward" for each of the x, y, and z axes). To achieve Doppler cooling, the lasers are slightly red-detuned from a highly excited transition, i.e., a transition that will quickly spontaneously decay. As before, the atoms experience a velocity- dependent force that results in effective drag, hence the term "optical molasses."

The final, and most effective, stage of the cooling process is evaporative cooling. Evaporative cooling is based on the simple idea – as seen in cooling coffee or in the mixing chamber of a dilution fridge – that if the most energetic particles are allowed to escape, the temperature of the remaining cloud becomes lower. The hottest atoms are allowed to escape by slowly lowering the depth of the trap, such that the hottest atoms are able to tunnel across the lip of the trap. Of course, this process decreases the number of atoms in the trap. However, if enough atoms were initially loaded and pre-cooled, the final number in the trap after evaporation can still be quite high (of order 10^7 [15]).

A crucial aspect of achieving a low-temperature thermal cloud via evaporative cooling is that the rate at which the gas thermalizes must be much higher than the rate of cooling. This is particularly problematic for producing ultracold fermions, where there are either no interactions for a hyperfine polarized cloud or only weak interspecies interactions for most values of the external magnetic field (see Sec. 2.1.2 for details). Therefore, it is very common when cooling a cloud of fermions to perform sympathetic cooling by placing it in contact with a cloud of bosons. The bosons effectively serve as an interaction medium that allows a much higher rate of thermalization; this is similar to how weakly interacting electrons in a solid state system are cooled by contact with the phonon reservoir.

2.1.2 Using cold atoms for condensed matter physics

Once atoms have been cooled and trapped, there are proposals to use them for a number of purposes, including metrology [19], chemistry [20], and quantum computation [21]. Here we will focus on another important application: using cold atoms as a clean, well-controlled system to simulate interesting models from condensed matter physics. To that end, there are three important properties of cold atoms which we will discuss: the ability to make artificial lattices using lasers, the ability to control the hyperfine polarization of the atomic cloud, and the ability to engineer a universal and tunable class of inter-atomic interactions.

Artificial optical lattices are based on the same idea as the optical dipole trap. For making a one-dimensional optical lattice, one creates a standing light wave using a retroreflected laser (a laser that is reflected off a mirror normal to its direction of propagation). The atoms then experience an effective sinusoidal potential (the optical lattice) added to the background parabolic confinement, with the strength of the potential set by the AC Stark shift. Whether the atoms want to reside at the nodes or the anti-nodes of the standing wave depends on whether the retroreflected beam is red- or blue-detuned from the relevant atomic transition. Two- and three-dimensional lattices can be made by adding retroreflected lasers along the other directions. While the lattice depth can be easily controlled by varying the intensity of the lattice laser, the lattice spacing a is fixed at $\lambda/2$, where λ is the wavelength of the laser, for these retroreflected lattices. Lattices with spacing $a > \lambda/2$ can be realized either by making a standing wave with a non-zero angle between the lasers or by holographic projection [22].

To control the population of individual hyperfine states, it is often useful to utilize a coupling to the magnetic dipole moment. As the hyperfine splittings are in the radio frequency (rf) regime, one can control the hyperfine state using an rf pulse, as in NMR. More explicitly, consider two hyperfine levels "$|\uparrow\rangle$" and "$|\downarrow\rangle$," connected by

a magnetic dipole coupling, with an unperturbed energy splitting ω_0. If we act on this system with a near-resonant rf wave at frequency ω with detuning $\delta = \omega - \omega_0$, then the effective Hamiltonian in the rotating wave approximation is [10]

$$H_{\text{eff}} = \frac{\delta}{2}\sigma^z + \frac{\Omega}{2}\sigma^x \ , \tag{2.2}$$

where Ω is the Rabi frequency characterizing the coupling strength between the photons and the two-level system. Starting from a spin-up polarized state, one could reverse the spin direction through an NMR-like pi-pulse, i.e., by turning on a resonant interaction ($\delta = 0$) for time $T_\pi = \pi/\Omega$. However, a more robust choice is "adiabatic rapid passage," wherein the detuning is slowly tuned from $\delta \ll 1$ to $\delta \gg 1$. As long as this is done sufficiently slowly ($\dot{\delta} \ll \Omega^2$), the system remains in its ground state, which goes from $|\uparrow\rangle$ in the limit $\delta \ll 1$ to $|\downarrow\rangle$ in the limit $\delta \gg 1$. One can also prepare a partially polarized state by sweeping the detuning non-adiabatically, in which case the polarization is set by the ramp rate through a standard Landau-Zener process [23, 24]. While this in principle creates a coherent spin mixture, the state generally decoheres – giving a classical mixture of spins up and down with stable polarization (long T_1) – because gradients of the magnetic field within the MOT lead to position-dependent shifts in the detuning [10].

Finally, interactions between the atoms can be controlled through use of a Feshbach resonance [26], which occurs near a point where two hyperfine states of the atoms energetically wish to form a bound state (a molecule). The classic way of the seeing this is to solve the scattering problem between two atoms, using an atomic interaction that is generally well-described by a Lennard-Jones potential, with characteristic range on the order of ten angstroms [10]. However, we can think instead about the simpler case of a square well potential of some depth V and radius R, the solution to which can be found in any undergraduate quantum mechanics textbook.

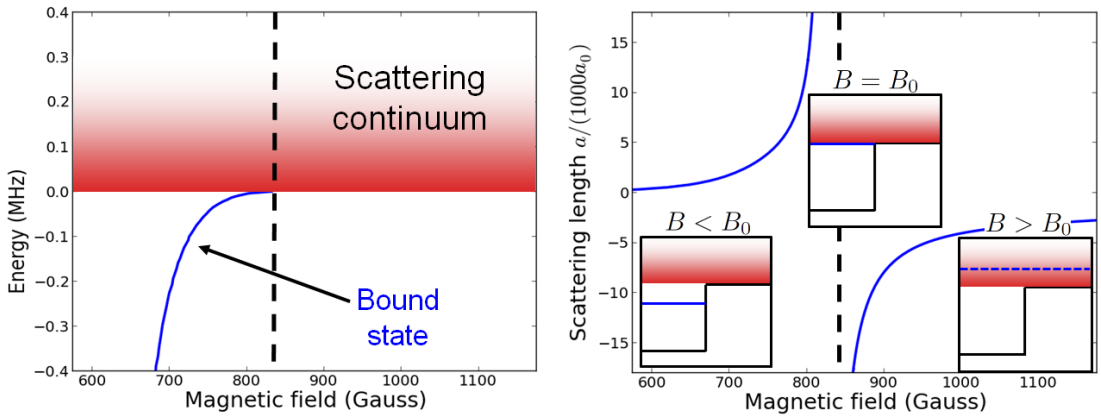

Figure 2.3: A Feshbach resonance occurs when two particle scattering is modified near the formation of a bound state (Left), leading to a tunable scattering length (Right) that diverges at the resonance position (data from Ref. 25, $a_0 = 0.53$Å is the Bohr radius).

As the depth V is tuned to be increasingly negative, at some $V = V_0$ it will admit a bound state – we will call this point the "Feshbach resonance." In the vicinity of V_0, one can solve the scattering problem to extract the scattering phase shift δ_s. An expansion of δ_s at low momentum yields two parameters: the s-wave scattering length a and the effective range r_{eff}. For characteristic values of the inter-atomic potential one finds that $a \gg r_{\text{eff}}$, and thus the phase shift can be approximated by its $r_{\text{eff}} \to 0$ limit, in which it depends solely on the scattering length a and not on the details of the scattering potential. Therefore, in this limit, our assumption of solving the square well in place of the Lennard-Jones potential is justified.

An important property of this scattering length a is that it becomes infinite at the Feshbach resonance, where the zero energy scattering state is resonant with the bound state. In fact, the scattering length changes sign at this point as well, from positive (repulsive) on the side of the Feshbach resonance where the resonant level is well-defined, to negative (attractive) on the other side of the resonance. A key point is that, as the scattering occurs between hyperfine states, the effective scattering potential V is tuned by an external magnetic field. Therefore, a static magnetic field

can be used as a control knob that allows one to tune the scattering length across resonance. As an example, the scattering length in ^{6}Li is shown in Fig. 2.3.

The tools described above have been combined to do a vast amount of exciting experimental research, even just within the last few years. The ability to tune the interactions has been exploited to greatly extend our understanding of superfluidity, allowing investigation of the Bose-Einstein condensate (BEC) to Bardeen-Cooper-Schrieffer (BCS) superfluid crossover in both spin balanced [27–29] and imbalanced [30–32] situations. A particularly interesting aspect of this crossover is the behavior of the quantum gas at the point where the scattering length a diverges, known as the unitary scattering limit or unitarity. There, since both the values of a and the potential range R cannot enter into the physics of the problem, the sole length scale is set by the density, or equivalently by the Fermi wavevector k_F. Thus, in this limit thermodynamics becomes universal, and a very successful recent investigation of thermodynamics at unitarity has shown remarkable agreement between experiment and theory [33].

Another major avenue of research has been realizing model Hamiltonians inspired by condensed matter physics using cold atoms on optical lattices, including the Bose-Hubbard model [11], variants of the transverse-field Ising model [2, 13, 14], and initial work on fermionic spin and Hubbard-like models [34]. The major advantage of simulating these systems with cold atoms is that the tunability of cold atoms allows for realization of nearly exact versions of condensed matter models (i.e., the Bose-Hubbard model) originally derived as low energy effective theories of condensed matter systems. In addition to improved control, cold atoms have the potential to visualize the states in a novel manner, replacing solid state techniques such as transport and neutron scattering with position-, momentum-, and/or spin-resolved imaging of the ultracold atoms [10].

2.2 Cold atoms as isolated quantum systems

While cold atoms have been used quite successfully for simulating condensed systems, there is one aspect in which they differ from most traditional condensed matter experiments; cold atoms are, to a high degree, isolated quantum systems. A general solid state physics experiment is rarely performed in a true vacuum. For example, a crystal may be attached to a sample holder, which in turn sits in thermal contact with a cryostat. The electrons thermalize not through interactions between each other, but by interactions with the phononic environment, which in turn exchanges heat with the cryostat. For a traditional transport experiment, the electrons in the crystal are connected through leads to some external perturbation – say a voltage source – that causes electrons to flow between the leads and the sample. Clearly the conduction electrons have many points of contact with their environment, and while out-of-equilibrium behavior is possible (transport being, for example, a linear non-equilibrium response) it is generally not reasonable to treat the sample as a closed quantum system, uncoupled from its environment.

In cold atoms, on the other hand, environmental interactions are fewer and of a different nature. The experiment takes place in ultra-high vacuum with the atomic cloud physically separated from the walls or other potential heat sources. While there do exist some small sources of environmental decoherence, which we detail in Sec. 2.2.2, most experiments on short – but still physically relevant – time scales can be well approximated as an isolated quantum system.

2.2.1 Differences from open quantum systems

The distinction between "open" quantum systems (those in contact with an environment) and isolated or "closed" quantum systems (those not interacting with an environment) is manifest primarily in the dynamics. We know how to find the exact

dynamics of a closed quantum system described by wavefunction $|\psi\rangle$ in a generic time-dependent Hamiltonian $H(t)$: solve the Schrödinger equation

$$i\frac{d|\psi(t)\rangle}{dt} = H(t)|\psi(t)\rangle \ . \tag{2.3}$$

This leads to unitary evolution; the wavefunction is given by $|\psi(t)\rangle = U(t)|\psi(0)\rangle$, where $U(t)$ is the unitary operator

$$U(t) = T\exp\left[-i\int_0^t H(t')dt'\right] \ . \tag{2.4}$$

We can also extend these ideas to initial mixed states, which can be created, for example, if the system and environment were allowed to interact in the past. During the period in which the system and environment do not interact, the system's density matrix again undergoes unitary evolution: $\rho(t) = U(t)\rho(0)U^\dagger(t)$.

An open quantum system can in principle be treated as a closed quantum problem if the environment is treated as part of the Hilbert space. However, if we only have experimental access to the state of subsystem A, its properties are governed by a reduced density matrix given by tracing out the environment E. Then the evolution of A may be described by solving for

$$\rho_A(t) = \mathrm{Tr}_E\big[|\psi(t)\rangle\langle\psi(t)|\big] \ . \tag{2.5}$$

Unlike the complete quantum system, the evolution of the subsystem A is not generally given by unitary evolution.

One interesting aspect of isolation is that it is not in general possible to describe the system in terms of the thermodynamic Gibbs ensemble, $\rho_G = \exp(-\beta H)$, where H is the Hamiltonian and $\beta = 1/T$ is the inverse temperature. This comes from the fact that the population in each energy eigenstate of the isolated system is conserved

during unitary evolution, so the energy distribution can't redistribute to that of the Gibbs ensemble. This is a topic about which we have more to say (see Chp. 7) but, for the time being, we would like to briefly address the experimental implications of thermalization in closed quantum system.

In open quantum systems, thermalization is generally thought of in the context of surrounding the system by a heat bath at some temperature T. Then, over some time scale τ_{therm}, the system exchanges heat with the environment to reach a new equilibrium in which the system temperature approaches T and the system is well described by the Gibbs ensemble.

In closed quantum systems, thermalization occurs *within* the system, rather than through interactions with the external environment. Therefore, if we prepare the system in a non-equilibrium state at some total energy E, but with said energy distributed among energy eigenmodes in a manner incommensurate with the Gibbs ensemble, a loose picture of closed system thermalization is that interactions within the quantum system act to redistribute the energy density. If these interactions do their job, we again expect that after some time τ_{therm}, observables of a finite subsystem will be well described by those of the Gibbs ensemble with β chosen to give the correct total energy. In some sense, the remainder of the quantum system acts as a heat bath for the subsystem of interest, restoring the conventional ideas of statistical mechanics. As a corollary, in order to achieve thermalization on reasonable timescales, it is important to maximize interactions within the cloud, so that thermalization happens faster than, say, the rate of atom loss due to evaporative cooling.

2.2.2 Mechanisms for loss and heating

While cold atoms can be largely be treated as isolated quantum systems, there remain a few environmental source of decoherence when looking on a long time scale (generally milliseconds to seconds). Which heating mechanism dominates depends on

the particular experimental setup of interest. Here we mention three major sources of decoherence – spontaneous emission via interactions with the photons, three body interactions, and background gas collisions – and discuss when they may be applicable. All of these mechanisms lead to atom loss, and usually result in heating (increase in entropy) of the remaining atomic cloud. It is also worth noting that, while environmental interactions are generally thought of as contributing to heating, this need not always be the case. For example, evaporation is associated with atom loss, but clearly results in cooling.

Whenever an atom is in an excited state that is coupled by an electric dipole moment to other, lower-lying states, it is always possible for that excited state to emit a photon and relax to the lower atom level; this process is known as spontaneous emission. Cold atom experiments work predominately in the lowest orbital manifold, where the atomic levels are stable or quasi-stable (lifetime much longer than the experimental timescale) to spontaneous emission. However, the AC Stark effect requires some small dressed state population in the higher orbital manifold, in which there is a non-zero rate of spontaneous emission. From such an excited state, there are two distinct decay paths: back down to the original hyperfine state or to some other hyperfine state. In the first case, the atom will generally remain in the trap, but by momentum conservation should pick up a kick to match the momentum of the emitted photon, thus contributing to heating. In the second case, it is often true that the hyperfine state after spontaneous emission is not trapped by the MOT potentials; therefore, the atom is usually lost from the trap, often scattering from atoms in the trap as it leaves and increasing their heat [35].

Spontaneous emission, while omnipresent, can be controlled via the detuning δ between the control laser and the atomic transition. At first order in perturbation theory, the atomic orbitals hybridize and the probability of finding the atom in its excited orbital state scales as I/δ^2, where I is the intensity of the laser. This also

gives the scaling of the spontaneous emission rate. Given that the magnitude of the AC Stark shift scales as I/δ, experimentally one generally tries to maximize the intensity while simultaneously using as high a detuning as possible, to minimize the spontaneous emission rate for a given magnitude of the AC Stark shift. Therefore, with sufficiently far-detuned lasers, spontaneous emission can often be neglected [18].

Another source of decoherence is atom loss due to three body recombination in the vicinity of a Feshbach resonance [10]. In deriving the scattering length near a Feshbach resonance, we neglected the actual formation of long-lived bound states on the grounds of energy conservation, instead treating the bound states as short-lived excitations that modify the scattering process. However, if three particles simultaneously interact, it *is* possible to form a long-lived molecule by putting the excess energy in the motional state of the third atom. The amount of energy transferred in this case is usually sufficient to kick the third atom out of the trap, so all three atoms are generally lost since the MOT is not designed to trap molecules. Again, as the atoms leave the trap, they will often transfer some of their momentum to the trapped atoms, thus contributing to heating. To avoid three body loss, it is important to work with dilute gases, since the rate of three body losses – set by the probability of seeing three atoms in an interaction volume – scales as the cube of the density.

Finally, in a realistic experiment, the atoms are not trapped in a perfect vacuum. In addition to the trapped atoms, there remain a number of thermal atoms bouncing around in the chamber, which will immediately remove an atom from the trap upon scattering. In principle, this can be controlled experimentally by improving the vacuum; however, once other heating mechanisms are under control, background gas collisions still remain a vital experimental concern.

There are some other effects which can contribute to heating in these cold atom systems, such as electrical or mechanical noise in various parts of the experimental apparatus [36], e.g., fluctuations in laser power. However, despite the sources of

decoherence, cold atomic gases can still be treated as nearly perfectly isolated on time scales long enough to see some interesting effects [35, 37]. Understanding some of the properties of such isolated quantum systems is a major theme of this thesis.

Chapter 3

Numerical approaches to ground state and dynamical properties

To solve many strongly-interacting problems exactly, one must often resort to numerical methods when no analytical techniques are available. The simplest such method is exact diagonalization, where one writes out the Hamiltonian matrix in some finite Hilbert space and proceeds to solve the full spectrum using a Hermitian eigensolver [38] or the ground state using a sparse matrix eigensolver such as the Lanczos algorithm [39]. Exact diagonalization returns all the information about the energy eigenstates, in the form of a vector in the Hilbert space, from which it is straightforward to find the value of arbitrary observables. In this sense, exact diagonalization is the simplest and most flexible numerical technique available for most problems, and is usually good for initial numerical investigation of some problem.

However, the efficiency of exact diagonalization scales very poorly with system size. Consider the case of an L-site spin-1/2 chain. The size of this problem's Hilbert space is $N = 2^L$, meaning the Hamiltonian operator can be represented as a $2^L \times 2^L$ matrix. Since exact diagonalization of dense matrices requires computational time of order N^3 (which reduces to order N^2 for Lanczos on sparse matrices), the

computational time scales exponentially with L, which quickly becomes a problem for systems of size $L \gtrsim 20$. Similar scaling holds for evaluating observables, assuming one can efficiently generate the matrix corresponding to the observable of interest, while the situation becomes even worse for a time-dependent Hamiltonian. Even some clever tricks, such as exploiting symmetries to block diagonalize the Hamiltonian into smaller matrices corresponding to the various symmetry sectors, only generally reduce N by a polynomial function of L, leaving the scaling exponential. Given this limitation of exact diagonalization, a number of numerical techniques have been developed that improve the scaling properties to solve larger systems. We will discuss two such techniques, quantum Monte Carlo and the density matrix renormalization group (DMRG), in the remainder of this chapter.

Before getting into the details of how these numerical methods work, we would like to briefly discuss some generic principles regarding numerics. First, how can one in general skirt the exponential scaling problem, and what is lost in the process? While many different numerical techniques approach this question from many different angles, the methods described in this chapter boil down to two basic concepts. First, Monte Carlo methods are based on the idea that one does not need to look at every basis element in the Hilbert space to calculate meaningful physical quantities. Instead, one can often sample a very small fraction of the Hilbert space at random, obtaining a representative "snapshot" of the wavefunction instead of the entire thing. While the Hilbert space grows exponentially, the number of basis elements that one samples need not. Second, rather than working within the entire Hilbert space, solving the problem within a subset of Hilbert space with the "important physics" can provide a great reduction in the computational power needed. Usually, such a technique has a variational knob that allows one to gradually increase the fraction of Hilbert space utilized, to help controllably remove the approximation. Such a variational approach is used in DMRG and certain Monte Carlo methods.

The second, more philosophical question, is how numerics fit into the traditional landscape of analytical and experimental physics? Views on this topic range from "they don't" to "who cares?", but nevertheless I would like to give a brief perspective on this topic. As we've mentioned, one major plus for computational approaches is their potential to yield exact answers for problems that cannot be solved analytically. However, such answers are only as good as one's ability to interpret them, which generally relies on approaching the system through a series of analytical approximations. In many ways, the situation is similar to the relationship between experimental and theoretical physics, where some unexpected signal in an experimental setup may require a new theoretical description (cf. superconductivity). One major advantage of numerical tools is their tunability; while it might be hard to change the Hubbard U in some high-T_c cuprate, it is comparatively easy to do so in Monte Carlo code. Therefore, analytical, experimental, and numerical methods can in principle work together. In this context, numerical methods can be thought of as another very important tool in the theorist's toolkit. However, like any other tool, one must be careful in properly applying numerical techniques, to avoid accidentally bashing a hole through the drywall.

3.1 Quantum Monte Carlo

Quantum Monte Carlo is based on the idea that, when solving a problem involving a complicated sum or integral, it is often not necessary to examine every term in the sum. Instead, as we will see, one can sample a small fraction of the terms and, as long as the sum is not too singular, the answer can sometimes be obtained (to fixed precision) much faster than by performing a direct summation. To illustrate this point, imagine that you wanted to know how often you would win a game of Solitaire, playing by some specific strategy. In principle, you could sit down, write out all the

possible scenarios, and solve for the winning percentage. Alternatively, you could simply simulate 1000 games of Solitaire to figure out the winning percentage (within a statistical error bar); on a modern computer, this would take much less time to run than to, say, write the code. Clearly, sampling a small part of the space of all possible deck shuffles – only considering 1000 out of $52! = 8 \times 10^{67}$ permutations – is much more efficient than doing the full "summation."

A more informative example is to consider numerically evaluating some complicated integral in d dimensions, such as

$$I = \int_V e^{-r^2/2} |\sin(r^2)| d^d r \ . \tag{3.1}$$

Here, for simplicity, we consider integrating over the hypercube V given by $r_i \in (0, 1)$ for $i = 1, 2, \ldots d$.

A naive method for solving this integral numerically is to perform Riemann integration; break up the space into n^d hypercubes with side length $h = n^{-1}$ and then sum up the values at the centers of the cubes. A quick Taylor expansion tells us that the error in this case is roughly

$$\epsilon_{\text{Riem}} \approx \frac{\langle \nabla^2 f \rangle}{2} h^2 \ , \tag{3.2}$$

where $f = \exp(-r^2/2)|\sin(r^2)|$ is the function we wish to integrate and $\langle \nabla^2 f \rangle$ is the average value of the Laplacian within the integration volume V. To achieve error ϵ_0, we must use $h \sim \sqrt{\epsilon_0}$ independent of d. Then, as there are $n^d \sim \left(\frac{1}{\sqrt{\epsilon_0}}\right)^d$ terms in the sum, the number of steps to perform Riemann integration with a fixed accuracy ϵ_0 scales exponentially with number of dimensions.

In Monte Carlo (MC), we instead sample this integrand stochastically. To do so, assume that we have some method to sample from an arbitrary probability distribu-

tion, for example

$$p(r) = \frac{e^{-r^2/2}}{\mathcal{N}} \ , \tag{3.3}$$

where $\mathcal{N}$ is some (possibly unknown) normalization.[1] The standard trick in MC is to rewrite the integral as the expectation value of some observable over this distribution.

$$\begin{aligned} I &= \int_V e^{-r^2/2} \left|\sin(r^2)\right| d^d r & (3.4)\\ &= \frac{\overbrace{\left(\int_V e^{-r^2/2} d^d r\right)}^{\mathcal{C}} \left(\int_V e^{-r^2/2} \left|\sin(r^2)\right| d^d r\right)}{\int_V e^{-r^2/2} d^d r} & (3.5)\\ &= \mathcal{C}\, \frac{\mathcal{N}^{-1} \int_V e^{-r^2/2} \left|\sin(r^2)\right| d^d r}{\mathcal{N}^{-1} \int_V e^{-r^2/2} d^d r} & (3.6)\\ &= \mathcal{C}\, \frac{\int_V p(r) \left|\sin(r^2)\right| d^d r}{\int_V p(r) d^d r} & (3.7)\\ &= \mathcal{C}\, \langle \left|\sin(r^2)\right| \rangle_{p(r)} \ , & (3.8) \end{aligned}$$

where

$$\mathcal{C} = \int_V e^{-r^2/2} d^d r = \left(\sqrt{\frac{\pi}{2}} \mathrm{erf}\left[\frac{1}{\sqrt{2}}\right]\right)^d \ . \tag{3.9}$$

So to get I along with the standard error σ_I, we must sample a large number of configurations r from the probability distribution $p(r)$ and then compute the statistics of the observable $|\sin(r^2)|$.

An analysis of the errors shows why this helps. Whereas the (deterministic) error in Riemann integration scales exponentially with number of dimensions, the error in MC is purely statistical, meaning it scales as $N^{-1/2}$, where N is the number of configurations sampled. The MC error also depends on the statistical variance of the observable and the correlation time (the number of steps between independent samples); empirically, for a wide range of integrands, these quantities do not depend

[1] A more standard technique for doing Monte Carlo integration such as this is to instead sample from flat distribution $p(r) = \text{const}$. We use Gaussian sampling here as a precursor to VMC, motivating the introduction the Metropolis algorithm.

Number dim. d	$\epsilon_{\text{Riem}} = \lvert I - I_{\text{Riem}}\rvert / I$	$\epsilon_{\text{MC}} = \lvert I - I_{\text{MC}}\rvert / I$	$\epsilon_{\text{MC}}/\epsilon_{\text{Riem}}$
2	1.014×10^{-7}	5.342×10^{-4}	5.267×10^{3}
4	2.603×10^{-5}	1.249×10^{-3}	4.797×10^{1}
10	2.668×10^{-2}	1.181×10^{-3}	4.427×10^{-2}
20	1.547×10^{-1}	1.360×10^{-3}	8.791×10^{-3}

Table 3.1: Relative errors of Riemann summation and Monte Carlo integration for solving the integral in Eq. 3.1. When varying the number of dimensions d, the Riemann hypercube size h is adjusted so that both Riemann and Monte Carlo integration utilize $2^{20} \approx 1$ million computational steps.

strongly on dimensionality. Therefore, one expects that the statistical error in MC should be essentially independent of dimensionality, making it generally preferable to Riemann summation for higher-dimensional integrals. These expectations are confirmed for the integral from Eq. 3.1, as seen in Table 3.1.

There is one important aspect of this argument that we have not addressed, namely how to sample from some arbitrary distribution $p(r)$.[2] This can be done by the Metropolis algorithm [40], which works as follows.

1. Start in some random configuration C from configuration space Ω. In the example above, each configuration C is a d-dimensional coordinate r and the configuration space is $\Omega = \{r \,|\, r \in V\}$.

2. Sample some other configuration C' according to a known probability distribution $q(C'|C)$. For example, one could select C' at random from Ω, in which case $q(C'|C) = \frac{1}{N_\Omega - 1}$, where N_Ω is the number of configurations in Ω.

3. Accept the move from $C \to C'$ (i.e., change configuration from C to C') with probability $p_a(C'|C) = \min\left[1, \frac{p(C')q(C|C')}{p(C)q(C'|C)}\right]$. If the move is not accepted, keep the configuration C.

[2] In this example, the Gaussian distribution is special in that very efficient analytical tools exist to sample from it. However, the Metropolis algorithm can be applied to a much more general class of distributions.

4. Measure the value of the desired observable(s) in the current configuration, and accumulate statistics over multiple MC steps. In the example above, this means computing $|\sin(r^2)|$ for r at each step.

5. Repeat from step 2. Each iteration of this procedure is known as a MC step.

The Metropolis algorithm guarantees that, after a few steps required to forget the initial conditions, the configuration C will be sampled proportional to the probability distribution $p(C)$. Note that, as the accept probability only involves the ratio $p(C')/p(C)$, it is not necessary that $p(C)$ be normalized. It is also not necessary that $q(C'|C)$ be non-zero for all $C' \neq C$. It is, however, important that any configuration C' can be reached from any other configuration C in a finite number of steps, to avoid getting stuck in a corner of configuration space.

To motivate how the Metropolis algorithm works, we show that $p(C)$ is a fixed point of this procedure. Assume, then, that C is already distributed according to $p(C)$. Following the above procedure, the probability of ending up at configuration C after one step is

$$\begin{aligned}
p_1(C) &= \overbrace{p(C)\sum_{C'\neq C} q(C'|C)(1-p_a(C'|C))}^{\text{Probability to stay in } C} + \overbrace{\sum_{C'\neq C} p(C')q(C|C')p_a(C|C')}^{\text{Probability to move to } C} \\
&= p(C)\sum_{C'\neq C} q(C'|C) + \sum_{C'\neq C}\Big\{p(C')q(C|C')p_a(C|C') - p(C)q(C'|C)p_a(C'|C)\Big\} \\
&= p(C) + \sum_{C'\neq C}\Bigg\{p(C')q(C|C')\min\left[1, \frac{p(C)q(C'|C)}{p(C')q(C|C')}\right] - \\
&\qquad\qquad p(C)q(C'|C)\min\left[1, \frac{p(C')q(C|C')}{p(C)q(C'|C)}\right]\Bigg\} \\
&= p(C)\ . \quad \square
\end{aligned} \tag{3.10}$$

In addition to performing complicated multidimensional integrals, MC methods are also very useful in simulating classical thermodynamic systems; in this case, con-

figurations i (the positions and momenta of N particles, for example) are sampled from the partition function Z, with weight proportional to $\exp(-\beta E_i)$. From here out, however, we focus instead on applying MC to problems in quantum mechanics (quantum Monte Carlo, a.k.a. QMC) to get both variational and exact values for ground state observables.

3.1.1 Variational Monte Carlo

In variational Monte Carlo (VMC), one solves for the expectation value of physical observables with respect to some trial wavefunction $|\psi_T\rangle$. An observable $\mathcal{O}$ has the quantum mechanical expectation value $\langle\mathcal{O}\rangle = \frac{\langle\psi_T|\mathcal{O}|\psi_T\rangle}{\langle\psi_T|\psi_T\rangle}$. Inserting a complete set of states $\{|x\rangle\}$, this becomes

$$\begin{aligned}
\langle\mathcal{O}\rangle &= \frac{\sum_x \langle\psi_T|x\rangle\langle x|\mathcal{O}|\psi_T\rangle}{\sum_x \langle\psi_T|x\rangle\langle x|\psi_T\rangle} && (3.11)\\
&= \frac{\sum_x \langle\psi_T|x\rangle\langle x|\psi_T\rangle \frac{\langle x|\mathcal{O}|\psi_T\rangle}{\langle x|\psi_T\rangle}}{\sum_x \langle\psi_T|x\rangle\langle x|\psi_T\rangle} && (3.12)\\
&= \frac{\sum_x |\langle\psi_T|x\rangle|^2 \frac{\langle x|\mathcal{O}|\psi_T\rangle}{\langle x|\psi_T\rangle}}{\sum_x |\langle\psi_T|x\rangle|^2} \ . && (3.13)
\end{aligned}$$

After such a rewriting, the physical expectation value now has the form of a *statistical* expectation value over a configuration space $\Omega = \{|x\rangle\}$:

$$\langle\mathcal{O}\rangle = \langle\mathcal{O}_L(x)\rangle_{p(x)}, \tag{3.14}$$

where $p(x) \propto |\langle\psi_T|x\rangle|^2$ and $O_L(x) \equiv \frac{\langle x|\mathcal{O}|\psi_T\rangle}{\langle x|\psi_T\rangle}$. This allows us to solve for $\langle\mathcal{O}\rangle$ using MC.

The word "variational" in VMC comes from the fact that $|\psi_T\rangle$ is often chosen from some variational class of wavefunctions. The results of the VMC simulation are then used to tune the variational parameters in order to, for example, minimize the

energy. This can be done using standard techniques for function optimization [38] with modifications to deal with statistical noise, often taking advantage of VMC's ability to calculate derivatives of the observable with respect to variational parameters for use as input to the optimization routines [41, 42].

3.1.2 Ground state projection Monte Carlo

There are a number of techniques to move beyond a variational approach and stochastically sample the true ground state of a real-valued Hamiltonian H. One particular class of methods involve projection to the ground state, starting from the realization that

$$\lim_{\beta\to\infty} e^{-\beta(H-E_0)}|\psi_T\rangle = |\psi_0\rangle \ , \tag{3.15}$$

where E_0 ($|\psi_0\rangle$) is the ground state energy (wavefunction) and $|\psi_T\rangle$ is an arbitrary initial state with the condition that $\langle\psi_T|\psi_0\rangle \neq 0$.

Ground state QMC methods perform the imaginary time projection $e^{-\beta H}$ stochastically. To do this, one starts the algorithm with N_w walkers distributed in configuration space, such that the average number of walkers at configuration x is proportional to $\langle x|\psi_T\rangle$. In this way, the walkers are a stochastic "snapshot" of the initial wavefunction $|\psi_T\rangle$. Assuming the Hamiltonian has a bounded spectrum (e.g., for a lattice Hamiltonian), one can break up the exponential into small time steps:

$$e^{-\beta H} = \underbrace{e^{-\tau H}e^{-\tau H}\cdots}_{\text{repeated } N=\beta/\tau \text{ times}} \approx (1-\tau H)^N \ . \tag{3.16}$$

Then at each time step, one stochastically applies the operator $U = 1-\tau H$, such that after N steps of the algorithm, the walkers represent a snapshot of $U^N|\psi_T\rangle$, i.e., the average number of walkers in configuration x is proportional to $\langle x|U^N|\psi_T\rangle$. For large enough N, these walkers correspond to a stochastic snapshot of the ground state, $|\psi_0\rangle$.

To illustrate how such a problem is solved in practice, consider finding the ground state of a single particle hopping on a one-dimensional L-site lattice with open boundary conditions. The Hamiltonian is

$$H = -\sum_{j=1}^{L-1} (c_j^\dagger c_{j+1} + \text{h.c.}) \ , \tag{3.17}$$

in units where the hopping energy is 1. We start the algorithm by sampling N_w walkers proportional to $\langle j|\psi_T\rangle$. For the example shown in Fig. 3.1, the initial state is chosen to seed all walkers at the middle site, $\langle j|\psi_T\rangle = \delta_{j,L/2}$. We wish to apply $U = 1 - \tau H$ stochastically, which can be accomplished as follows. For a walker at site j, there are three moves allowed: stay in place with probability $1/\mathcal{N} \sim \langle j|U|j\rangle$, move left one site with probability $\tau/\mathcal{N} \sim \langle j-1|U|j\rangle$, and move right with probability $\tau/\mathcal{N} \sim \langle j+1|U|j\rangle$. The normalization $\mathcal{N} = 1 + 2\tau$ is chosen to make these probabilities add up to one. If these moves cause the walker to move outside of the box, then we throw that walker out of the simulation. Therefore, after moving all the walkers and tossing away those that leave the box, we must resample to keep the walker number fixed at N_w, such that the expected number of walkers at site j after resampling is proportional to the number of walkers at j before resampling. For example, if there are two walkers on site j and one on site j' and we wish to resample to $N_w = 9$ walkers, then we replace the current walkers by a randomly-generated set of 9 walkers at sites j and j', such that the expected number on site j is $(2/3) \times 9 = 6$. These rules define a random walk for the N_w walkers, such that after N steps, the average number of walkers at site j is proportional to $\langle j|U^N|\psi_T\rangle$. Thus, the walkers give a snapshot of the wavefunction $U^N|\psi_T\rangle$, as desired. An illustration of such a random walk can be seen in Fig. 3.1.

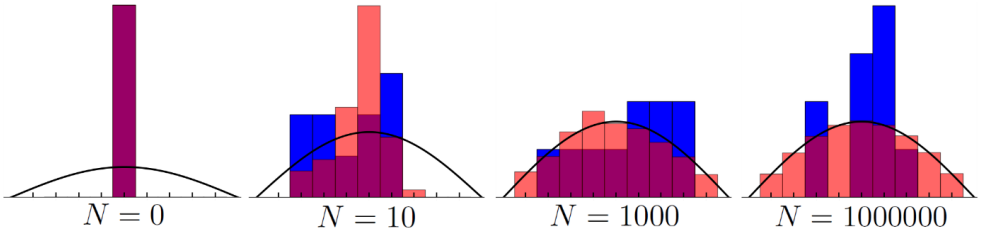

Figure 3.1: Illustration of the simple QMC example given in the text, with time step $\tau = 0.1$ and $N_w = 10$ walkers. Dark blue bars are a snapshot of the walkers at various steps of the algorithm, while light red bars show cumulative histograms of walker locations, which eventually approach the ground state (black line).

3.1.3 The sign problem

In the previous example, we have swept one very important issue under the rug: how can one sample the ground state wavefunction if it is not positive valued, as is allowed in quantum mechanics? What does it mean to sample from a negative "probability?" This issue, known as the "sign problem" or "Fermi sign problem," is one of the most insidious problems in the field of QMC.

The heart of the sign problem is that MC methods are designed to sample from a positive probability distribution, but wavefunctions need not have strictly non-negative components.[3] Sampling negative values is possible by expanding configuration space to include both the original configuration C and a sign $s = \pm$. One then samples configuration C with weight $|\langle C|\psi_0\rangle|$ and sign $s = \mathrm{sgn}\langle C|\psi_0\rangle$. However, this "fix" opens up its own can of worms, as we will now show.

Consider two non-interacting particles hopping on a one-dimensional L-site lattice, with the same Hamiltonian as Eq. 3.17. Treating the particles as distinguishable, if we start the system with particle A at site 1 and particle B at site 2 (denoted $|12\rangle$) and evolve in imaginary time, the particles will move independently according to H. Histogramming their positions (Fig. 3.2), each approaches the (single-particle)

[3] In fact, wavefunctions need not have strictly *real* components. However, for a real-valued Hamiltonian matrix, all eigenstates can be written as real-valued vectors in the Hilbert space.

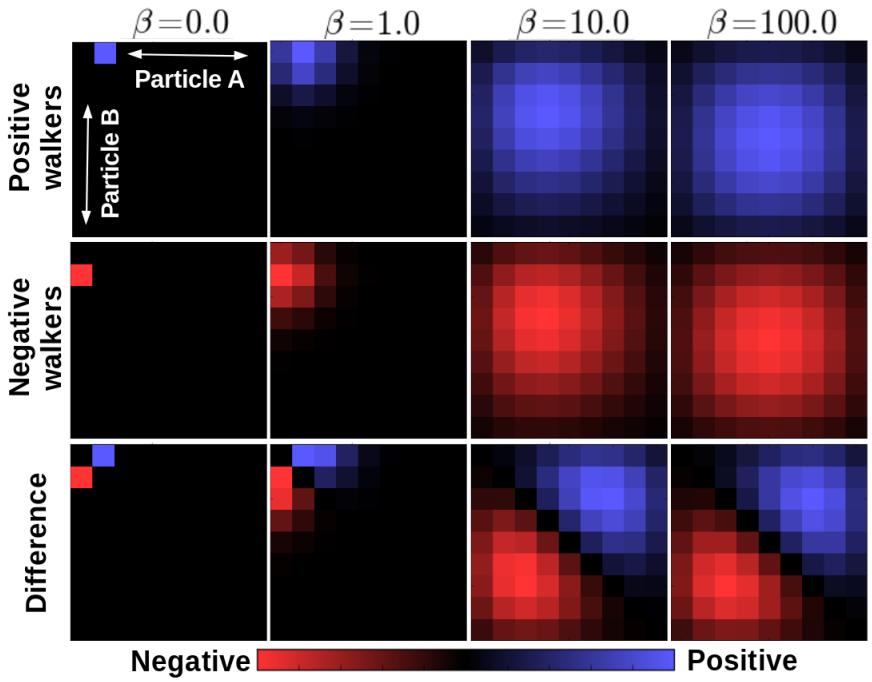

Figure 3.2: Probability distribution of positive and negative walkers for the example of the sign problem given in the text. The distributions of positive and negative walkers independently approach the bosonic ground state, while their difference approaches the fermionic ground state.

ground state of H: $\langle j|\psi_0^{\rm sp}\rangle = \sin\left(\frac{\pi j}{L+1}\right)$. Physically, this is acceptable if the particles are bosons, as such a ground state is symmetric under particle exchange.

However, if the particles are fermions, they should not both end up in the single particle ground state. In this case, we would have to start the simulation from an anti-symmetric state, such as

$$|\psi_T\rangle = |12\rangle - |21\rangle \ . \tag{3.18}$$

Algorithmically, this corresponds to starting with two walkers, one at $|12\rangle$ with a positive sign and the other at $|21\rangle$ with a negative sign. However, the argument above still works separately for the positive walkers and the negative walkers: each individually approaches the bosonic ground state, while their difference approaches the fermionic ground state (see Fig. 3.2).

To see how the fermionic component emerges, we rewrite the positive initial state as

$$|\varphi_+\rangle \equiv |12\rangle = \frac{1}{2}\big[|12\rangle - |21\rangle\big] + \frac{1}{2}\big[|12\rangle + |21\rangle\big] \ , \tag{3.19}$$

in which the first term is anti-symmetric (fermionic) and the second term is symmetric (bosonic). If the single particle orbitals of H are ϕ_0 (ground state), ϕ_1 (first excited state), etc., then the many-body ground state of H is the bosonic state

$$|b\rangle = |\phi_0(A)\phi_0(B)\rangle \tag{3.20}$$

with energy E_b. For simplicity, consider $E_b = 0$, which can be achieved by adding a constant offset to H. The fermionic ground state is

$$|f\rangle = \frac{1}{\sqrt{2}}\Big(|\phi_0(A)\phi_1(B)\rangle - |\phi_0(B)\phi_1(A)\rangle\Big) \ , \tag{3.21}$$

with energy $E_f > 0$. Since H conserves symmetry, after long time β the positive walkers sample from the (positive-valued) wavefunction

$$e^{-\beta H}|\varphi_+\rangle = \frac{1}{2}\left[|b\rangle + e^{-\beta E_f}|f\rangle\right] \ . \tag{3.22}$$

Since $E_f > 0$, the bosonic piece dominates in the $\beta \to \infty$ limit. Similarly, the negative walkers sample from

$$e^{-\beta H}|\varphi_-\rangle = \frac{1}{2}\left[-|b\rangle + e^{-\beta E_f}|f\rangle\right] \ . \tag{3.23}$$

Adding these two together, we see that the physical solution is that $e^{-\beta H}|\psi_T\rangle$ approaches the fermionic state $e^{-\beta E_f}|f\rangle$, as it must. However, in the actual simulation the positive and negative walkers evolve independently without talking to each other; this means that the physical cancellation of the bosonic parts of the wavefunction does not occur in the simulation. Therefore, while the wavefunction that we sample

is fermionic, there remains bosonic noise whose weight, as a function of β, increases exponentially faster than the fermionic signal. This is what we mean by the Fermi sign problem.

A more detailed discussion of a particular ground state QMC implementation can be found in Chp. 4. In particular, we address ways of minimizing the Fermi sign problem and develop a series of new tools for full configuration-interaction QMC, a relatively new QMC algorithm developed for quantum chemistry applications [43]. We also discuss a more general version of the sign problem, and its connection to (basis-dependent) properties of the Hamiltonian matrix.

3.2 Density matrix renormalization group

A second useful numerical technique that we will discuss is the density matrix renormalization group (DMRG), invented by Steve White [44, 45] as a more efficient tool to replace Wilson's numerical renormalization group [46]. As with any renormalization procedure, DMRG attempts to coarse-grain the system, throwing away the least important degrees of freedom.

The efficiency of DMRG comes from working with a special class of variational wavefunctions known as *matrix product states*. Indeed, many contemporary approaches to DMRG treat it primarily as an efficient technique for searching within the variational subspace of matrix product states. This is the approach we will take in the following sections, based on the review of these methods by Schollwöck [47]. As we will see, matrix product states allow efficient numerical optimization and measurement of a wide range of wavefunctions. The matrix product state representation of a wavefunction is closely tied to the study of entanglement, providing a physical connection between this ansatz and many topics of current interest.

Given the efficiency of matrix product states in variational ground state searches, DMRG has become the technique of choice for solving one-dimensional problems [48, 49]. It has recently been adapted quite successfully for a number of two-dimensional problems [50, 51], although the scaling of entanglement in two dimensions (and thus, as we'll see, the efficiency of DMRG) is exponentially worse than that of one dimension.

3.2.1 Matrix product states

A matrix product state (MPS) is a representation of an arbitrary wavefunction $|\psi\rangle$ as a product of matrices. Consider the Hilbert space $\mathcal{H} = \{|\vec{\sigma}\rangle = |\sigma_1 \dots \sigma_L\rangle\}$ of an L-site spin system, where σ enumerates d spin components (for example, $d = 2$ in a spin-1/2 system). Then the standard MPS ansatz with open boundary conditions is

$$|\psi\rangle = \sum_{\vec{\sigma}} M_1^{\sigma_1} M_2^{\sigma_2} \cdots M_L^{\sigma_L} |\vec{\sigma}\rangle \ . \tag{3.24}$$

Here, for each site i and spin σ_i (except $i = 1, N$; see below), we represent $|\psi\rangle$ by a $D \times D$ matrix $M_i^{\sigma_i}$, where D is an arbitrary integer known as the *bond dimension*. Since the coefficient $c_{\vec{\sigma}} = M^{\sigma_1} \cdots M^{\sigma_L}$ is a complex number, the first and last matrices M^{σ_1} and M^{σ_L} need to be $1 \times D$ row and $D \times 1$ column vectors respectively. There also exists an ansatz more suited to periodic boundary conditions [52], which we will not address here.

For the remainder of this chapter, we use two pieces of notational shorthand:

- We write M^{σ_i} instead of $M_i^{\sigma_i}$.
- We utilize the Einstein summation convention (summing over repeated indices).

A useful graphical representation of an MPS is shown in Fig. 3.3a. The matrix is represented by a black circle, with one leg corresponding to each index needed to

specify a matrix element M^{σ}_{ij}. The product of two matrices M and N is

$$\left[M^{\sigma}N^{\sigma'}\right]_{ij} = M^{\sigma}_{ik}N^{\sigma'}_{kj} \ ; \tag{3.25}$$

this is represented graphically by connecting the summed-over legs of M^{σ} and $M^{\sigma'}$ (Fig. 3.3b), leaving four floating legs σ, σ', i, and j. In this representation, the only floating parts of a matrix product state are its spin indices; the internal matrix indices are summed over ("contracted"), as in Fig. 3.3c.

Efficient manipulation of MPSs involves two basic operations: proper ordering of contractions and singular value decompositions. Ordering of contractions can be illustrated by considering the overlap between wavefunctions represented as MPSs:

$$\langle\psi'|\psi\rangle = \sum_{\sigma_1\ldots\sigma_L} (M'^{\sigma_L})^{\dagger}\cdots(M'^{\sigma_1})^{\dagger}M^{\sigma_1}\cdots M^{\sigma_L} \ . \tag{3.26}$$

A naive contraction in which, for each $\vec{\sigma}$, we perform the matrix multiplication and then add the $\vec{\sigma}$ together, requires a prohibitive number of steps (of order $d^L LD^3$). However, distributing the σ sums we can rewrite Eq. 3.26 as

$$\langle\psi'|\psi\rangle = \sum_{\sigma_L}\left[(M'^{\sigma_L})^{\dagger}\sum_{\sigma_{L-1}}\left[\cdots(M'^{\sigma_2})^{\dagger}\sum_{\sigma_1}\left[(M'^{\sigma_1})^{\dagger}M^{\sigma_1}\right]M^{\sigma_2}\cdots\right]M^{\sigma_L}\right] \ . \tag{3.27}$$

Each σ sum requires $O(dD^3)$ steps (d matrix multiplications). Thus, by evaluating the overlap through contracting these sums from the inside out, this method requires $O(dLD^3)$ steps, rendering it computationally feasible. Graphically, this corresponds to evaluating the vertical contractions before the horizontal ones, sweeping across the MPS from left to right (Fig. 3.3e).

The second tool in the MPS toolkit is the singular value decomposition. Consider an arbitrary matrix M of size $r\times c$, with $r \leq c$ for simplicity. M can always be written in the form $M = USV^{\dagger}$ [38], known as its singular value decomposition (SVD), where

U, S, and V are of size $r \times r$, $r \times r$, and $c \times r$ respectively. U and V satisfy the canonicity condition $U^\dagger U = V^\dagger V = \mathbb{1}_r$, while S is a non-negative real-valued diagonal matrix, known as the singular value matrix.

One illustration of the power of the SVD is using it to write an arbitrary wavefunction

$$|\psi\rangle = \sum_{\vec{\sigma}} c_{\sigma_1 \cdots \sigma_L} |\vec{\sigma}\rangle \tag{3.28}$$

as an MPS. We can think of the coefficient $c_{\sigma_1 \cdots \sigma_L}$ as a rank-L tensor, or by combining indices we can reshape it into a matrix C_{σ_1, α_1}, where $\alpha_1 = (\sigma_2, \ldots, \sigma_L)$ is a super-index that iterates over all spins besides σ_1. Performing an SVD on C_{σ_1, α_1}, we get

$$C_{\sigma_1, \alpha_1} = U_{\sigma_1 j} S_j V^*_{\alpha_1 j} \ , \tag{3.29}$$

using S_j to denote the j-th diagonal element of the singular value matrix S. We can reshape U to get d row vectors $A^{\sigma_1}_{1j} = U_{\sigma_1 j}$, which will be the matrices at site 1 in the MPS.

Similarly, we can multiply the S and V matrices and reshape to get new "coefficient" $c_{j,\sigma_2,\ldots,\sigma_L} = S_j V^*_{(\sigma_2,\ldots,\sigma_L)j}$. We can then continue performing the SVD site-by-site, reshaping the U matrices to get the desired A^σ matrices for the MPS. At first glance it would appear that the matrices at site $\ell + 1$ are d times larger than those on site ℓ; for example, at site 2 we would reshape $c_{j,\sigma_2,\ldots,\sigma_L}$ to $C_{(j,\sigma_2),(\sigma_3,\ldots)}$, yielding a $d^2 \times d^2$ matrix for U. However, a closer analysis reveals that the size increases for a bit, but eventually must come back down to give a column vector on site L. Therefore, the size of the A matrices follows the progression: $1 \times d$ on site 1, $d \times d^2$ on site 2, ..., $d^{L/2-1} \times d^{L/2}$ on site $L/2$, $d^{L/2} \times d^{L/2-1}$ on site $L/2+1$, ..., $d \times 1$ on site L. Note that, while A^{σ_1} never needs to be larger than $1 \times d$, we can pad it with zeros to get a $1 \times D$ matrix if desired.

The matrices generated by such an SVD have some useful properties. Consider the canonicity condition, $U^\dagger U = \mathbb{1}$. After reshaping, this corresponds to

$$\sum_\sigma (A^\sigma)^\dagger A^\sigma = \mathbb{1} \ . \tag{3.30}$$

Matrices satisfying this property are referred to as left canonical. The above procedure generates left canonical matrices on each site except for site L.[4] Similarly, one can define right canonical matrices satisfying

$$\sum_\sigma B^\sigma (B^\sigma)^\dagger = \mathbb{1} \ . \tag{3.31}$$

These are the matrices that one would generate by sweeping from site L to site 1 (right to left) and appropriately reshaping the V matrix. For the rest of this chapter, matrices denoted A are left canonical and those denoted B are right canonical.

There is a special class of MPSs called "mixed canonical" matrix products states, which are of the form $\cdots AA\Psi BB\cdots$ where only one site ℓ has non-canonical matrices Ψ^{σ_ℓ}. If $\Psi^{\sigma_\ell}_{ij}$ is a $D \times D$ matrix, consider an SVD on the $D \times dD$ matrix $\Psi_{i(j,\sigma_\ell)} = U_{i\alpha} S_\alpha V^*_{(j,\sigma_\ell),\alpha}$. We can define $A'^{\sigma_{\ell-1}} = A^{\sigma_{\ell-1}} U$; it is easy to check that A' is also left canonical. We can also reshape V to get a right canonical matrix $B^{\sigma_\ell}_{\alpha,j} = V_{(j,\sigma_\ell),\alpha}$. Therefore, our MPS has been rewritten as $\cdots AA'SBB\cdots$, where S is the singular value matrix which we can think of as residing on the bond between sites $\ell - 1$ and ℓ. The wavefunction can be rewritten as

$$|\psi\rangle = \sum_\alpha S_\alpha |\alpha_A\rangle |\alpha_B\rangle \ , \tag{3.32}$$

[4] A careful analysis of the procedure shows that the site L is not left canonical because no SVD is ever performed at site L.

where

$$\begin{aligned} |\alpha_A\rangle &= \sum_{\sigma_1 \ldots \sigma_{\ell-1}} \left[A^{\sigma_1} \cdots A^{\sigma_{\ell-2}} A'^{\sigma_{\ell-1}}\right]_{1\alpha} |\sigma_1 \cdots \sigma_{\ell-1}\rangle \\ |\alpha_B\rangle &= \sum_{\sigma_\ell \ldots \sigma_L} \left[B^{\sigma_\ell} \cdots B^{\sigma_L}\right]_{\alpha 1} |\sigma_\ell \cdots \sigma_L\rangle \ . \end{aligned}$$

Eq. 3.32 is precisely a Schmidt decomposition if the $|\alpha\rangle$'s are orthonormal; this is where we utilize canonicity.

$$\begin{aligned} \langle \beta_B | \alpha_B \rangle &= \left(\sum_{\sigma_\ell} B^{\sigma_\ell} \left[\cdots \sum_{\sigma_L} B^{\sigma_L} \left(B^{\sigma_L}\right)^\dagger \cdots \right] \left(B^{\sigma_\ell}\right)^\dagger \right)_{\alpha\beta} \\ &= \left(\sum_{\sigma_\ell} B^{\sigma_\ell} \left[\cdots \sum_{\sigma_L} B^{\sigma_{L-1}} \mathbb{1} \left(B^{\sigma_{L-1}}\right)^\dagger \cdots \right] \left(B^{\sigma_\ell}\right)^\dagger \right)_{\alpha\beta} \\ = \ \ldots &= \delta_{\alpha\beta} \ . \end{aligned}$$

A similar proof shows that the $|\alpha\rangle_A$ states are orthonormal as well.

The Schmidt decomposition in Eq. 3.32 is the basis for physical insight into properties of the MPS. For instance, consider an MPS $|\psi\rangle$ of bond dimension D. In addition, assume that it is normalized, so that $\sum_\alpha S_\alpha^2 = 1$. Then the von Neumann entanglement entropy of $|\psi\rangle$ across the bond between sites $\ell - 1$ and ℓ is

$$S = -\sum_\alpha S_\alpha^2 \log\left(S_\alpha^2\right) \ . \tag{3.33}$$

As there are D values of α in the sum, the maximum entanglement entropy of such an MPS is $S_{\max} = \log D$.

Another consequence of the Schmidt decomposition is to enable compression. Suppose we wish to approximate an MPS $|\psi'\rangle$ with bond dimension D' by some $|\psi\rangle$ with smaller bond dimension D. The Schmidt decomposition suggests a good way to do this: perform the SVD, then throw out terms in the Schmidt decomposition corre-

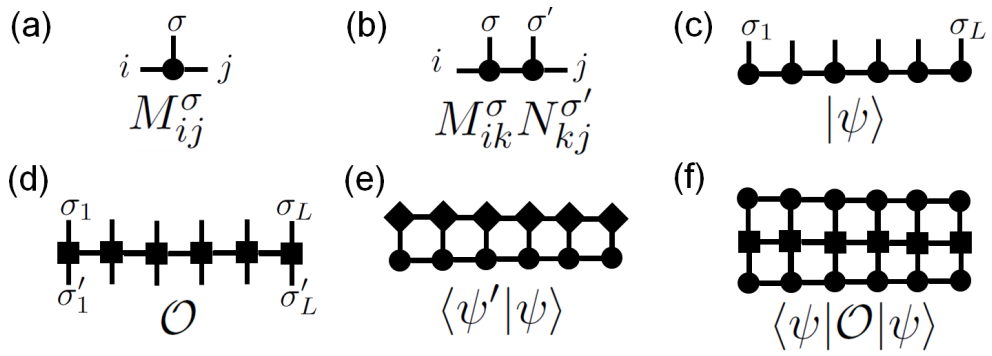

Figure 3.3: Graphical representation of (a) a matrix within the MPS, (b) multiplication of two matrices, (c) a matrix product state, (d) a matrix product operator, (e) a wavefunction overlap, and (f) an expectation value. Circles/squares/diamonds represent matrices, lines represent indices, and lines connecting different matrices are summed over.

sponding to the smallest values of S_α.[5] This procedure is an optimal compression of the matrices on that one site, in that it minimizes the magnitude of the difference between states, $\big|\big||\psi\rangle - |\psi'\rangle\big|\big|^2$. We refer to this procedure as SVD compression of the MPS from bond dimension D' to D.

3.2.2 Matrix product operators

In addition to matrix product states, one can define a matrix product operator (MPO), whose pictorial representation is shown in Fig. 3.3d. Noting that operators are an outer product of a ket with a bra, an MPO is very similar to an MPS, except that matrices in an MPO have two spin indices per site. Thus, a generic MPO is written

$$\mathcal{O} = \sum_{\vec{\sigma}\vec{\sigma}'} M^{\sigma_1\sigma'_1} \cdots M^{\sigma_L\sigma'_L} |\vec{\sigma}\rangle\langle\vec{\sigma}'| \ . \tag{3.34}$$

Since one can rewrite the MPO in the form of an MPS by wrapping the spin indices σ, σ' into a single super-index (σ, σ') of spin dimension d^2, all of the tricks developed for MPSs can be applied to MPOs, including SVD compression. While it is not

[5] From an entanglement perspective, this corresponds to removing the components of $|\psi'\rangle$ with the highest entanglement energy [53].

obvious that Hamiltonians or other physical observables can be expressed in the form of an MPO with small bond dimension, it turns out to be the case that most local observables can be written as MPOs of bond dimension $D_{\mathcal{O}} \lesssim 5$ [47].

An important operation involving an MPO $\mathcal{O}$ is acting with it on some MPS $|\psi\rangle$ to get $|\psi'\rangle = \mathcal{O}|\psi\rangle$. As with wavefunction overlaps, this can be done very efficiently by appropriate ordering of contractions, which is easily seen in the graphical MPS/MPO representation (Fig. 3.3f). Mathematically, this operation corresponds to taking tensor products of the MPO matrices $M^{\sigma_j \sigma'_j}$ with those of the MPS $\Psi^{\sigma'_j}$ and summing over the spin index σ'_j. An important consequence of such an operation is to increase the bond dimension of the resultant MPS; if $|\psi\rangle$ has bond dimension D and $\mathcal{O}$ has bond dimension $D_{\mathcal{O}}$, then $|\psi'\rangle$ will have bond dimension $D_{\mathcal{O}}\, D$.

Having established the basic techniques used to manipulate and compress matrix product states and operators, we proceed to explain how they can be used for time evolution and ground state calculations.

3.2.3 Time evolution via MPS

Time evolution of an MPS turns out to be a rather straightforward operation. There are many variants to implement this idea, including time-evolving block decimation (TEBD [54, 55]), time-dependent DMRG (tDMRG [56, 57]), and the more recently developed time-dependent MPS (tMPS [58]) algorithm that we discuss below. While there are some technical differences between these algorithms, they all operate on a similar premise.

We wish to use MPSs to solve for the time evolution of some initial state $|\psi(0)\rangle$. For simplicity, we consider the case of a time-independent Hamiltonian H. Then, at time t, according to the Schrödinger equation,

$$|\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle \ . \tag{3.35}$$

In order to use the MPS formalism that we have developed, we must show how to express the unitary evolution operator $U(t) = \exp(-iHt)$ as an MPO. While this is not in general a trivial thing to do exactly, it can be accomplished to arbitrary accuracy using a Trotter decomposition [59].

More explicitly, consider some one-dimensional Hamiltonian H with at most nearest neighbor interactions, meaning we can break it up into a sum of terms acting on even bonds and terms acting on odd bonds,

$$H = H_{\text{odd}} + H_{\text{even}} = \sum_{j\,\text{odd}} H_j + \sum_{j\,\text{even}} H_j \ , \tag{3.36}$$

where H_j acts only on sites j and $j+1$. Since they act on different sites, the terms within H_{odd} or H_{even} commute with each other, i.e., $[H_j, H_{j+2n}] = 0$ for $n \in \mathbb{Z}$. This is useful because time evolution becomes separable; looking at the odd sites, for example,

$$U_{\text{odd}}(t) = e^{-iH_{\text{odd}}t} = e^{-i\sum_{j\,\text{odd}} H_j t} = \prod_{j\,\text{odd}} e^{-iH_j t} \ . \tag{3.37}$$

As we will soon see, this product can be easily expressed as an MPO.

However, H_{odd} and H_{even} do not commute with each other. Therefore, in order to apply, $e^{-iHt} = e^{-i(H_{\text{odd}}+H_{\text{even}})t}$, we must first perform a Trotter decomposition. To do so, consider a small time step $\tau = t/N$, such that $e^{-iHt} = \left(e^{-iH\tau}\right)^N$. For small τ,

$$e^{-iH\tau} = e^{-i(H_{\text{odd}}+H_{\text{even}})\tau} = e^{-iH_{\text{odd}}\tau}e^{-iH_{\text{even}}\tau}e^{-i[H_{\text{odd}},H_{\text{even}}]O(\tau^2)} \ . \tag{3.38}$$

Dropping the $O(\tau^2)$ commutator term, this is the first order Trotter breakup. Higher order Trotter breakups give higher order corrections in τ, but regardless the general principle is to use such an approximation and extrapolate to the $\tau = 0$ limit to solve for the correct time evolution.

As an example of how to construct an MPO corresponding to $U_j(\tau) \equiv e^{-iH_j\tau}$, we use the case of an Ising interaction $H = J\sum_j s_j^z s_{j+1}^z$ where $s^{x,y,z}$ are the Pauli matrices; more general nearest neighbor interactions can be done in a similar fashion. Consider $U_j(\tau) = e^{-iJs_j^z s_{j+1}^z \tau}$. Since this is diagonal in the s^z basis,

$$\langle \uparrow_j, \sigma_{j+1} | U_j | \downarrow_j, \sigma'_{j+1} \rangle = 0 \tag{3.39}$$

for all values of spins σ_{j+1} and σ'_{j+1} on site $j+1$. In terms of the MPOs, this implies that $M_j^{\uparrow\downarrow} = M_{j+1}^{\uparrow\downarrow} = M_j^{\downarrow\uparrow} = M_{j+1}^{\downarrow\uparrow} = 0$. To get the correct matrix elements, we need the remaining matrices to satisfy

$$M_j^{\uparrow\uparrow} M_{j+1}^{\uparrow\uparrow} = M_j^{\downarrow\downarrow} M_{j+1}^{\downarrow\downarrow} = e^{-iJ\tau}, \quad M_j^{\uparrow\uparrow} M_{j+1}^{\downarrow\downarrow} = M_j^{\downarrow\downarrow} M_{j+1}^{\uparrow\uparrow} = e^{iJ\tau} \ . \tag{3.40}$$

Note that the product of these matrices is a complex number, meaning that we want M_j to be a row vector and M_{j+1} to be a column vector. The solution to these equations is not unique, but one choice that works is

$$M_j^{\uparrow\uparrow} = \left(e^{iJ\tau}, \, e^{-iJ\tau} \right) \quad , \quad M_{j+1}^{\uparrow\uparrow} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \tag{3.41}$$

$$M_j^{\downarrow\downarrow} = \left(e^{-iJ\tau}, \, e^{iJ\tau} \right) \quad , \quad M_{j+1}^{\downarrow\downarrow} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \ . \tag{3.42}$$

Due to the product form of $U_{\text{odd}} = \prod_{j\,\text{odd}} U_j$ (similarly for U_{even}), the entire MPO for U_{odd} is constructed by using the above ansatz for each odd bond. This gives an MPO of bond dimension $D_U = 2$, with matrices of size 1×2 on the odd sites and size 2×1 on the even sites.[6]

[6] One can shape this into the more conventional form of 2×2 square matrices by padding the matrices with zeros to the bottom right.

The procedure is then, starting from an MPS $|\psi(0)\rangle$ at time zero, apply the MPO versions of $U_{\text{even}}(\tau)$ and $U_{\text{odd}}(\tau)$ sequentially to get

$$|\psi(t)\rangle \approx U_{\text{odd}}(\tau)U_{\text{even}}(\tau)\cdots U_{\text{even}}(\tau)|\psi(0)\rangle \ , \tag{3.43}$$

where the approximation comes from the Trotter breakup error, becoming exact in the limit $\tau \to 0$. However, if $|\psi(0)\rangle$ has bond dimension D, then each application of an MPO corresponds to a tensor product, multiplying the bond dimension by a factor of 2 in the example above. So after time $t = N\tau$, the bond dimension blows up to $4^N D$ if the MPS is stored exactly. The solution to this problem has already been suggested; at each step apply the SVD compression procedure to reduce the bond dimension back down to a manageable quantity D_{feasible}. However, as the entanglement grows in time, we need to continually increase D_{feasible} to accurately capture the quantum state. At the cost of this controlled approximation, the time evolution is simulated efficiently, becoming exact in the limit of $\tau \to 0$ and $D_{\text{feasible}} \to \infty$.

3.2.4 Ground states via MPS

In place of real time evolution, one can use the methods of the previous section to perform imaginary time evolution $|\psi(\beta)\rangle = e^{-\beta H}|\psi(0)\rangle$. As we saw in the context of QMC, taking the limit $\beta \to \infty$ of this expression is one method for finding the ground state of H. However, for matrix product states, a much more efficient algorithm exists: the density matrix renormalization group (DMRG) [44, 45]. While not precisely equivalent to DMRG, we will discuss a closely-related MPS algorithm for finding the ground state of some Hamiltonian H expressed in MPO form: single-site ground state MPS (gsMPS [47]). As we do not utilize this algorithm elsewhere in the thesis, we only briefly describe it here for the sake of completeness and to highlight the connection between MPS algorithms and DMRG.

The idea behind single-site gsMPS is to note that, if we fix the matrices on all but site j, then the problem of optimizing the matrices on site j is linear, and thus easily solved exactly. The algorithm is then to sweep back-and-forth across the sites, from site 1 to site L and back again, at each step optimizing the matrices on a single site. While these local optimization steps are not guaranteed to give the global minimum for a fixed bond dimension D, empirically they often not only do so, but do so much more efficiently than imaginary time evolution [47].

Consider a mixed canonical state $|\psi\rangle = \cdots A^{\sigma_{\ell-1}} \Psi^{\sigma_\ell} B^{\sigma_{\ell+1}} \cdots$ and a Hamiltonian H expressed as an MPO with matrices M. We wish to optimize the energy $E = \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle}$ or, using a Lagrange multiplier λ, to optimize $\langle\psi|H|\psi\rangle - \lambda\langle\psi|\psi\rangle$ where λ will give the energy. The second term is particularly clean for a mixed canonical state, namely

$$\langle\psi|\psi\rangle = \Psi^{\sigma}_{ij} \left(\Psi^{\sigma}_{ij}\right)^* \ , \tag{3.44}$$

since the A's and the B's contract to give identities as we saw earlier. To avoid some math, we evaluate the first term using the graphical representation, which tell us that, using optimal contractions, the expectation value within part A (sites 1 to $\ell - 1$) is just a rank-3 tensor L_{ijk} "=" $\langle\psi_A|H_A|\psi_A\rangle$ (see Fig. 3.4). Similar, part B contracts to give a tensor $R_{i'j'k'}$. While we do not detail this here, these tensors can be efficiently created and updated during the gsMPS sweeps. Contracting with the matrices at site ℓ, we get

$$\langle\psi|H|\psi\rangle - \lambda\langle\psi|\psi\rangle = L_{ijk} \Psi^{\sigma'}_{ii'} \left(\Psi^{\sigma}_{kk'}\right)^* M^{\sigma\sigma'}_{jj'} R_{i'j'k'} - \lambda \Psi^{\sigma}_{ij} \left(\Psi^{\sigma}_{ij}\right)^* \ . \tag{3.45}$$

Note that the matrix elements of Ψ appear strictly quadratically; this is the hallmark of a linear problem that can be easily solved. Taking the derivate with respect $(\Psi^{\sigma}_{kk'})^*$

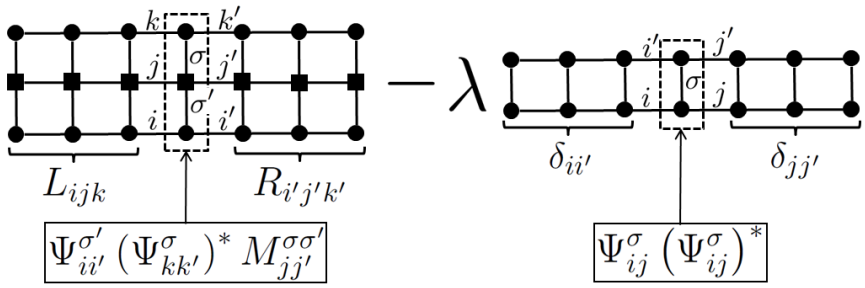

Figure 3.4: Graphical representation of the ground state MPS algorithm, in which one optimizes $\langle\psi|H|\psi\rangle - \lambda\langle\psi|\psi\rangle$. The expectation value of H is split up into two rank-3 tensors, L and R, which can be efficiently updated during the sweeps. The overlap $\langle\psi|\psi\rangle$ is even more straightforward, as the contraction of appropriately canonized matrices gives an identity matrix.

and setting it to zero, we get the linear equations

$$L_{ijk}\Psi^{\sigma'}_{ii'}M^{\sigma'\sigma}_{jj'}R_{i'j'k'} - \lambda\Psi^{\sigma}_{kk'} = 0\ . \tag{3.46}$$

We can see that this is just a linear eigenvalue problem by combining indices: defining eigenstate $\psi_{(\sigma,k,k')} = \Psi^{\sigma}_{kk'}$ of matrix $\mathbf{H}_{(\sigma,k,k'),(\sigma',i,i')} = L_{ijk}M^{\sigma'\sigma}_{jj'}R_{i'j'k'}$, we see that $\mathbf{H}\psi = \lambda\psi$. The ground state of the $dD^2 \times dD^2$ matrix $\mathbf{H}$ is then generally found via the Lanczos algorithm, and subsequently reshaped to give a new matrix Ψ' of strictly lower energy. One then performs an SVD to make site ℓ canonical, moves onto the next site, and continues to sweep back and forth until the ground state is reached.

While very clean, the above single-site gsMPS procedure is prone to getting stuck in local energy minima. A number of procedures exist that are able to skirt this issue, including standard DMRG and two-site gsMPS [60]. One state-of-the-art procedure, known as multi-site targeting, can be easily implemented on top of single-site gsMPS [61]. Essentially, in addition to the optimal single-site matrix Ψ', one mixes in some higher excitations by adding states that can be reached via a single application of H to the optimized MPS with Ψ replaced by Ψ' . As the algorithm progress, the weight of these excitations is decreased, until eventually one is just doing single-site gsMPS.

As this procedure is not used in this work, we refer the reader elsewhere for details [47, 61].

Chapter 4

Understanding and improving full configuration-interaction QMC

Full configuration-interaction quantum Monte Carlo (FCIQMC) [43] is a recently-developed ground state QMC method, similar to those discussed in Sec. 3.1.2, which is based on the chemistry version of exact diagonalization known as the full configuration-interaction method. It improves in two major respects over similar projection techniques such as diffusion Monte Carlo (DMC [62], essentially the continuous-space analogue of the particle-in-a-box example in Sec. 3.1.2). First, FCIQMC works in a discrete as opposed to a continuous basis; originally it was conceived for use with Slater determinants of atomic or molecular orbitals, but it is naturally extensible to other discrete bases such as momentum eigenstates of a finite box [63, 64]. While not strictly necessary, FCIQMC is also often performed in a second-quantized basis, i.e., Slater determinants for fermionic particles, whereas DMC traditionally works in a first-quantized formalism (see Sec. 4.1.3 for more details on this distinction).

The second distinguishing characteristic of FCIQMC, enabled by the discrete basis, is *annihilation*. As usual, the walkers are stored as signed basis elements, but if

two walkers end up the same basis element with opposite sign during the course of the simulation, those two walkers annihilate and are removed from the simulation. In the case of our simple example in Sec. 3.1.3, such an annihilation process could in principle diminish or remove the bosonic noise. As this annihilation process removes noise from the simulation, it is able to mitigate the sign problem [43], albeit at a cost that still scales exponentially with system size.

The remainder of this chapter is split into two sections. In Sec. 4.1 we introduce the FCIQMC algorithm and discuss aspects of the FCIQMC sign problem in the context of annihilation. As an important example, we discuss the differences between first and second quantized formalisms with respect to the FCIQMC sign problem, understanding how and when their sign problems differ. In Sec. 4.2 we present a number of improvements to the FCIQMC algorithm, using its application to the Fermi polaron as a test case. We discuss the advantages and limitations of these extensions, speculating on their applicability to more general problems in many-body theory.

4.1 Annihilation and the FCIQMC sign problem

While a number of QMC methods have been developed that attempt to solve fermion ground state problems exactly by cancelling some of the "bosonic" noise terms (e.g., Ref. 65), FCIQMC is one of the first to show real success in pushing past existing exact methods for some problems in quantum chemistry [43, 66]. The key aspect leading to the success of FCIQMC is annihilation, which enables filtering out some sign-problem-affiliated noise without affecting sampling of the fermionic ground state. Another notable aspect of FCIQMC is its capacity to work in a second-quantized basis of Slater determinants. In this section, we introduce the FCIQMC algorithm [43] and show how to understand the sign problem in terms of properties of the basis-

dependent Hamiltonian matrix [63]. We then analyze the effect of working in a basis of second quantized determinants, comparing it to the more standard first quantized formalism used in DMC. We show that from a standpoint of the sign problem, second quantization is at least as good as first quantization. Finally, we prove what conditions must hold for the first and second quantized sign problems to differ.

4.1.1 The FCIQMC algorithm

In FCIQMC, the configuration space consists of signed basis elements: $\Omega = \{(s, |\mathbf{i}\rangle)\}$, where $s = \pm$ and $|\mathbf{i}\rangle$ is a member of some many-body basis. We wish apply the operator $U = 1 - \tau[H - S]$ (see Sec. 3.1.2), where S is a tunable parameter that is used to control the number of walkers. FCIQMC starts with a set of walkers sampled from some initial state $|\psi_T\rangle$, such that configuration $C = (s, |\mathbf{i}\rangle)$ is sampled with probability proportional to $|\langle\mathbf{i}|\psi_T\rangle|$ and sign $s = \mathrm{sgn}\langle\mathbf{i}|\psi_T\rangle$. Given some population of walkers, a single stochastic application of the operator U consists of three basic steps: 1) off-diagonal spawning, 2) diagonal cloning/killing, and 3) annihilation. We store all walkers in a list $\mathcal{L}_{\mathrm{walkers}}$, with the exception of walkers spawned during a given step of FCIQMC, which are temporarily stored in $\mathcal{L}_{\mathrm{spawned}}$.

In the off-diagonal spawning step, we stochastically apply the off-diagonal components of U. Put another way, for each configuration $(s, |\mathbf{i}\rangle) \in \mathcal{L}_{\mathrm{walkers}}$, we wish to sample some other configuration $(s', |\mathbf{j}\rangle)$ with weight $w(\mathbf{j}|\mathbf{i}) = ss'\langle\mathbf{j}|U|\mathbf{i}\rangle = -\tau ss'\langle\mathbf{j}|H|\mathbf{i}\rangle$, where s' is chosen to make this weight positive. To implement this, first propose a move by selecting basis element $|\mathbf{j}\rangle \neq |\mathbf{i}\rangle$ at random from some known, normalized probability distribution $q(\mathbf{j}|\mathbf{i})$. Unlike in the Metropolis algorithm, q here must be chosen such that $q(\mathbf{j}|\mathbf{i}) \neq 0$ if $\langle\mathbf{j}|H|\mathbf{i}\rangle \neq 0$. Then choose s' such that $-\tau ss'\langle\mathbf{j}|H|\mathbf{i}\rangle \geq 0$. To make the algorithm sample with net weight $w(\mathbf{j}|\mathbf{i})$, we spawn with weight

$$w_S(\mathbf{j}|\mathbf{i}) = \frac{w(\mathbf{j}|\mathbf{i})}{q(\mathbf{j}|\mathbf{i})} = \frac{|\tau\langle\mathbf{j}|H|\mathbf{i}\rangle|}{q(\mathbf{j}|\mathbf{i})} \ , \tag{4.1}$$

meaning that we spawn an average of w_S walkers on $|\mathbf{j}\rangle$; if $w_S > 1$, then we deterministically spawn $\lfloor w_S \rfloor$ walkers and spawn one additional walker stochastically with probability $w_S - \lfloor w_S \rfloor$. All spawned walkers are added to $\mathcal{L}_{\text{spawned}}$.

Similarly, the diagonal clone/kill step involves (stochastically) multiplying each configuration $(s, |\mathbf{i}\rangle) \in \mathcal{L}_{\text{walkers}}$ by weight $w_D(\mathbf{i}) = \langle \mathbf{i}|U|\mathbf{i}\rangle = 1 - \tau(\langle \mathbf{i}|H|\mathbf{i}\rangle - S)$. There are two possible scenarios for w_D: either $w_D > 1$ or $1 \geq w_D \geq 0$.[1] If $w_D > 1$, then we wish to clone $w_D - 1$ walkers in the same configuration $(s, |\mathbf{i}\rangle)$. As before, this is done by deterministically spawning $\lfloor w_D - 1 \rfloor$ walkers and stochastically spawning one more with probability $(w_D - 1) - \lfloor w_D - 1 \rfloor$. If $w_D < 1$, then we kill the walker, i.e., remove it from the list $\mathcal{L}_{\text{walkers}}$, with probability $1 - w_D$. During the simulation, the offset S is adjusted through standard feedback routines [43] to keep the walker number close to some desired value N_w.

Having stochastically applied both the diagonal and off-diagonal components of U, the list $\mathcal{L}_{\text{spawned}}$ is merged into $\mathcal{L}_{\text{walkers}}$, which now contains a snapshot of $U^N|\psi_T\rangle$. We then perform the crucial final step of FCIQMC: annihilation. Annihilation proceeds by examining the list $\mathcal{L}_{\text{walkers}}$ for any pairs on the same basis element with opposite sign, i.e., $(+, |\mathbf{i}\rangle)$ and $(-, |\mathbf{i}\rangle)$. As these two walkers add together to zero, and thus contribute nothing but noise, they are removed from the simulation. Annihilation proceeds until all such pairs have been removed from the list.

Finally, in order to retrieve properties of ground state wavefunction sampled by the simulation, we must measure the value of observables with respect to the wavefunction snapshot $|\psi_n\rangle$ after n steps of the algorithm; the most common such observable is the energy $E = \langle H \rangle$. Noting that H commutes with U, the energy is given by

$$E = \frac{\langle \psi_n|H|\psi_n\rangle}{\langle \psi_n|\psi_n\rangle} = \frac{\langle \psi_T|U^n H U^n|\psi_T\rangle}{\langle \psi_T|U^n U^n|\psi_T\rangle} = \frac{\langle \psi_T|HU^{2n}|\psi_T\rangle}{\langle \psi_T|U^{2n}|\psi_T\rangle} = \frac{\langle \psi_T|H|\psi_{2n}\rangle}{\langle \psi_T|\psi_{2n}\rangle} \ . \tag{4.2}$$

[1] While it is possible to treat the case where $w_D < 0$, we choose τ and S such that this scenario does not occur.

Since $|\psi_{2n}\rangle$ is the wavefunction we sample after $2n$ steps, it is stochastically given by a sum of walkers w at step $2n$: $|\psi_{2n}\rangle \to \sum_w s_w |\mathbf{i}_w\rangle$. Thus

$$E = \frac{\sum_w s_w \langle \psi_T | H | \mathbf{i}_w \rangle}{\sum_w s_w \langle \psi_T | \mathbf{i}_w \rangle} \ . \tag{4.3}$$

4.1.2 The FCIQMC sign problem

How can a sign problem emerge in the algorithm above? By construction, the only possibility for flipping the sign of a walker is through off-diagonal spawning; thus, the sign problem comes from spawning loops in which a walker, starting at $|\mathbf{i}\rangle$ with positive sign, spawns via the off-diagonal elements of U and eventually returns to $|\mathbf{i}\rangle$ with a negative sign. In the absence of such loops (for a given basis), the Hamiltonian does not have a sign problem, and the problem is relatively easy to solve. However, many problems of physical interest, including most fermionic Hamiltonians, do in fact have a sign problem.

For a more quantitative analysis of the sign problem, we turn to the work of Spencer et al. [63]. We will reproduce the important aspects of their argument here, as it is necessary for our later discussion. Consider the operator U in some fixed basis $\{|\mathbf{i}\rangle\}$, with matrix elements $\mathbf{U_{ij}} = \langle \mathbf{i} | U | \mathbf{j} \rangle$ and with S and τ chosen such that $\mathbf{U_{ii}} > 0$ for all $\mathbf{i}$. Break this matrix up into two components, $\mathbf{U} = \mathbf{U}^+ - \mathbf{U}^-$, such that $\mathbf{U}^{+(-)}$ has strictly non-negative elements. Then if $\mathbf{p}$ is a vector representing the positive walkers and $\mathbf{n}$ is a vector of negative walkers, the evolution described above is a stochastic version of

$$\begin{pmatrix} \mathbf{p}_{n+1} \\ \mathbf{n}_{n+1} \end{pmatrix} = \begin{pmatrix} \mathbf{U}^+ & \mathbf{U}^- \\ \mathbf{U}^- & \mathbf{U}^+ \end{pmatrix} \begin{pmatrix} \mathbf{p}_n \\ \mathbf{n}_n \end{pmatrix} \ . \tag{4.4}$$

The equations can be decoupled by considering the "in-phase" component $(\mathbf{p}+\mathbf{n})/2$ and the "out-of-phase" component $(\mathbf{p}-\mathbf{n})/2$, which go to the eigenstates with maxi-

mum eigenvalues of $\mathbf{U}^+ + \mathbf{U}^-$ and $\mathbf{U}^+ - \mathbf{U}^-$ respectively. We refer to this maximum eigenvalue as the "ground state" of the $\mathbf{U}^+ \pm \mathbf{U}^-$ matrix, as it is related to the ground state of the Hamiltonian matrix. The out-of-phase component is the ground state of H that we are after, while the in-phase component is the undesired noise. One can prove that the maximum eigenvalue $U^p_{\max}$ of $\mathbf{U}^+ + \mathbf{U}^-$, is at least as large as $U^n_{\max}$, the maximum eigenvalue of $\mathbf{U}^+ - \mathbf{U}^-$. Therefore the noise always grows faster than the signal, giving the sign problem. The sign problem's magnitude is set by the energy difference $\Delta U_{\max} = U^p_{\max} - U^n_{\max}$, since the signal to noise ratio grows exponentially at a rate $\Delta U_{\max}/\tau$.

Annihilation is slightly trickier in this language, but one important point can be seen fairly easily. Consider adding a constant noise vector $\mathbf{v}$ to both $\mathbf{p}$ and $\mathbf{n}$ simultaneously. In the absence of annihilation, this will affect the in-phase component, giving $(\mathbf{p} + \mathbf{n} + 2\mathbf{v})/2$, but will not affect the out-of-phase component. With perfect annihilation, the noise component $\mathbf{v}$ cancels itself out, so the in-phase component returns to $(\mathbf{p} + \mathbf{n})/2$ while the out-of-phase component remains unchanged. This is an illustration of a general point: annihilation impacts the evolution of the in-phase component (noise), but has no effect on the out-of-phase component (signal). A more in-depth analysis shows that, in certain limits, the effectiveness of annihilation depends on the same energy scale as does the sign problem, $\Delta U_{\max}$, becoming more effective (i.e., requiring fewer walkers to succeed) when the gap $\Delta U_{\max}$ is small [63].

4.1.3 First vs. second quantization in FCIQMC

Spencer et al.'s analysis of the sign problem, as discussed in the previous section, did not assume anything about the choice of basis. In particular, it did not require that the basis be first quantized (as in DMC [67]) or second quantized (as in the original FCIQMC calculations of Ref. 43). In this section, we will address this issue, answer-

ing two questions. First, when are the sign problems of first and second quantized algorithms different? Second, when they differ, which is better?

In order to be clear about what basis a quantity is expressed in, we will use the subscripts F, D, and P to indicate a first quantized (Hartree product) basis, a determinant (second quantized) basis, and a permanent basis respectively. A quantity without any subscript is general and can be considered in any of the three bases.

Let us consider a first quantized basis of Hartree products, $\{|\mathbf{i}\rangle\}$, for a system of N fermions, where $|\mathbf{i}\rangle = \phi_{i_1}\phi_{i_2}\cdots\phi_{i_N}$ and multiple occupancy of any single-particle basis function, ϕ_j, is forbidden. The second quantized basis is a set of determinants $\{|\mathbf{D_i}\rangle\}$. We use the notation $|\mathbf{i}\rangle \in |\mathbf{D_i}\rangle$ to indicate that $|\mathbf{i}\rangle$ is in the determinant expansion of $|\mathbf{D_i}\rangle$, regardless of the sign with which it appears. If $|\mathbf{i}\rangle \in |\mathbf{D_i}\rangle$, then $|\mathbf{D_i}\rangle$ can be obtained from $|\mathbf{i}\rangle$ using the anti-symmetrizer operator $\mathcal{A}$:

$$\mathcal{A}|\mathbf{i}\rangle = \frac{(-1)^\sigma}{\sqrt{N!}}|\mathbf{D_i}\rangle \ , \tag{4.5}$$

where σ is the number of permutations required to bring $|\mathbf{i}\rangle$ into coincidence with $|\mathbf{D_i}\rangle$. We utilize a similar notation for the (bosonic) basis of permanents, $\{|\mathbf{P_i}\rangle\}$, where $\mathcal{S}|\mathbf{i}\rangle = \frac{1}{\sqrt{N!}}|\mathbf{P_i}\rangle$, using the symmetrizer $\mathcal{S}$.

In the general case where the first quantized Hamiltonian may have a sign problem, the FCIQMC algorithm without annihilation is unstable with respect to the ground state of the $\mathbf{U}_F^+ + \mathbf{U}_F^-$ matrix. By the variational principle, this ground state is symmetric with respect to Hartree products related by permutation symmetry. The ground state of $\mathbf{U}_F^+ + \mathbf{U}_F^-$ is hence identical to the ground state of the same matrix expressed in a permanent basis. The second quantized FCIQMC algorithm is unstable to the ground state of $\mathbf{U}_D^+ + \mathbf{U}_D^-$.

When is the sign problem the same for these two different cases? In other words, when do $\mathbf{U}_F^+ + \mathbf{U}_F^-$ (in a permanent basis) and $\mathbf{U}_D^+ + \mathbf{U}_D^-$ (in a determinant basis)

have the same ground state energy? It is clear that one possible way for this to occur is if these are simply the same matrix, i.e., if

$$\langle \mathbf{P_i}| \left(\mathbf{U}_F^+ + \mathbf{U}_F^-\right) |\mathbf{P_j}\rangle = \langle \mathbf{D_i}| \left(\mathbf{U}_D^+ + \mathbf{U}_D^-\right) |\mathbf{D_j}\rangle \tag{4.6}$$

for all $\mathbf{i}$ and $\mathbf{j}$. What is true, but not immediately apparent, is that if these matrices do not match, then their ground state energies differ; we prove this in Appendix A. We also show that these conditions can be recast as a requirement that any two off-diagonal matrix elements of U taking a walker from $|\mathbf{i}\rangle$ to $|\mathbf{j}\rangle, |\mathbf{j}'\rangle \in |\mathbf{D_j}\rangle$ must do so with the same sign, after anti-symmetrization; we now discuss this alternative viewpoint of the sign problem.

One can compare the first quantized walkers to the second quantized walkers by mapping to second quantization using the anti-symmetrizer. This "filtering" process is naturally performed in a mixed estimator measurement of observable $\mathcal{O}$,

$$\langle \mathcal{O} \rangle_{\mathrm{ME}} = \frac{\langle \psi_T | \mathcal{O} | \psi \rangle}{\langle \psi_T | \psi \rangle} \ , \tag{4.7}$$

where $|\psi_T\rangle$ is some anti-symmetric trial wavefunction and $|\psi\rangle$ is the walker wavefunction. This is, for example, how energies are measured in FCIQMC (see Eq. 4.2). Any physical observable must conserve symmetry, so

$$[\mathcal{O}, \mathcal{A}] = 0 \implies \langle \psi_T | \mathcal{O} | \psi \rangle = \langle \psi_T | \mathcal{A}\mathcal{O} | \psi \rangle = \langle \psi_T | \mathcal{O}\mathcal{A} | \psi \rangle \ . \tag{4.8}$$

Thus only the "filtered" walker distribution, $\mathcal{A}|\psi\rangle$, enters into calculation of observables, so for understanding the sign problem we need only look at the distribution of first quantized walkers after anti-symmetrization.

Consider a first quantized walker starting on $|\mathbf{i}\rangle$ that can spawn to both $|\mathbf{j}\rangle$ and $|\mathbf{j}'\rangle$, where $|\mathbf{j}\rangle, |\mathbf{j}'\rangle \in |\mathbf{D_j}\rangle$. After a single step of the FCIQMC algorithm, it is possible

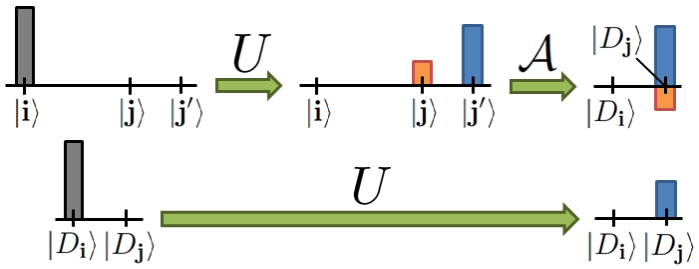

Figure 4.1: Example of a single step in which the first and second quantized FCIQMC algorithms differ. (Top) In the first quantized algorithm, a walker on $|\mathbf{i}\rangle$ spawns to $|\mathbf{j}\rangle, |\mathbf{j}'\rangle \in |\mathbf{D_j}\rangle$. After anti-symmetrizing, the walkers can add incoherently. (Bottom) Same step in the second quantized algorithm. The cancellation on $|\mathbf{D_j}\rangle$ is automatically accounted for by the sign and magnitude of $\langle \mathbf{D_i}|U|\mathbf{D_j}\rangle$.

for the antisymmetrizer to map these walkers to $|\mathbf{D_j}\rangle$ with opposite signs, resulting in a worse sign problem (see Fig. 4.1). FCIQMC in a determinant basis performs this cancellation automatically. Therefore, one would expect that the sign problems will only be identical if the first quantized walkers map to the same sign under the anti-symmetrizer. This argument also suggests that, if the sign problems differ, the second quantized algorithm will be better. We prove this as well in Appendix A, showing that the conditions above uniquely characterize when the first and second quantized sign problems differ.

There are a number of physical situations where the equivalence conditions in Eq. 4.6 can hold, including:

- the two-particle operator is diagonal in the single-particle basis;
- off-diagonal Hamiltonian matrix elements in the space of determinants only involve a single integral, i.e., the Hamiltonian matrix elements between determinants differing by a single orbital do not have a contribution from the two-particle operator and the Hamiltonian matrix elements between determinants differing by two orbitals contain only a Hartree term or only an exchange term.

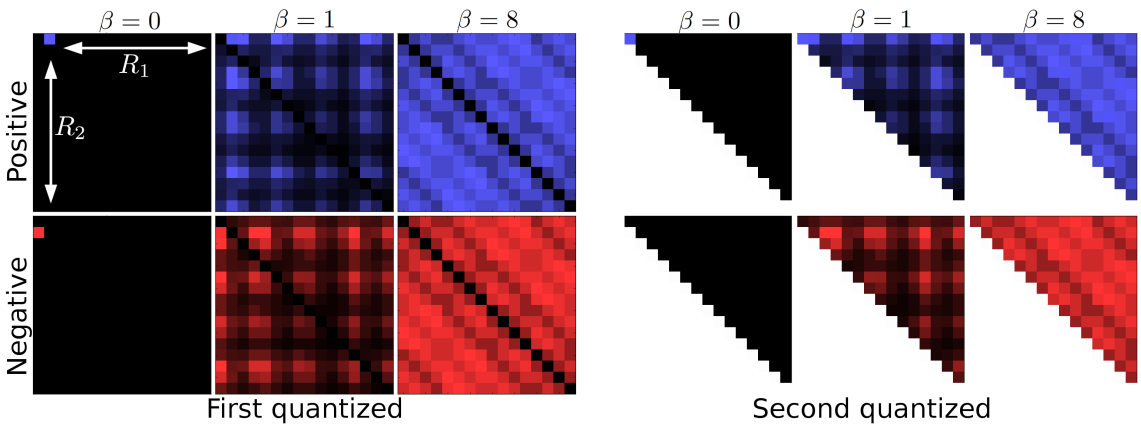

Figure 4.2: Probability distribution of walkers for a fermionic Hamiltonian where the first and second quantized problems are identical, $H = -\sum_{\langle ij \rangle} \left[c_i^\dagger c_j + c_j^\dagger c_i - n_i n_j \right]$, where the sum is over nearest neighbors of a two-dimensional square lattice. The upper right triangles of the first and second quantized distributions match, and the "missing" triangles, which can be reconstructed by expanding out the Slater determinants, will match as well. Positions are plotted as $R_i = x_i + L y_i$ for particle i at position (x_i, y_i).

Two important examples of the first case are the Hubbard model in an arbitrary basis and the DMC algorithm, where walkers accumulate a (diagonal) weight according to the electronic potential and move (off-diagonally) according to the one-body kinetic operator. To illustrate the relationship between first and second quantized walkers in such a case, we plot their distributions for a simple Hamiltonian – two particles hopping with nearest neighbor interactions – in Fig. 4.2. Notice that the first and second quantized walker distributions match step-by-step.

However, there are a number of Hamiltonians for which the first and second quantized sign problems are not identical. For instance, an arbitrary double excitation matrix element of $\mathbf{U}_F^+ + \mathbf{U}_F^-$ in a permanent basis is given by [68]

$$\frac{1}{N!} \sum_{\substack{|\mathbf{i}\rangle \in |\mathbf{P_i}\rangle \\ |\mathbf{j}\rangle \in |\mathbf{P_j}\rangle}} \left| \langle \mathbf{i}|U|\mathbf{j}\rangle \right| = \left| \langle ij|\hat{u}_{12}|ab\rangle \right| + \left| \langle ij|\hat{u}_{12}|ba\rangle \right| , \tag{4.9}$$

$M_{\mathbf{ij}}$	Lowest eigenvalue (a.u.)
$(\mathbf{U}_F)_{\mathbf{ij}} = \langle \mathbf{i}\|U\|\mathbf{j}\rangle$	0.436866981
$(\mathbf{U}_F^+ + \mathbf{U}_F^-)_{\mathbf{ij}} = -\|\langle \mathbf{i}\|U\|\mathbf{j}\rangle\|$	0.473717218
$(\mathbf{U}_D)_{\mathbf{D_i D_j}} = \langle \mathbf{D_i}\|U\|\mathbf{D_j}\rangle$	0.436866981
$(\mathbf{U}_D^+ + \mathbf{U}_D^-)_{\mathbf{D_i D_j}} = -\|\langle \mathbf{D_i}\|U\|\mathbf{D_j}\rangle\|$	0.465099719
$(\mathbf{U}_F^+ + \mathbf{U}_F^-)_{\mathbf{P_i P_j}}$ $= -\frac{1}{4!}\sum_{\substack{\|\mathbf{i}\rangle\in\|\mathbf{P_i}\rangle \\ \|\mathbf{j}\rangle\in\|\mathbf{P_j}\rangle}} \|\langle \mathbf{i}\|U\|\mathbf{j}\rangle\|$	0.473717218

Table 4.1: Largest eigenvalues of matrices related to the 3D uniform electron gas and consisting of elements $M_{\mathbf{ij}}$. Periodic boundary conditions were applied to a simulation cell of dimensions $L \times L \times L$ containing four electrons at a density of $r_s = 1$ a.u.. The plane-wave basis set of 38 plane waves contains all plane waves with momentum less than $2\pi/L$. The resulting determinant and permanent Hilbert spaces each contain 567 functions with $M_s = 0$ and momentum $\mathbf{k} = \mathbf{0}$. The time step and offset are $\tau = 0.1$ and $S = 0$ respectively.

where $\hat{u}_{12}$ is the two-body-interaction component of the U operator. The equivalent element in $\mathbf{U}_D^+ + \mathbf{U}_D^-$ is

$$\langle \mathbf{D_i}|U|\mathbf{D_j}\rangle = \Big|\langle ij|\hat{u}_{12}|ab\rangle - \langle ij|\hat{u}_{12}|ba\rangle\Big|. \tag{4.10}$$

It is clear that the two matrix elements are not, in general, equivalent. Hence the sign problem in FCIQMC can be different in a first and second quantized bases. As a concrete example where the sign problems differ, we consider the uniform electron gas (UEG, see Ref. 64). Results in Table 4.1 show that, even with just four electrons in a small plane-wave basis set, the sign problem in a determinant basis is inherently different from that in a first quantized basis.

In summary, we have shown that the sign problem in FCIQMC in a second quantized formalism is generally different from the the sign problem in a first quantized representation, except in certain cases such as the conventional instability of Diffusion Monte Carlo. For situations where the sign problem differs, we have shown in Appendix A that the second quantized algorithm is better. Thus, in addition to the technical improvements offered by second quantized FCIQMC (i.e., ease of annihi-

lation), it improves a more fundamental issue in fermion QMC: the sign problem. Finally, we note that, while all of our arguments and proofs have been focused on the case of fermions, they can be readily extended to the case where the particles are bosons to show that working in a second quantized basis (of permanents) is at least as good as a first quantized basis.

4.2 Improving FCIQMC as applied to the Fermi polaron

Having detailed the properties of FCIQMC and in particular its sign problem, we will now discuss a number of methods we have developed to improve the FCIQMC algorithm. It should be noted that other such techniques exist, most notably the initiator method of Ref. 66, but for the problems we have considered (i.e., the Fermi polaron) we have not found that method to be particularly successful. Thus, we have set out to develop extensions to improve the convergence of FCIQMC, utilizing the Fermi polaron (described below) as a test bed for these improvements. We discuss a number of improvements that we have made to the algorithm, including the implementation of importance sampling and release node, the development of a variant of fixed node with a tunable sign approximation, and an extension allowing FCIQMC to work in the thermodynamic limit.

4.2.1 Introduction to the Fermi polaron

The polaron is a classic problem in condensed matter physics, as it is one of the simplest problems that shows strongly-interacting many body effects. The most well-known version is the Fröhlich polaron [69, 70], involving a single electron interacting with a bath of phonons. We will be interested in the Fermi polaron as seen in cold atomic gases [1], defined instead for a three-dimensional system of two fermion

species, denoted spin up and spin down. Starting with a non-interacting Fermi sea of spin up particles, a single spin down "impurity" is added that interacts with the sea of spin ups. One wishes to solve for the ground state properties of this impurity, as dressed by the spin up interactions. Many possible ground states can result, including a well-defined polaronic quasiparticle with mass renormalized by the interaction, a tightly-bound bosonic molecule formed by binding a single spin-up particle, and other more exotic possibilities.

Since experimentally, near a broad s-wave Feshbach resonance, the fermions interact by a short-range potential with scattering length a, the Hamiltonian that we consider is [71]

$$H = \sum_{k\sigma} \epsilon_k c^\dagger_{k\sigma} c_{k\sigma} + \frac{g}{\mathcal{V}} \sum_{kpq} c^\dagger_{k+q,\uparrow} c^\dagger_{p-q,\downarrow} c_{p,\downarrow} c_{k,\uparrow} \ , \tag{4.11}$$

where $\epsilon_k = k^2/(2m)$ is the kinetic energy, $\mathcal{V}$ is the volume of the system, and m is the mass of the particles. We regularize the short-range interaction by introducing a high-momentum cutoff Λ on the spin-ups. The interaction strength g is cutoff-dependent [72]:

$$g^{-1} = \frac{1}{8\pi a} - \frac{\Lambda}{4\pi^2} \ . \tag{4.12}$$

We solve this Hamiltonian for a cube of side length $L = \mathcal{V}^{1/3}$ with periodic boundary conditions. With N spin up particles, the Fermi wavevector and energy are $k_F = (6\pi^2 N/\mathcal{V})^{1/3}$ and $E_F = k_F^2/(2m)$ respectively.

In the non-interacting limit ($a = g = 0$), the ground state of this Hamiltonian is an undressed polaron,

$$|D_0\rangle \equiv |\mathrm{FS}_\uparrow, \mathbf{0}_\downarrow\rangle \ , \tag{4.13}$$

where $|\mathrm{FS}_\uparrow, \mathbf{0}_\downarrow\rangle$ denotes a Fermi sea of spin-up particles with a single spin down at zero momentum. Upon turning on a weak attractive interaction (small negative a), the polaron remains a well-defined quasiparticle with non-zero quasiparticle residue $Z = |\langle\psi_0|D_0\rangle|^2$, where $|\psi_0\rangle$ is the ground state of the interacting Hamiltonian. Note

that Z can be written as the ground state expectation value $\langle\psi_0|\mathcal{P}_0|\psi_0\rangle$, where $\mathcal{P}_0 = |D_0\rangle\langle D_0|$ projects onto the undressed polaron.

Off of unitary on the BEC side ($a^{-1} > 0$), the polaron eventually becomes unstable to formation of a tightly-bound molecule. If we allow the masses of the spin up and spin down particles to differ, then an even more exotic panoply of ground states has been predicted, including a molecule at non-zero ground state momentum and a trimer, in which the molecule binds a second spin-up particle [72]. As the primary focus of this section is to investigate and improve the numerical method of FCIQMC and not to further map out the fascinating phase diagram of the Fermi polaron, from here out we restrict ourselves to the "mass-balanced" case: $m_\uparrow = m_\downarrow$.

The mass-balanced Fermi polaron has been most thoroughly investigated at unitarity, where $(k_F a)^{-1} = 0$. Chevy found that the ground state at unitarity is well-described by a variational wavefunction that includes all single particle-hole excitations on the spin up Fermi sea, giving a ground state energy of $-0.6066E_F$ [71, 73]. These results were later extended by Combescot and Mora [73], who introduced a procedure for calculating the variational ground state energy at an arbitrary number of particle-hole excitations and proceeded to solve for the two particle-hole pair variational energy, $-0.6156E_F$.

The Fermi polaron has also been approached via Monte Carlo techniques, most notably diagrammatic Monte Carlo [74]. Monte Carlo results for the ground state energy agree well with the variational solutions. Given the agreement of the ground state polaron energy between variational expansions, Monte Carlo, and experiment [1], the energetics of the polaron at unitarity are a good test case for improvements to the FCIQMC algorithm.

One remaining source of contention is the molecular binding transition, which is theoretically predicted to occur near $(k_F a)^{-1} = 0.9$ [72]. However, experimentally the quasiparticle residue vanishes at $(k_F a)^{-1} \approx 0.75$ [1]. It is believed that this difference

may be due to the small but finite density of the spin down particles in the experiment. Once we have developed the techniques to solve the Fermi polaron in this limit, we will briefly address this issue in Sec. 4.2.7.

4.2.2 Applying basic FCIQMC to the Fermi polaron

Given the effectiveness of the variational expansions in describing the polaron, we choose as our basis all momentum space determinants with at most M particle-hole pairs dressing the spin-up Fermi sea.[2] The Chevy ansatz [71] is then, for example, the exact solution for $M = 1$. In this basis, all off-diagonal elements of H come from the interaction and are given by the constant value $g/\mathcal{V}$, with their sign set by Fermi statistics. We fix the number of spin-up particles N and momentum cutoff Λ. Note that we eventually want to take the physical limit $N \to \infty$ and $\Lambda \to \infty$, for which the basis size becomes infinite. We work in the sector with zero total momentum, which is always valid for a polaron-like quasiparticle [72]. N will be limited to 7, 19, 27, ... to ensure that the spin up Fermi sea is a closed shell for our cubic volume with periodic boundary conditions.

The walkers are initialized in the undressed polaron state with positive sign, $|\psi_i\rangle = |D_0\rangle$. We write general determinants as

$$|D\rangle = |\mathbf{k}_1, \mathbf{k}_2, \ldots, \mathbf{k}_\ell; \mathbf{q}_1, \ldots, \mathbf{q}_\ell\rangle = c^\dagger_{\mathbf{k}_1} \cdots c^\dagger_{\mathbf{k}_\ell} c_{\mathbf{q}_1} \cdots c_{\mathbf{q}_\ell} |D_0\rangle \ , \tag{4.14}$$

where $\mathbf{k}$'s label the particle excitations, $\mathbf{q}$'s label the holes, $\mathbf{k}_1 < \mathbf{k}_2 < \ldots < \mathbf{k}_\ell$, and $\mathbf{q}_1 < \ldots < \mathbf{q}_\ell$ for an arbitrary momentum ordering. The values of $\mathbf{k}_i$ and $\mathbf{q}_i$ are restricted to momenta on a discrete grid set by the size of the box L. The energy

[2] As we work here in a basis with a restricted number of many-body excitations, the method is more accurately denoted configuration-interaction QMC (CIQMC), with FCIQMC being the special case with $M = N$; note that we used the notation "CIQMC" in Ref. 4. However, as the methods are algorithmically identical, we continue to use the notation "FCIQMC" for this dissertation, to avoid confusion.

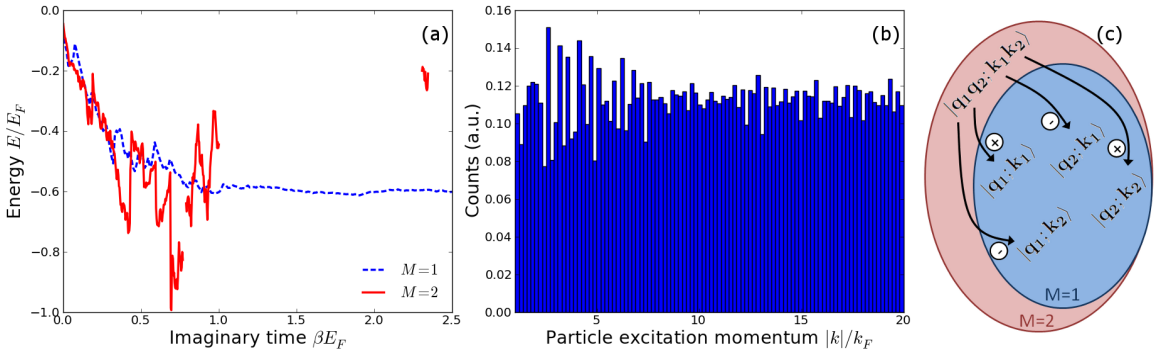

Figure 4.3: Issues with naive FCIQMC. (a) Polaron energy traces at unitarity for FCIQMC with $M = 1$ (dashed blue) and $M = 2$ (solid red) particle-hole excitations. Discontinuities occur for $M = 2$ when the denominator of the energy metric (Eq. 4.15) is zero. (b) Histogram of the particle excitation momentum for $M = 1$, computed stochastically. Momenta are uniformly occupied from the edge of the Fermi sea ($k = k_F$) up to the cutoff ($k = 20k_F$). (c) Graphical representation of sign issues for the polaron (see text for details). Panels (a-b) use $N = 33$ spin up particles and momentum cutoff $\Lambda/k_F = 20$.

metric in this basis is

$$E(\beta) = \frac{\langle D_0|H|\psi_w(\beta)\rangle}{\langle D_0|\psi_w(\beta)\rangle} = \frac{\sum_w S_w \langle D_0|H|D_w\rangle}{\sum_w S_w \langle D_0|D_w\rangle} , \tag{4.15}$$

where the zero of energy is defined such that $\langle D_0|H|D_0\rangle = 0$. For sufficiently large β, the statistical average of $E(\beta)$ converges to the fermionic ground state energy.

Using the original FCIQMC algorithm in this basis, we are able to find the ground state energy for $M = 1$, $N = 33$ and $\Lambda = 20k_F$, where the sign problem is weak (Fig. 4.3a). For $M = 2$, however, the sign problem prevents convergence with fixed number of walkers $N_w = 10^6$ (Fig. 4.3a). The $M = 2$ sign problem can be understood by starting from a determinant with two particle-hole pairs, $|D_2\rangle = |\mathbf{k}_1, \mathbf{k}_2; \mathbf{q}_1, \mathbf{q}_2\rangle$. Off-diagonal matrix elements allow hopping from $|D_2\rangle$ to four single particle-hole determinants (see Fig. 4.3c). Of these, two come with positive sign, two with negative sign, and all have the same weight $g/\mathcal{V}$. However, one can show that the signs within the single particle-hole shell should all be positive in the ground state, so

the determinants with two particle-hole pairs rain down nearly-random noise. This results in a large sign problem unless the noise is efficiently annihilated.

We emphasize that our approach to the Fermi polaron differs from previous problems where FCIQMC has been applied because almost all determinants are "important"; this fact arises because the interaction is short-range in real space, i.e., long-range in momentum space. Fig. 4.3b shows a histogram of the walkers in the first particle-hole shell for the $M = 1$ ground state. The particle excitations are evenly occupied from the Fermi sea all the way up to the cutoff, in contrast to situations as in Ref. 43 where the ground state consists of a few large-weight determinants; this makes annihilation more difficult.

FCIQMC for quantum chemistry problems minimizes these issues through an intelligent choice of single particle orbitals. However, for the Fermi polaron, there is no clear choice of single-particle basis that minimizes the impact of unimportant determinants. In the next section we show that using importance sampling can have a similar effect. This not only improves the variance, as usual in importance sampling, but also appears to improve the annihilation properties by favoring certain determinants.

4.2.3 Importance sampling

Importance sampling is by now a standard Monte Carlo technique that has been used to decrease statistical noise. Importance sampling consists of sampling walker $|D\rangle$ from a probability distribution proportional to $|\langle D|\psi_T\rangle\langle D|\psi_0\rangle|$ instead of $|\langle D|\psi_0\rangle|$, where $|\psi_T\rangle$ is a trial wavefunction that we choose. This is implemented by reweighting the off-diagonal moves, and can be thought of as simply acting with the (non-Hermitian) effective "Hamiltonian" H_{is} given by

$$\langle D'|H_{\text{is}}|D\rangle = \langle D'|H|D\rangle \frac{\langle \psi_T|D'\rangle}{\langle \psi_T|D\rangle} \ . \tag{4.16}$$

The energy metric becomes

$$E(\beta) = \frac{\sum_w S_w \langle D_0|H|D_w\rangle / \langle \psi_T|D_w\rangle}{\sum_w S_w \langle D_0|D_w\rangle / \langle \psi_T|D_w\rangle} \ . \tag{4.17}$$

We also calculate quasiparticle residue through the method of mixed estimators [75],

$$Z \approx 2\langle \psi_T|\mathcal{P}_0|\psi_0\rangle - \langle \psi_T|\mathcal{P}_0|\psi_T\rangle \ . \tag{4.18}$$

As H_{is} is not Hermitian, we must be careful to only act with it on kets.

Choosing an appropriate trial wavefunction plays a major role in the success of importance sampling. A naive guess, the free fermion wavefunction, would give $|\psi_T\rangle = |D_0\rangle$; this would not allow sampling of any particle-hole excitations. Therefore, we must construct some $|\psi_T\rangle$, which in general we would like to be as close as possible to the ground state $|\psi_0\rangle$.

To motivate our choice of $|\psi_T\rangle$, consider the Schrödinger equation $H|\psi_0\rangle = E_0|\psi_0\rangle$. If $H = T + V$, where the kinetic term T is diagonal in the momentum basis, then

$$\langle D|\psi_0\rangle = -\frac{1}{\langle D|T|D\rangle - E_0} \sum_{D'} \langle D|V|D'\rangle \langle D'|\psi_0\rangle \ . \tag{4.19}$$

For $|D\rangle$ with n particle-hole pairs, the interaction term connects it to determinants $|D'\rangle$ with $n-1$, n, and $n+1$ particle-hole pairs. The strength of the interaction is $g/\mathcal{V}$ with sign $\mathcal{S}_{D'} \equiv \text{sgn}(\langle D|V|D'\rangle)$. While these coupled linear equations are generally hard to solve, we can make the simplifying approximation of only considering $|D'\rangle$ with $n-1$ particle-hole pairs in constructing $|\psi_T\rangle$. To allow some freedom, we will define variational coefficients α and β to replace g and E_0 respectively, giving

$$\langle D_n|\psi_T\rangle = -\frac{\beta/\mathcal{V}}{\langle D_n|\hat{T}|D_n\rangle - \alpha} \sum_{\substack{D'_{n-1}\text{s.t.} \\ \langle D'_{n-1}|V|D_n\rangle \neq 0}} \mathcal{S}_{D'_{n-1}} \langle D'_{n-1}|\psi_T\rangle \ , \tag{4.20}$$

Λ	M	N	$1/(k_F a)$	α	β	Basis size
20	1	33	0	-0.6	20	1.09×10^6
20	2	33	0	-0.6	20	2.88×10^{11}
20	3	33	0	-0.6	20	3.28×10^{16}
20	4	33	0	-0.6	20	2.03×10^{21}
20	2	33	0.5	-1.22	29.6	2.88×10^{11}
20	2	33	0.9	-2.22	37.3	2.88×10^{11}
20	2	∞	0	-0.6	20	∞

Table 4.2: Basis size and variational parameters α and β used for different Hamiltonians defined by $[\Lambda, M, N, 1/(k_f a)]$, as described in the text.

where $|D_n\rangle$ denotes a determinant with n particle-hole pairs. This gives a recursive definition for $|\psi_T\rangle$ which, since normalization is irrelevant, we seed from $\langle D_0|\psi_T\rangle = 1$. The parameter α is chosen to be approximately the energy. Then, for each shell of n particle-hole pairs, we calculate $Z_n = \langle\psi_T|\mathcal{P}_n|\psi_T\rangle$ using variational Monte Carlo, where $\mathcal{P}_n$ projects onto the subspace of determinants with n excitations. β is chosen to minimize the maximum deviation of Z_n from its average value, $1/(M+1)$. Some characteristic values of α and β are shown in Table 4.2.

There is one minor issue with this ψ_T: due to the discreteness of the Hilbert space, Eq. 4.20 will potentially yield $\langle D|\psi_T\rangle = 0$ for some small set of determinants $|D\rangle$, which become increasingly scarce as $N \to \infty$. This prevents any weight from being put on these determinants, introducing a bias. To avoid this bias, we introduce a third parameter γ such that if Eq. 4.20 yields $|\langle D|\psi_T\rangle| < \gamma$, we set $\langle D|\psi_T\rangle \to \gamma$. The biased simulation with $\gamma = 0$ gives a variational upper bound on the energy, which should be fairly close to correct. As an example, at unitarity for the relatively small basis with $N = 7$, $\Lambda = 10k_F$, and $M = 3$, the difference in energy is $-0.6279(3)E_F$ for $\gamma = 0$ versus $-0.6310(17)E_F$ for $\gamma = 10^{-4}$. We use $\gamma = 10^{-4}$ throughout the remainder of this paper, unless otherwise specified.

Importance sampling greatly improves the effectiveness of FCIQMC. For example, at unitarity FCIQMC with importance sampling is able to solve the polaron ground

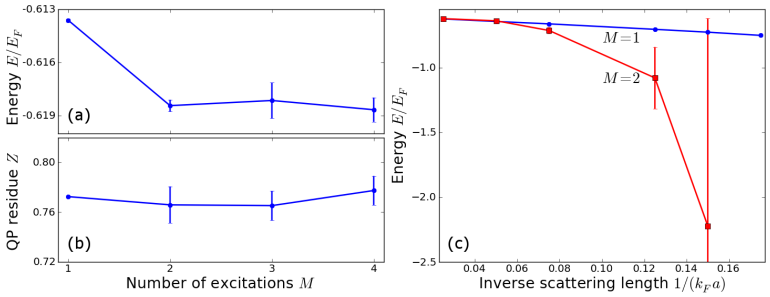

Figure 4.4: Limits of importance sampling. At unitarity with $N = 33$ and $\Lambda/k_F = 20$, importance sampled FCIQMC can solve the polaron energy (a) and quasiparticle residue (b), with $\gamma = 0$ to improve statistics. (c) Energies off of unitarity, with error bars computed for runs of fixed imaginary time $\Delta\beta = 35/E_F$. For $(k_F a)^{-1} > 0.15$ with $M = 2$, the simulation failed to converge.

state for $\Lambda = 20k_F$, $N = 33$, and $M = 4$ – corresponding to a basis size of 2.03×10^{21} – using only 10^6 walkers. Solutions for the polaron energy and quasiparticle residue at unitarity are shown in Fig. 4.4a-b.

We next push to positive values of a^{-1}, where for $(k_F a)^{-1} > 0.9$ the polaron is expected to become unstable to formation of a molecule. As the interaction strength $1/(k_F a)$ is increased, the weight of the wavefunction is pushed to higher momenta, making annihilation more difficult. As seen in Fig. 4.4c, for $M = 2$ particle-hole pairs, the error bar gradually increases as we push off of unitarity. Eventually, for $1/(k_F a) > 0.15$, FCIQMC with importance sampling fails to find the ground state. To proceed, we now introduce modifications that allow us to further stabilize the algorithm.

4.2.4 Partial node approximation

Fixed node quantum Monte Carlo is a method that removes the exponential cost of solving a fermionic Hamiltonian by approximating the Hamiltonian with a related one that has no sign problem. In this section we develop a way to extrapolate to the exact solution of our (restricted basis) Hamiltonian using a fixed-node-inspired

starting point. In fixed node algorithms, such as the lattice formulation of van Bemmel et. al. [76], one must specify a "correct" sign for each determinant. For simplicity we choose to use $|\psi_T\rangle$ as determining our sign structure. Given this choice, an off-diagonal element $\langle D'|H_{\rm is}|D\rangle$ will be called sign-violating if $\langle D'|H_{\rm is}|D\rangle > 0$, since in this case applying $-\tau H_{\rm is}$ to hop from D to D' would flip the sign, violating the sign structure of $|\psi_T\rangle$.

The fixed node Hamiltonian $H_{\rm fn}$ is given by

$$\begin{aligned}
\langle D'|H_{\rm fn}|D\rangle &= \begin{cases} 0 & \text{if sign viol. (s.v.)} \\ \langle D'|H_{\rm is}|D\rangle & \text{if not s.v. (n.s.v.)} \end{cases} \text{, where } D \neq D' \\
\langle D|H_{\rm fn}|D\rangle &= \langle D|H_{\rm is}|D\rangle + \sum_{D' \text{S.V.}} \langle D'|H_{\rm is}|D\rangle \ . \qquad (4.21)
\end{aligned}$$

Note that the sign-violating off-diagonal matrix elements have been removed and "dumped" onto the diagonal. $H_{\rm fn}$ is guaranteed to give a variational upper bound on the ground state energy, yielding the correct ground state energy if $|\psi_T\rangle = |\psi_0\rangle$.

Lattice fixed node was originally designed for real space lattices where there are very few off-diagonal elements connected to any given configuration. For our Hamiltonian, there will in general be many sign-violating matrix elements connected to a given determinant $|D\rangle$ (of order $> 10^7$ for our parameters), and the sums required to modify the diagonals of $H_{\rm fn}$ become analytically intractable. We therefore initially work with a modified Hamiltonian $H'_{\rm fn}$ – the fixed node Hamiltonian without diagonal dumping – given by

$$\begin{aligned}
\langle D'|H'_{\rm fn}|D\rangle &= \begin{cases} 0 & \text{if sign viol.} \\ \langle D'|H_{\rm is}|D\rangle & \text{if not s.v.} \end{cases} \text{, where } D \neq D' \\
\langle D|H'_{\rm fn}|D\rangle &= \langle D|H_{\rm is}|D\rangle \ , \qquad (4.22)
\end{aligned}$$

whose ground state energy is no longer guaranteed to be a variational upper bound on that of H. In Sec. 4.2.6 we will describe the steps needed to reintroduce the variational upper bound.

Due to the presence of annihilation in FCIQMC, we are able to introduce a less-severe approximation, which we call the partial node approximation. In partial node, we interpolate between the importance sampled and fixed node Hamiltonians

$$H_{\text{pn}} = \epsilon_{\text{pn}} H'_{\text{fn}} + (1 - \epsilon_{\text{pn}}) H_{\text{is}} \ . \tag{4.23}$$

Surprisingly, we find that this algorithm works up to a very high fraction of the original sign violating terms (small ϵ_{pn}). To obtain the exact answer for $M = 2$, we extrapolate to the original Hamiltonian ($\epsilon_{\text{pn}} = 0$), heuristically using a quadratic fit (Fig. 4.5b).

To quantify the improvement gained by using the partial node approximation, we measure the average flip rate of the sign of the walkers on the reference determinant D_0 (which typically should contain approximately 30% of the walkers for a given snapshot). The flipping of this sign is a signature of an instability of the algorithm coming from the sign problem, and the average flip rate must be significantly below the inverse correlation time for useful data to be garnered from the simulation. Fig. 3d shows the average flip rate using the partial node approximation. Notice that the flip rate for $\epsilon_{\text{pn}} = 0$ is close to the inverse correlation time, presenting a problem, but by $\epsilon_{\text{pn}} = 0.4$ the flip rate has significantly decreased.

4.2.5 Release node FCIQMC

In FCIQMC without annihilation, the sign problem manifests as a variance that scales exponentially with both β and the energy difference between the fermionic and bosonic ground states (see Sec. 3.1.3). This exponential scaling in β holds even in the

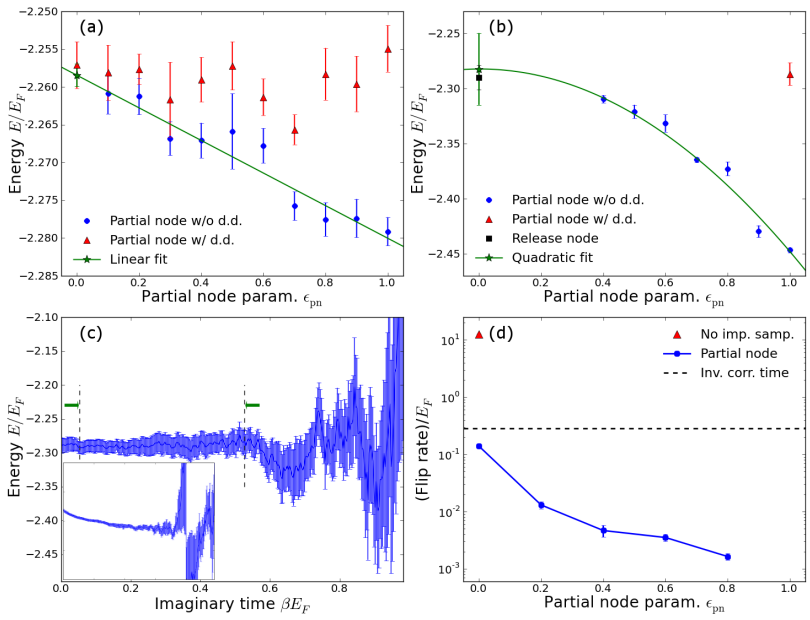

Figure 4.5: Partial and release node FCIQMC for $1/(k_F a) = 0.9$, $\Lambda = 20k_F$, and $N = 33$, both with and without diagonal dumping (d.d.). These two techniques give consistent ground state energies for $M = 1$ (a) and $M = 2$ (b). Note that at $\epsilon_{\text{pn}} = 0$, the partial node method reduces to importance sampling. In (b), the release node is initialized from the $\epsilon_{\text{pn}} = 0.4$ partial node walkers, the trace of which is shown in (c), including standard error across 100 trials. A less optimal initial condition is shown in the inset. (d) Flip rate of the denominator of the energy metric as a function of ϵ_{pn} (blue dots), compared to the flip rate without importance sampling (red triangle) and the inverse of the Monte Carlo correlation time (black dashed line). Panels (a-c) use 10^6 walkers, while (d) uses 10^4.

presence of annihilation at finite walker number, so minimizing the total β needed to reach the ground state wave-function with a certain fidelity also minimizes the effect of the sign problem. Thus, one can improve the convergence time of the importance sampled FCIQMC algorithm ($\epsilon_{\text{pn}} = 0$) by starting very close to the ground state; this is called release node.

One approach to starting near the true ground state is to use the ground state of the partial node Hamiltonian H_{pn} with a small value of ϵ_{pn}. We aren't able to explicitly represent this wavefunction, but we can prepare it stochastically via partial node FCIQMC. The set of N_w walkers prepared thusly is allowed to evolve under the true Hamiltonian for fixed imaginary time β. When the wavefunction is "released"

to evolve under the exact Hamiltonian, it quickly relaxes to the exact ground state; this is what we call release node FCIQMC [77].

In release node, we use the same energy metric as in importance sampled FCIQMC. In this case, since we have now propagated with two non-commuting Hamiltonians (H_{pn} and H_{is}), the arguments in Eq. 4.2 do not work, so at early times after "release" the energy metric has no meaning. The energy metric becomes meaningful again when the wavefunction reaches the ground state, at which point it gives the ground state energy. As seen in Fig. 4.5c, the energy converges to the ground state energy before the errors blow up if we start close enough to the ground state. The release node energies agree well with the partial node extrapolation (Fig. 4.5a-b), which provides a check of the validity of our extrapolation.

4.2.6 Diagonal dumping in partial node FCIQMC

The partial node approximation that we have discussed thus far does not give a variational upper bound on the energy, since we removed the diagonal dumping to get the effective Hamiltonian H'_{fn} (Eq. 4.22). To restore the variational upper bound, we now discuss how one can put the effects of the diagonal dumping back into the partial node algorithm.

While the sum involved in diagonal dumping is in general not analytically tractable, it can be done stochastically in much the same way as off-diagonal spawning. During the diagonal create/kill step, we would like to stochastically apply

$$U_{\text{diag}} = 1 - \tau\langle D|H_{\text{is}}|D\rangle - \tau \sum_{D'\text{S.V.}} \langle D'|H_{\text{is}}|D\rangle \tag{4.24}$$

to a walker on determinant $|D\rangle$. The first two terms can be done exactly, while the sum – denoted ΔK – can be sampled as follows.

1. Pick another determinant $|D'\rangle$ according to some normalized probability function $q(D'|D)$.

2. If $\langle D'|H_{\text{is}}|D\rangle$ is not sign-violating, then let $\Delta K = 0$.

3. Otherwise, let $\Delta K = \langle D'|H_{\text{is}}|D\rangle / q(D'|D)$.

However, as this stochastically sampled ΔK is now potentially very large, multiplying $|D\rangle$ by weight U_{diag} has the potential to be disastrous if $U_{\text{diag}} \ll 0$. Therefore, we find it necessary when stochastically dumping the diagonal to also use the approximation

$$1 - \tau\langle D|H_{\text{is}}|D\rangle - \tau\Delta K \approx e^{-\tau(\langle D|H_{\text{is}}|D\rangle - \Delta K)} \ . \tag{4.25}$$

This introduces a time step error which one must extrapolate to zero; empirically a quadratic extrapolation for $\tau < 10^{-3}$ appears to be quite effective. Partial node results with stochastic dumping of the diagonal are shown in Fig. 4.5a-b.

4.2.7 Results of partial and release node FCIQMC

Using partial node FCIQMC followed by release node, we are now able to determine values for the energy and quasiparticle residue of the Fermi-polaron on the BEC side of the interaction, which are shown in Fig. 4.6a-b for $\Lambda = 20k_F$ and $N = 33$. The energies are compared to polaron energies from diagrammatic Monte Carlo [74]. Our results match well for small values of $1/(k_F a)$, while for larger values of $1/(k_F a)$ our energies deviate from the diagrammatic results. This is a consequence of working at fixed $\Lambda = 20k_F$, as opposed to diagrammatic Monte Carlo, which is done in the limit $\Lambda \to \infty$. Our results agree with the diagrammatic solution after extrapolating to $\Lambda = \infty$.

Fig. 4.6b shows the quasiparticle residue Z compared to experimental results [1], where there is a small but finite density of spin-down atoms. We find that, as with the

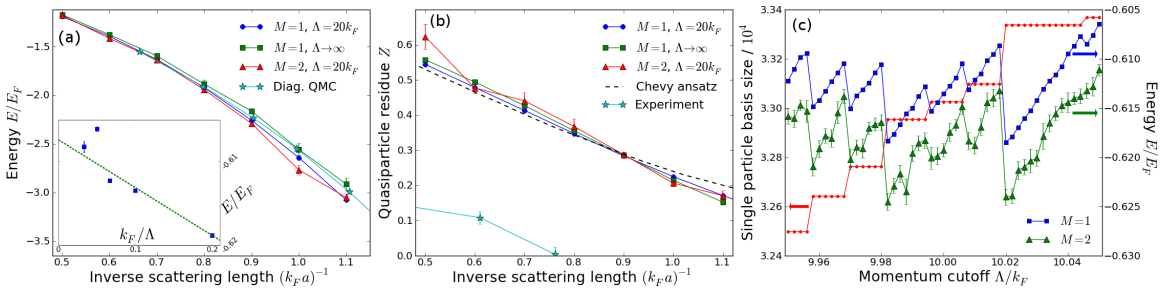

Figure 4.6: Ground state energy (a) and quasiparticle residue (b) at $\Lambda = 20k_F$, $N = 33$. Energies are compared against the $N \to \infty$, $\Lambda \to \infty$ results from diagrammatic QMC [74], while Z is compared to experimental measurements [1] and Chevy's analytical solution for $M = 1$ [71]. Differences in Z between FCIQMC and analytic results come from finite size effects and our use of a mixed estimator. The inset to (a) shows finite size effects in extrapolating to the TDL at unitarity. In (c), we show a correlation between jumps in these energies and single particle basis size.

energy, an $M = 1$ variational expansion provides a good estimate of Z. Therefore, as noted elsewhere [1], our theoretical model differs significantly from the experiment, possibly as a result of the finite spin-down density used experimentally.

In extrapolating our results to the physical limit, $\Lambda \to \infty$, we encountered an unexpected problem. Using a linear fit over a wide range of $1/\Lambda$ (Fig. 4.6a inset), we find that individual data points have error well outside the line, with no discernible pattern. Zooming into a very small region of Λ (Fig. 4.6c), we discover the reason behind this: small fractional increases in basis size as new shells become available cause large jumps in the ground state energy. A similar issue occurs when attempting to extrapolate in particle number N. We nevertheless attempt such an extrapolation in Fig. 4.6a-b. As a result, the size of the error bars reflects not the accuracy of the data points at individual values of Λ, but rather these inherent shell effects.

Shell effects vanish as $N \to \infty$. Therefore, we next address the possibility of extending FCIQMC into the thermodynamic limit.

4.2.8 Extension to the thermodynamic limit

All the quantum Monte Carlo simulations described so far have been done for a finite number of particles. In this section, we describe how to modify FCIQMC to work directly in the thermodynamic limit (TDL, $N \to \infty$). One complication is that, in the TDL, the momenta form a continuous set of k-points instead of being limited to a discrete grid. Additionally, it is important to work in a representation where instead of enumerating the momenta for all $N = \infty$ particles, we instead store only their excitations above a known state, in this case $|D_0\rangle$. Crucially, in order for the spectrum of the Hamiltonian to remain bounded, allowing application of $1 - \tau H$ instead of $e^{-\tau H}$, the TDL only works with a cutoff M on the number of excitations allowed. This requirement is similar in spirit to the diagrammatic Monte Carlo method of imposing a cutoff on diagram order, in an attempt to avoid a divergence.

As momentum space is no longer discretized in the TDL, one might assume that annihilation is no longer possible, and therefore that the algorithm is doomed to fail. However, there is one key exception: due to the discrete choice of $|D_0\rangle$, annihilation can still occur at that one determinant.

Given this, the TDL algorithm is in practice nearly identical to the algorithm with finite N; in fact, one can think of the TDL algorithm as just the limit of the finite algorithm for larger and larger N. We have carefully defined $|\psi_T\rangle$ such that all factors of $\mathcal{V}$ cancel out when determining relevant quantities, such as spawning probability or energy. This is important because taking the limit $N \to \infty$ while at fixed spin-up density (constant k_F) means taking $\mathcal{V} \to \infty$ as well.

The TDL algorithm is remarkably effective at finding the ground state with $M = 1$ at unitarity, where the sign problem is weak. These energies at various values of Λ are shown in Fig. 4.7a, where an extrapolation to $\Lambda = \infty$ is much smoother than for finite N due to the absence of shell effects.

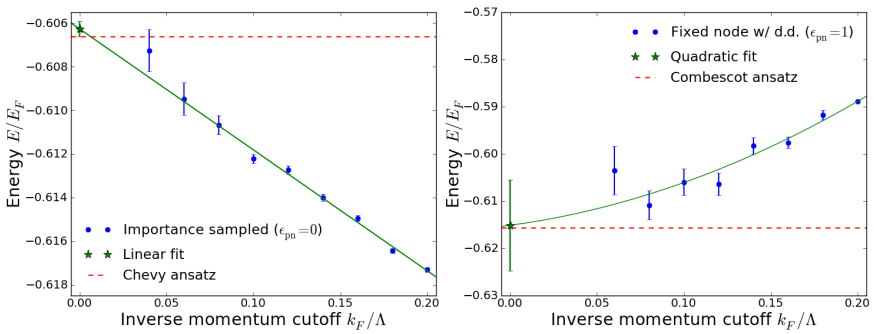

Figure 4.7: (Left) Ground state polaron energy in the TDL for $M = 1$, $(k_F a)^{-1} = 0$ as a function of Λ (blue dots). A linear extrapolation in $1/\Lambda$ (green star) agrees with the exact $M = 1, \Lambda \to \infty$ energy as calculated by Chevy [71, 73]. (Right) Fixed node ($\epsilon_{\rm pn} = 1$) energies for $M = 2$ compared against the exact solution by Combescot and Giraud [73].

For the sign-problem-heavy $M = 2$ case, we instead show data in the sign-free fixed node limit ($\epsilon_{\rm pn} = 1$), using the diagonal dumping method described in Sec. 4.2.6. This gives a good energy for the $M = 2$ polaron, which is a variational upper bound on the known energy of $-0.6156E_F$ within error bars. We believe this latter approach of combining fixed node with working directly in the thermodynamic limit will have wide applicability even for the approximate calculations that currently dominate the fermion QMC literature.

4.3 Conclusions

In this chapter, we have discussed many aspects of the FCIQMC algorithm, noting both its limitations with respect to the sign problem and ways in which it improves on certain other techniques. In particular, we showed in the first half of this chapter that the predominate use of FCIQMC as a second quantized algorithm for finding fermionic ground state has a sign problem that is at least as good as equivalent methods that work in a first quantized formalism. We showed that when they differ, the sign problem in second quantization is strictly better, and explicated the conditions under

which this occurs. An open question is how much of an sign problem improvement second quantization gives for systems of physical interest, and how this compares to other techniques to improve the sign problem, such as an intelligent choice of single particle basis and/or the developments we described in the second half of this chapter.

Second, we have shown that the FCIQMC algorithm can solve the Fermi polaron problem, after improving the algorithm with an appropriate importance sampled wave function, introducing partial and fixed node approximations, and utilizing release node methods. We believe these improvements demonstrate two main points that should be useful as FCIQMC and its variants are applied to future problems in condensed matter physics. First, for strongly correlated condensed systems where many determinants are occupied in any single particle basis, we have demonstrated that physical insight – in the form of a good choice of trial wavefunction – can significantly improve the behavior of the FCIQMC algorithm. Combining the improved statistics of importance sampling with sign-attenuating approximations such as partial node FCIQMC, we showed a significant increase in the effectiveness of FCIQMC in solving the polaron problem. Finally, we showed that these methods can be made exact, up to statistical noise, via extrapolation or release node QMC.

Second, we have shown that under certain conditions the FCIQMC algorithm can be extended to work with systems in the thermodynamic limit. Furthermore, we have introduced a fixed node algorithm in this limit, which has historically been an important tool for solving fermionic systems with QMC. We anticipate that these new developments will open a variety of problems in condensed matter physics to be approached using these new developments.

Chapter 5

Kibble-Zurek scaling of the quantum Ising chain

The previous chapter addressed the properties of the Fermi polaron ground state, assuming the atoms could be cooled to a sufficiently low temperature for such properties to be visible. The remainder of this thesis will primarily focus on the properties of isolated cold atom and model systems out of thermal equilibrium. As an initial foray into this field, in this chapter we discuss dynamics of one of the simplest but most useful models in condensed matter theory: the one-dimensional transverse-field Ising chain.

The transverse-field Ising (TFI) chain is one of the simplest models that exhibits a quantum phase transition [78]. The theory of scaling in the vicinity of a second-order phase transition is one of the most successful paradigms of modern condensed matter theory. Scaling theory was originally developed for omnipresent thermodynamic phase transitions, where thermal fluctuations drive the transition from an ordered to a disordered phase. However, nearly identical concepts apply to zero-temperature quantum phase transitions, where the order-to-disorder transition is driven by quantum fluctuations, i.e., non-commutativity of the quantum operators.

The ideas of critical scaling theory are codified by the renormalization group, in which one "coarse grains" the system, essentially zooming out physically by integrating out the highest energy (shortest distance) modes in the system. One then solves for how various terms in the action "flow" as coarse graining proceeds, in order to identify the properties of physics at the longest length scales (lowest energy scales). The flows are characterized by certain numbers known as critical exponents, which describe the asymptotic behavior of various observables as one approaches the critical point. These critical exponents and associated scaling functions are universal properties of the critical point, meaning that two physically distinct models that exhibit the same symmetry breaking will have the same critical exponents.

Critical exponents and scaling functions characterizing the zero-temperature (ground state) quantum critical point in the TFI chain have been understood for more than sixty years, since the seminal work of Onsager [79]. However, with the recent advent of isolated quantum dynamics as found in cold atom systems, a question has arisen whether such a critical point exhibits universal scaling of non-equilibrium dynamics, particularly within the framework of the Kibble-Zurek (KZ) mechanism, which describes the dynamics of a system as it is ramped across a phase transition at a finite rate. Kibble first introduced this idea to model symmetry breaking during cooling of the early universe [80], and later Zurek showed that condensed matter systems exhibit similar behavior in the context of slowly ramping across the superfluid phase transition in liquid ^{4}He [81]. Both proposals looked at ramps from deep in the disordered phase to deep in the ordered phase, which has historically been the primary focus of research on the KZ mechanism (cf. [82, 83]). Recent work has also explored ramps that end at the critical point [5, 84, 85].

However, the scaling theory of the KZ mechanism applies more generally than these two limits. Near the quantum critical point (QCP), observables measured during the ramp are postulated to exhibit universal scaling forms [86, 87]. In this chapter, we

solve for the scaling functions of excess heat, ground state log-fidelity, and equal-time spin correlations in the TFI chain for two scaling directions – time t during the ramp and finite system size L – and investigate the qualitative features of these observables.

In many cases, the quantum dynamics of a closed non-integrable quantum system lead to a thermal state. In contrast, we show here that the spin correlation functions are qualitatively different from those of any thermal state. We also find a ramp protocol where the state at large time not only does not thermalize, but also does not dephase to the diagonal ("generalized Gibbs") ensemble. Finally, we provide evidence for the universality of the dynamics by using time-dependent matrix product states (tMPS) to simulate the experimentally-realizable, non-integrable model of Mott insulating bosons in a tilted potential [2, 88].

5.1 Non-equilibrium dynamics of the transverse-field Ising chain

The Hamiltonian for the TFI chain on an L-site one-dimensional lattice is

$$H = -J \sum_{j=1}^{L} \left[(1-\lambda)\, s_j^x + s_j^z s_{j+1}^z \right], \tag{5.1}$$

where λ is a tunable transverse magnetic field and $s^{x,z}$ are Pauli matrices with periodic boundary conditions ($s_1 = s_{L+1}$). We work in units where $J = 1/2$ and, for technical simplicity, restrict ourselves to L even.

This model has two quantum phase transitions (QPTs) at $\lambda_c = 0$ and 2 [78]. The ground state of the TFI chain is given by a disordered paramagnetic (PM) phase for $\lambda < 0$ and $\lambda > 2$, while for $0 < \lambda < 2$, the system is in a ferromagnetic (FM) phase with two degenerate ordered ground states in the limit $L \to \infty$. For the remainder of this chapter, we will focus on the QCP at $\lambda_c = 0$.

The Hamiltonian of the TFI chain can be diagonalized by applying a Jordan-Wigner transformation [78, 89]:

$$s_j^x = 1 - 2c_j^\dagger c_j, \quad s_j^z = -(c_j + c_j^\dagger) \prod_{m<j} s_m^x, \tag{5.2}$$

where $c_j^\dagger$ creates an effective fermion at site j. This Hamiltonian conserves the fermion number modulo 2, and for an even length chain the ground state has an even number of fermions [89]; thus, we restrict ourselves to that sector. Within the even fermion parity sector, the "string" operator ($\prod_{m<j} s_m^x$) takes the periodic boundary conditions on the spins and translates them into anti-periodic boundary conditions for the fermions. So, Fourier transforming,

$$H = \sum_k \left[(1 - \lambda - \cos k) c_k^\dagger c_k + \frac{\sin k}{2} (c_k^\dagger c_{-k}^\dagger + c_{-k} c_k) \right], \tag{5.3}$$

where the available momenta are $k_n = \frac{2\pi(n+1/2)}{L}$ for $n = 0, 1, \ldots, L-1$.

This Hamiltonian can only excite fermions in momentum pairs $(k, -k)$, so we introduce a pseudo-spin σ corresponding to whether the $(k, -k)$ pair is filled ($\sigma^z = 1$) or unfilled ($\sigma^z = -1$). Then in the sector where all fermions are either occupied or unoccupied in $(k, -k)$ pairs (which contains the ground state), the Hamiltonian becomes $H = \sum_{n=0}^{L/2-1} H_{k_n}$, where

$$H_k = (1 - \lambda - \cos k)\sigma^z + (\sin k)\sigma^x \ . \tag{5.4}$$

Note that H is a free-fermion Hamiltonian, so the TFI chain is integrable.

As a prelude to understanding the dynamic scaling of this model, let's examine the basic scaling properties of its ground state. As mentioned earlier, renormalization is a procedure to understand the long-wavelength properties of the system near the critical point. These limits are particularly straightforward to see in the TFI model,

by just taking the $|\lambda|, k \ll 1$ limit of Eqs. 5.3 and 5.4:

$$H \approx \sum_k \left[\lambda \sigma_k^z + k \sigma_k^x\right] . \tag{5.5}$$

As we will see later, this limit is taken naturally in the case of KZ dynamics.

Note that this Hamiltonian is a function of λ as well as the set of momenta $\{k\} = \frac{2\pi(n+1/2)}{L}$, i.e., $H = H(\lambda, \{k\})$. Consider the Hamiltonian closer to the critical point at $\lambda' = \lambda/2$. Comparing the Hamiltonians at λ and λ', we see that they satisfy the relation $H(\lambda, \{k\}) = 2H(\lambda', \{k/2\})$. In other words, the low-energy physics at λ and λ' is equivalent as long as we rescale the energy and momenta by a factor of 2. This is the essence of scaling theory.

We can associate critical exponents with the rescaling properties of energy and momentum (or equivalently time and distance) in the vicinity of the critical point. For instance, the correlation length critical exponent ν and the dynamic critical exponent z are defined such that, under renormalization, characteristic length (ξ) and time (ξ_t) scales diverge as

$$\xi \sim |\lambda|^{-\nu};\ \xi_t \sim |\lambda|^{-\nu z} . \tag{5.6}$$

From the above discussion, it is clear that $\nu = z = 1$ for the TFI model. Similarly, we can ask about the scaling properties of observables, such as spin-spin correlations, near the critical point. In this case, critical exponents are replaced by entire scaling functions, describing the low-energy long-distance behavior of the spin correlations in terms of scaling-dimensionless quantities such as L/ξ. For instance, the scaling form of unequal-time spin-spin correlations in an infinite length TFI chain is

$$G(x,t) = \langle \psi_0(\lambda) | S_j^z(0) S_{j+x}^z(t) | \psi_0(\lambda) \rangle = \lambda^{2\beta} g\left(\frac{x}{\xi} = x\lambda, \frac{t}{\xi_t} = t\lambda\right) , \tag{5.7}$$

where $|\psi_0(\lambda)\rangle$ is the ground state of the Hamiltonian in Eq. 5.5. We have introduced a new critical exponent β, corresponding to the scaling properties of the spin S^z. This "anomalous" critical exponent was found to be $\beta = 1/8$ by Onsager and Yang [79, 90], although its derivation is not as simple as that of ν and z. When considering the low-energy long-distance behavior of other models in the Ising universality class, not only are the critical exponents β, ν, and z universal, but so is the entire scaling function $g(x/\xi, t/\xi_t)$.

We are now in a position to discuss the non-equilibrium scaling properties of the TFI chain, in particular its scaling within the KZ criterion.

5.1.1 Kibble-Zurek scaling

In a KZ ramp, λ is varied as a function of time near the critical point at $\lambda = 0$. For simplicity, we focus primarily on the case of a linear ramp, $\lambda(t) = vt$, where v is the ramp rate and the QCP is at $\lambda = t = 0$. Importantly, we always work at times t such that $\lambda < 2$, so that the ramp only crosses a single critical point. The ramp begins at $t = -\infty$ deep in the disordered phase, where the wavefunction is initialized in the ground state of the instantaneous Hamiltonian $H(t = -\infty)$. For an infinite system, the Hamiltonian is gapless at the critical point. Therefore, it is impossible to ramp slowly enough through the QCP to avoid creating excitations and to produce true long range order in the ordered phase.

Near the QCP, the characteristic time scale $\xi_t \sim \lambda^{-\nu z}$ becomes arbitrarily large. Thus for a non-zero ramp rate, there exists a Kibble time, $-t_K$, at which the lowest momentum modes stop following the ramp adiabatically and become significantly excited. More explicitly,

$$\begin{aligned} |t_K| &= \xi_t(\lambda(|t_K|)) \sim |vt_K|^{-\nu z} \\ \Rightarrow t_K &\sim v^{-\nu z/(1+\nu z)} \ . \end{aligned} \tag{5.8}$$

For the TFI chain $\nu = z = 1$, so $t_K \sim v^{-1/2}$. One can similarly define a length $\ell_K \sim t_K^{1/z}$ such that t_K and ℓ_K set the characteristic time and length scales for the Kibble-Zurek scaling forms [86, 87].

For a KZ ramp of the TFI chain, the wavefunction $|\psi(t)\rangle$ can be written as a product $|\psi(t)\rangle = \otimes_k |\Psi_k(t)\rangle$, where each mode evolves independently as

$$\begin{aligned} id|\Psi_k\rangle/dt &= H_k(t)|\Psi_k\rangle \\ &= \left[(1 - vt - \cos k)\sigma_k^z + (\sin k)\sigma_k^x\right]|\Psi_k\rangle \ . \end{aligned} \tag{5.9}$$

The KZ scaling limit is defined as taking $v \to 0$ with τ, Λ, and κ constant [86, 87], where

$$\begin{aligned} \tau &\equiv tv^{1/2} \sim t/t_K \\ \Lambda &\equiv Lv^{1/2} \sim L/\ell_K \\ \kappa &\equiv kv^{-1/2} \sim k\ell_K \\ &= \frac{2\pi(n+1/2)}{\Lambda}. \end{aligned} \tag{5.10}$$

$n = 0, 1, ..., (L/2) - 1$ indexes the modes. Note that $L \to \infty$ in the scaling limit, as $\ell_K \to \infty$ and Λ is constant.

To take the KZ scaling limit, we can expand out the sine and cosine in Eq. 5.9 and multiply through by a factor of $v^{-1/2}$ to get

$$i\frac{d|\Psi_\kappa\rangle}{d\tau} = \left\{ \Big[-\tau - \frac{\kappa^2}{2}v^{1/2} + O(v)\Big]\sigma_\kappa^z + \Big[\kappa - \frac{\kappa^3}{6}v + O(v^2)\Big]\sigma_\kappa^x \right\} |\Psi_\kappa\rangle \ . \tag{5.11}$$

Notice that all quadratic and higher order terms in κ from the sine and cosine come with positive powers of v. Thus, in the KZ scaling limit ($v \to 0$), Eq. 5.11 simplifies to

$$id|\Psi_\kappa\rangle/d\tau = \left(-\tau\sigma_\kappa^z + \kappa\sigma_\kappa^x\right)|\Psi_\kappa\rangle \tag{5.12}$$

Each mode $\Psi_\kappa(\tau)$ in Eq. 5.12 has the form of a Landau-Zener (LZ) equation, which can be solved in terms of parabolic cylinder functions [24, 91]. Note that $|\Psi_\kappa(\tau)\rangle$ is expressed solely in terms of the scaled variables τ and κ; the existence of these scaled wavefunctions depending solely on scaled quantities constitutes the essence of the KZ scaling theory.

In order to visualize the scaled wavefunction from Eq. 5.12, we solve for the dynamic scaling functions of three observables of the TFI chain: the ground state log-fidelity F [92], the excess heat Q, and the spin-spin correlation function G:

$$F(t) = -\log\left(\left|\langle\psi^0(t)|\psi(t)\rangle\right|^2\right) \tag{5.13}$$

$$Q(t) = \langle\psi(t)|H(t)|\psi(t)\rangle - \langle\psi^0(t)|H(t)|\psi^0(t)\rangle \tag{5.14}$$

$$G(x,t) = \langle\psi(t)|s_j^z s_{j+x}^z|\psi(t)\rangle\ , \tag{5.15}$$

where $|\psi^0(t)\rangle$ is the instantaneous ground state of $H(t)$. Both log-fidelity and excess heat serve as good initial measures of non-adiabaticity, going to zero in limit of an adiabatic ramp in which the system remains in its ground state. Fidelity also holds a special place in the realm of quantum computation and simulation, particularly in adiabatic quantum computation, where it is proposed that many-body systems be ramped from a boring (unentangled) state to an interesting (entangled) state for use in computing. Such proposal often involve ramps across a quantum critical point [93, 94], and small ground state log-fidelity serves as herald for the ramp's success. Meanwhile, spin correlations provide additional insight into the properties of the wavefunction, which are of particular relevance both experimentally and in the context of scaling theory.

For the non-equilibrium ramp, excess heat and log-fidelity become

$$Q(t) = \sum_{\kappa} \left[\langle \Psi_\kappa(t) | H_\kappa(t) | \Psi_\kappa(t) \rangle - \langle \Psi^0_\kappa(t) | H_\kappa(t) | \Psi^0_\kappa(t) \rangle \right] \tag{5.16}$$

$$F(t) = -\sum_{\kappa} \log \left(\left| \langle \Psi^0_\kappa(t) | \Psi_\kappa(t) \rangle \right|^2 \right) \ . \tag{5.17}$$

where $|\Psi^0_\kappa(t)\rangle$ is the instantaneous ground state of $H_\kappa(t)$. The excess heat has the form of a sum over the excess heat in each mode, Q_κ, which is just the product of the excited state occupancy

$$p^{\rm exc}_\kappa = 1 - \left| \langle \Psi^0_\kappa | \Psi_\kappa \rangle \right|^2 \tag{5.18}$$

and the excited state energy

$$E^{\rm exc}_\kappa = 2\sqrt{t^2 + k^2} = 2v^{1/2}\sqrt{\tau^2 + \kappa^2} \ . \tag{5.19}$$

If we define a scaled form of this mode excess heat as

$$q^{\rm exc}_\kappa(\tau) \equiv \frac{Q_\kappa(\tau)}{v^{1/2}} = 2p^{\rm exc}_\kappa(\tau)\sqrt{\tau^2 + \kappa^2} \ , \tag{5.20}$$

then the scaled excess heat density is

$$q(\tau, \Lambda) \equiv \frac{Q(t)}{vL} = \frac{1}{\Lambda} \sum_{\kappa} q^{\rm exc}_\kappa(\tau) \ . \tag{5.21}$$

Similarly, the scaled log-fidelity density[1] is

$$\begin{aligned} f(\tau,\Lambda) &\equiv \frac{F(t)}{v^{1/2}L} \\ &= -\frac{1}{\Lambda}\sum_{\kappa}\log\left(1-p_{\kappa}^{\rm exc}(\tau,\Lambda)\right) \ . \end{aligned} \tag{5.22}$$

These scaling forms are obtained by making the observable scaling-dimensionless using factors of t_K and ℓ_K. So, for example, the excess heat density Q/L becomes

$$\frac{Qt_K}{(L/\ell_K)} = \frac{Qv^{-1/2}}{Lv^{1/2}} = q \ . \tag{5.23}$$

Some values of the scaled excess heat and log-fidelity are shown in Fig. 5.1c-d.

For the equal-time s^z-s^z correlation function, we postulate a non-equilibrium scaling form, and check for scaling collapse. Utilizing translation invariance and the value of the anomalous critical exponent for spin, $\beta = 1/8$, the scaled correlation function is

$$g(\tau,\Lambda,\chi) \equiv G(x,t,L,v)x^{1/4}, \tag{5.24}$$

where $\chi \equiv xv^{1/2}$ is the scaled distance. To solve for the spin correlations, we use a standard trick of breaking up each site of the spin chain into a pair of Majorana fermions [78]:

$$A_j = c_j^\dagger + c_j;\ B_j = c_j^\dagger - c_j \ . \tag{5.25}$$

In terms of these Majoranas, the spin-spin correlation function is

$$\langle s_j^z s_{j+x}^z\rangle = \langle B_j A_{j+1} B_{j+1}\cdots B_{j+x-1}A_{j+x}\rangle \ , \tag{5.26}$$

[1] We work with log-fidelity density because, as shown in Refs. 92 and 95, log-fidelity density is a natural quantity for which the thermodynamic limit ($L\to\infty$) may be taken. One way to see this is to assume we make an average "error" of ϵ per site, such that the fidelity is $\mathcal{F} = (1-\epsilon)^L$. This gives log-fidelity density of $F = -(\log\mathcal{F})/L = -\log(1-\epsilon)$, which remains well-defined as $L\to\infty$.

which can be broken up into a Pfaffian of pairwise expectation values of the form $\langle AA\rangle$, $\langle AB\rangle$, and $\langle BB\rangle$ [96] using Wick's theorem. In our momentum-space pseudo-spin notation, these expectation values are

$$\langle A_j A_{j+\ell}\rangle = \langle B_j B_{j+\ell}\rangle = -i\sum_k \left[\langle\sigma_k^y\rangle \sin(k\ell)\right] \tag{5.27}$$

$$\langle A_j B_{j+\ell}\rangle = \langle B_j A_{j+\ell}\rangle = \sum_k \left[\langle\sigma_k^z\rangle \cos(k\ell) - \langle\sigma_k^x\rangle \sin(k\ell)\right] . \tag{5.28}$$

We then evaluate the Pfaffian numerically[2], and find that good scaling collapse occurs for small v (Fig. 5.2a).

5.1.2 Diagonal observables: Finite size scaling and asymptotic limits

In a finite size system, a gap $\Delta \sim L^{-z}$ opens at the QCP. This has important effects on the dynamics, which are most easily understood for the diagonal observables – excess heat and log-fidelity. From Eqs. 5.21 and 5.22, it is clear that these observables depend only on the excitation probability $p_\kappa^{\rm exc}(\tau)$ and excess heat $q_\kappa^{\rm exc}(\tau)$ for each mode, some values of which are shown in Fig. 5.1a-b. Finite size effects enter by limiting the values of κ that are summed over; for example, the minimum value is $\kappa_{\rm min} = \pi/\Lambda$. As an initial example of finite size effects, in Fig. 5.1e-f we show the scaled diagonal observables as a function of scaled system size Λ for a snapshot taken right at the QCP ($\tau = 0$). In the regime of small scaled size $\Lambda \lesssim 1$, the finite-size gap dominates and the system remains near the ground state. The observables can be understood perturbatively by noting that they must be even under a sign reversal of the ramp rate [98]; thus, to lowest order, the excess heat and log-fidelity densities scale as $Q/L \sim F/L \sim v^2$ in the limit $\Lambda \ll 1$. So for fixed L, in which case $v^{1/2} \sim \Lambda$,

[2] We note that the theory of Toeplitz determinants [97] can, in principle, be used to analytically solve for the correlation function in the limit $x \gg 1$.

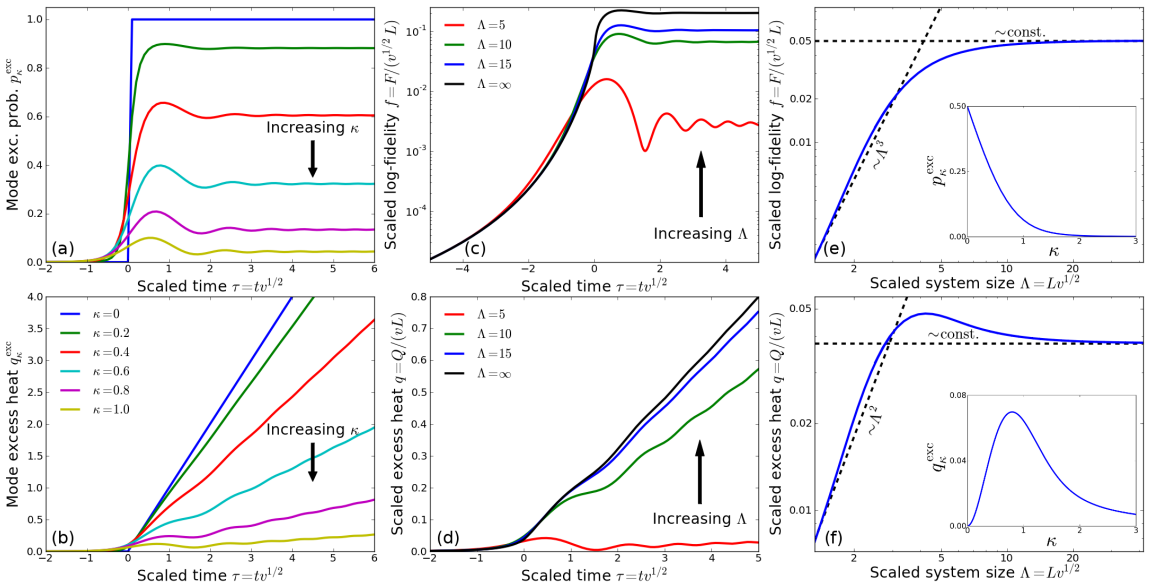

Figure 5.1: Diagonal observables (excess heat and log-fidelity) for the TFI model, which derive from the mode excitation probability $p_\kappa^{\rm exc}$ (a) and mode excess heat $q_\kappa^{\rm exc}$ (b). From these, one obtains the scaled log-fidelity (c) and excess heat (d) for various scaled system sizes $\Lambda = Lv^{1/2}$. (e,f) Finite size scaling for ramps to the QCP ($\tau = 0$). For comparison, the insets show κ dependence of $p_\kappa^{\rm exc}$ and $q_\kappa^{\rm exc}$ at $\tau = 0$. The limits of these scaling functions are detailed in the text.

we expect the scaled observables to asymptote to $q \sim \Lambda^2$ and $f \sim \Lambda^3$; this is confirmed in Fig. 5.1e-f.

In the opposite limit, $\Lambda \to \infty$, finite-size effects vanish; we will refer to this as the thermodynamic limit of KZ scaling (KZ-TDL). As we see in Fig. 5.1e-f, excess heat density and log-fidelity density approach a constant in this limit, thus confirming that they are intensive thermodynamic quantities. In between these two limits, the excess heat shows non-monotonic behavior, which is a manifestation of the non-monotonicity of the mode excess heat $q_\kappa^{\rm exc}$ at the QCP (Fig. 5.1f, inset). While Kibble-Zurek scaling theory is most elegant in the KZ-TDL, where finite size effects become negligible, most current experimental realizations of the TFI chain and related models [2, 13, 14] are for small finite systems. Therefore, understanding finite size effects will be crucial in connecting theory to experiment.

As a final comment on finite-size scaling, we point out that the $\Lambda \to 0$ limit of the KZ scaling form gives the ground state scaling form, since this is limit where the finite size gap is so large that no excitations occur. As such, the KZ scaling forms encompass equilibrium scaling theory. One corollary of this fact is that, if KZ scaling forms exist, then they must have the same critical exponents as the equilibrium scaling theory; otherwise, the equilibrium limit could not be recovered.

Turning now to time-dependence of the diagonal observables, during the initial part of the ramp ($\tau \ll -1$), the system is very nearly in equilibrium (Fig. 5.1c-d). Meanwhile, deep on the FM side of the ramp ($\tau \gg 1$), LZ physics tells us that the excitation probability of mode κ asymptotes to [24]

$$p_\kappa^{\text{exc}} = \exp(-\pi\kappa^2) \ . \tag{5.29}$$

In between, at finite positive τ, the excitation probability for each mode – given by the LZ equation and shown in Fig. 5.1a – oscillates as a function of scaled time before converging to Eq. 5.29. These oscillations show up in the log-fidelity (Fig. 5.1c), excess heat (Fig. 5.1d), and the correlation function (Fig. 5.2c-d). Note that the excess heat q continues to increase linearly as a function of time in the limit $\tau \gg 1$, while the log-fidelity f asymptotes to a constant. Thus the increasing excess heat is caused not by an increase in the occupation of excited states, but by an increase in the energy for these excitations.

5.1.3 Spin correlations: Athermal properties and decoherence

While diagonal observables have much cleaner (i.e., Pfaffian-free) formulations, spin correlations provide greater insight into the properties of the system near the QCP. For instance, the ground state correlations exhibit quasi-long-range order at the QCP,

with a power law tail $G(x) \sim x^{-1/4}$. However, at any finite temperature, these correlations become short range, decaying exponentially on a distance scale set by the thermal de Broglie wavelength for our integrable model. Looking at the correlations right at the QCP for the KZ ramp (Fig. 5.5b, black lines), we see the finite-size crossover from near-equilibrium power-law correlations for $\Lambda = 5$ to a non-equilibrium exponential decay of correlations for $\Lambda = 30$.

However, the TFI chain is integrable, so although a KZ ramp through the critical point creates excitations, the resulting excited states differ markedly from equilibrium thermal states at the same energies. From Eq. 5.29, we see that at long times the populations of modes with $\kappa^2 < (\log 2)/\pi$ are inverted, i.e., at a negative effective temperature with excitation probability greater than 50%. This leads to qualitatively athermal physics in the scaled correlation function, which goes substantially negative by $\tau = 5$ over a range of scaled distances (see Fig. 5.2c-d). This is qualitatively different from any thermal state (Fig. 5.2d), which would have a finite correlation length but never negative correlations. Similar behavior has been seen outside the scaling regime for both slow [99] and instantaneous [100] quenches of the TFI chain and the XY model [101].

There is one more limit that is worth discussing, namely what happens formally as $\tau \to \infty$ while remaining in the scaling regime? We have seen that the diagonal observables are relatively boring in this limit, but the spin correlations may undergo a process known as dephasing. To motivate dephasing, if one were to stop the ramp at some value λ_f and wait a very long time, the phase differences between modes would increase to the point where the phases are essentially random. In such an integrable system, dephasing results in the generalized Gibbs ensemble (GGE, see Ref. 102), given by removing all phase information from each mode; see Chp. 7 for a more detailed discussion of dephasing.

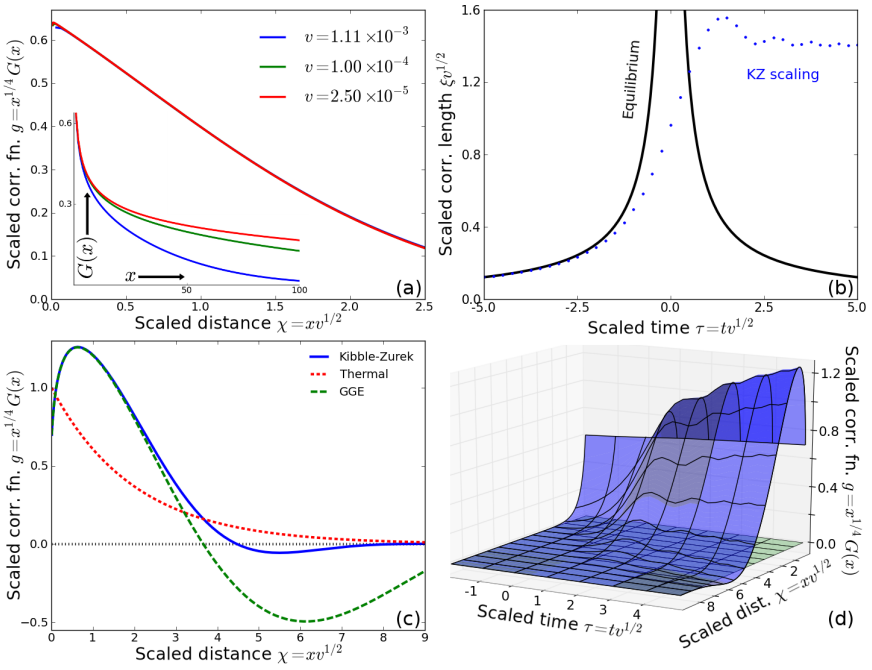

Figure 5.2: (a) Spin-spin correlation function of the TFI chain for a ramp to the QCP ($t = 0$), showing scaling collapse for a wide range of slow ramp rates. The inset shows correlation functions prior to scaling. (b) Scaled correlation length and (c-d) correlation function as a function of scaled time (b-d) and distance (c-d) during the ramp, where we define the correlation length as $\xi v^{1/2} = \sqrt{\int_0^\infty g(\chi)\chi^2 d\chi / \int_0^\infty g(\chi)d\chi}$. This correlation length deviates from equilibrium near $tv^{1/2} = -2.5$ and remains finite at the QCP and beyond. In (c), we compare the KZ correlation function at $\tau = 5$ to a representative thermal correlation function and the GGE correlation function (see text). All data in this figure are in the KZ thermodynamic limit, $Lv^{1/2} \gg 1$.

Since the Hamiltonian changes as a function of time, we must specify to which GGE we wish to compare our correlations. We define the GGE at time τ as the dephased ensemble that one would approach upon freezing the current Hamiltonian and waiting a long time (see Chp. 7):

$$\rho_{\mathrm{GGE}}(\tau) = \prod_{\kappa} \Big[(1 - p_\kappa^{\mathrm{exc}}(\tau)) |\Psi_\kappa^0(\tau)\rangle\langle\Psi_\kappa^0(\tau)| + p_\kappa^{\mathrm{exc}}(\tau) |\Psi_\kappa^1(\tau)\rangle\langle\Psi_\kappa^1(\tau)| \Big] , \tag{5.30}$$

where $|\Psi_\kappa^1(\tau)\rangle$ is the excited state of mode Hamiltonian H_κ. As an example, the GGE spin correlations at $\tau = 5$, which is on the FM side of the QCP, are shown in Fig. 5.2c. In the limit $\tau \to \infty$, the mode Hamiltonians asymptote to $H_\kappa \propto -\sigma_\kappa^z$, so the GGE approaches a fixed value with $|\Psi_\kappa^0\rangle \to |\uparrow\rangle$, $|\Psi_\kappa^1\rangle \to |\downarrow\rangle$, and $p_\kappa^{\text{exc}} \to e^{-\pi\kappa^2}$.

To see if $\tau \to \infty$ gives the dephased ensemble, we consider the phase difference $\Delta\varphi$ between characteristic modes $\kappa = 0$ and $\kappa = 1$, since excitations are exponentially suppressed for $\kappa \gtrsim 1$. Starting from some time $\tau_i \gg 1$, after which the dynamics is effectively adiabatic, the (scaled) energy difference ΔE and phase difference $\Delta\varphi$ are

$$\begin{aligned} \Delta E(\tau) &= \sqrt{\tau^2 + 1} - \tau \approx \frac{1}{2\tau} \\ \Delta\varphi &= \int_{\tau_i}^{\tau_f} \Delta E(\tau) d\tau \approx \frac{1}{2} \log(\tau_f/\tau_i) \end{aligned} \tag{5.31}$$

Since $\Delta\varphi \to \infty$ as $\tau_f \to \infty$, the phase information between modes is lost in the long time limit, so the observables approach those of the GGE.

We note that, for non-linear ramps (say cubic ramps, $\lambda \sim t^3$), a similar scaling theory can be written down [86]. Then the above argument again holds, except now

$$\Delta\varphi \approx \int_{\tau_i}^{\tau_f} \frac{1}{2\tau^3} \to \text{const.} \tag{5.32}$$

as $\tau_f \to \infty$, implying that phase information remains, and the cubic ramp does not approach the GGE.[3] To summarize, both linear and cubic ramps exhibit athermal behavior, such as negative correlations, coming from the inversion of low-momentum modes, but only the linear ramp dephases to the GGE in the long time limit.

[3]We note two points about observing dephasing in practice. First, our arguments only hold when taking the long time limit while remaining near the QCP ($\lambda \ll 1$) to stay in the scaling regime. This limit is difficult to achieve, and is compounded by disorder, which becomes more important in the slow ramp, long time limit. Second, one must know the exact location of the QCP for the cubic ramp; a slight "miss" would allow one to linearize about the actual QCP. The slower the ramp, the more accurately one needs to know the QCP to see cubic scaling.

5.2 Universality of Kibble-Zurek scaling

While the TFI chain is a beautiful theoretical model, it is difficult to realize in the lab. Therefore, if our results are to have experimental merit, we must address an important aspect of scaling theory: universality. The idea of universality is that, in the scaling limit, a large number of terms in the Hamiltonian can be irrelevant to the physics. We have already seen an example of this in the case of the TFI model, where in the scaling limit we could drop higher-order terms in momentum from the expansion of the sine and cosine. While those terms still exist, they do not affect the physics of the scaling limit. A more complicated example is a four-fermion interaction, such as $V = \gamma \sum_j n_j n_{j+1}$ (corresponding to an s^x-s^x interaction in the original spin model). Upon carrying out the same renormalization procedure as in Sec. 5.1, the coefficient γ of this interaction term *decreases*. Thus, in the scaling limit, the four-fermion interaction term flows to zero, and is dubbed "irrelevant." In fact, one can show that for all one-dimensional Hamiltonians breaking only a discrete two-fold (a.k.a. $\mathbb{Z}_2$) symmetry, the TFI critical point discussed above is the only critical point that can exist; every other perturbation is irrelevant [78]. Any one-dimensional model with such a symmetry-breaking is thus said to be in the TFI (or two-dimensional Ising[4]) universality class.

All of these arguments are well understood and accepted for equilibrium scaling. However, while the scaling arguments above can be applied heuristically to the case of non-equilibrium KZ scaling, there is some debate as to whether KZ scaling is universal [103, 104]. In this section, we provide strong evidence for the universality of KZ dynamics by comparing the TFI scaling functions derived in the previous section to those obtained numerically for an experimentally-realizable model in the same (equilibrium) universality class. There have been a number of recent experimental

[4] This name comes from the fact that the ground state properties of the one-dimensional quantum TFI chain can be mapped exactly to the thermodynamic properties of the two-dimensional classical Ising model [78].

advances toward realizing other, non-integrable models in the Ising universality class [2, 13, 14]. For the remainder of this chapter and the next, we focus on one, the Mott insulator in a tilted potential (MITP), realized experimentally in Ref. 2.

5.2.1 Mott insulator in a tilted potential

In Ref. 2, the authors experimentally create one-dimensional chains of interacting bosons, with sufficiently deep lattice and short-range interactions such that the behavior is well described by the repulsive Bose-Hubbard model. They then add a constant magnetic field gradient along the length of the chain, which appears to the hyperfine-polarized bosons as a constant potential gradient. Thus, the experimental Hamiltonian is

$$H_{\text{expt}} = \sum_j \left[-w \left(b_{j+1}^\dagger b_j + b_j^\dagger b_{j+1} \right) + \frac{U}{2} n_j (n_j - 1) - \mathcal{E} j n_j \right] , \tag{5.33}$$

where w is the hopping energy, U is the on-site repulsion, and $\mathcal{E}$ is the effective electric field (derivative of the external potential due to the magnetic field gradient). As we are interested in the Mott insulating phase, we consider $U \gg w$.

An interesting thing happens near $\mathcal{E} \approx U$, as discovered experimentally in Ref. 11 and explained theoretically in Ref. 88; the system begins to spontaneously form dipoles, which we will see are well described by an effective spin Hamiltonian that exhibits a paramagnetic to antiferromagnetic QPT. As it is crucial to understanding how this model fits into the Ising universality class, we summarize the derivation of the effective spin model below.

Consider the Mott insulator at unit filling (one boson per site). As we apply the tilt $\mathcal{E}$, a naive prediction might be that the atoms just fall downhill. However, while such a process might make sense in the presence of a heat bath, here it is forbidden

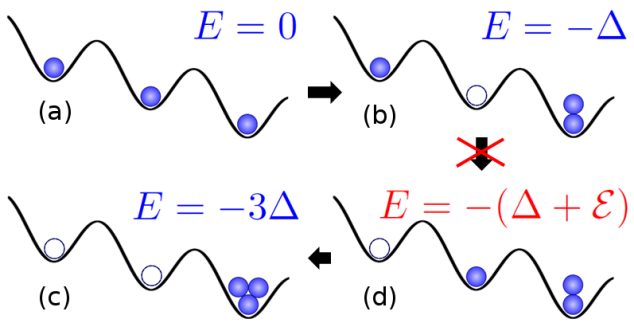

Figure 5.3: Example of states within (a-b) and outside of (c-d) the resonant manifold, showing their energies in the absence of hopping ($w = 0$). Starting from (a) a Mott insulator at unit filling with energy $E = 0$, a single boson can tunnel (b) with small energy gain Δ. However, the state with three atoms on a site (c), while energetically allowed ($E = -3\Delta$), cannot be reached as it requires tunneling through a state at much lower energy (d).

by energy conservation. Indeed, for a large range of tilt $\mathcal{E} \ll U$, the atoms simply remain in the Mott insulating state, localized by interactions.

Eventually, for $|\mathcal{E} - U| \approx w$, the bosons are energetically allowed to tunnel. However, once a boson has hopped downhill, nothing further can happen – it can't tunnel further down the chain, another boson can't fill the spot it vacated, etc. – as all of these result in large change in energy. If the diagonal energy gain upon tunneling a single boson is $\Delta \equiv \mathcal{E} - U$, then in the limit $|\Delta|, w \ll \mathcal{E}, U$ there exists a resonant manifold of states with energy close (within $O(w)$) to that of the Mott insulating state. Denoting the Mott insulating state as $|\cdots 11 \cdots\rangle$, it is clear from the above discussion that hopping a single boson to form a dipole is within the resonant manifold, so $|\cdots 02 \cdots\rangle$ is allowed. It would appear energetically as if states such as $|\cdots 003 \cdots\rangle$ are resonant as well; however such a state requires tunneling through the state $|\cdots 012 \cdots\rangle$, which is in a lower energy manifold (see Fig. 5.3). Therefore, the only states within the resonant manifold are those with ($|02\rangle$) or without ($|11\rangle$) a dipole on each bond. Note also that it is impossible to have dipoles on neighboring bonds.

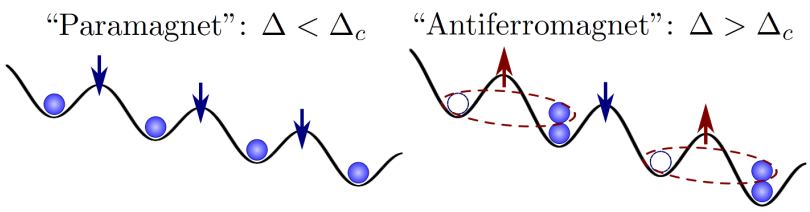

Figure 5.4: Illustration of the Mott insulator in a tilted potential (MITP), realized experimentally in Ref. 2 (see text for details).

Since the presence or absence of a dipole on each bond is a two-level system, we can recast the dipole Hamiltonian as a spin Hamiltonian, where each spin lives on a bond of the original lattice (see Fig. 5.4). Letting $\uparrow$ ($\downarrow$) correspond to the presence (absence) of a dipole, hopping terms correspond to spin rotations in this basis; thus $-wb^\dagger_{j+1}b_j$ maps to $\sqrt{2}wS^x$, where the factor of $\sqrt{2}$ comes from Bose statistics. Similarly, a diagonal energy gain of Δ occurs upon the formation of a dipole; thus we get $\Delta S^z/2$. Finally, as there cannot be two dipoles on neighboring bonds, this maps to a constraint disallowing neighboring up spins. So we see that the tilted boson system is well-described by the effective spin Hamiltonian

$$H_{\text{eff}} = \mathcal{P}\left\{-\sum_l \left[u\, S^x_l + \Delta\, \frac{S^z_l + 1}{2}\right]\right\}\mathcal{P}, \tag{5.34}$$

where l labels the bonds, $S^{x,z}$ are effective spin-1/2 Pauli matrices residing on the bonds, and $\mathcal{P}$ is a projector implementing the constraint that no two neighboring bonds are both spin up. The Mott insulator to dipole quantum phase transition is thus recast as a paramagnetic to antiferromagnetic phase transition (see Fig. 5.4), in which the projector $\mathcal{P}$ produces the antiferromagnetic interactions.

The energy scale of Eq. 5.34 is set by $u = \sqrt{2}w$; we work in units with $u = 1$. In analogy with the TFI chain, we define the parameters $\lambda = \Delta - \Delta_c$ and $v = \partial_t \lambda$, where $\Delta_c \cong 1.31$ is the location of the QCP [88]. The antiferromagnetic correlation function is given by $G(x) = (-1)^x \left(\langle S^z_j S^z_{j+x}\rangle - \langle S^z_j\rangle\langle S^z_{j+x}\rangle\right)$.

5.2.2 Comparison of Mott insulator and Ising dynamics

In order to numerically simulate a ramp in this non-integrable system, we use time-dependent matrix product states (tMPS, see Sec. 3.2.3). We take advantage of the fact that for the Hamiltonian Eq. 5.34, the term within the projectors (call it $H_d(t)$) is a sum of single site Hamiltonians. We thus evolve in time using $|\psi(t+\epsilon)\rangle \approx \mathcal{P}\exp(-i\epsilon H_d(t))|\psi(t)\rangle$ by first exactly applying the single site Hamiltonian $H_d(t)$, an MPO of bond dimension 1, then projecting out configurations with neighboring dipoles using the projection operator $\mathcal{P}$, which is broken up into odd and even sites as MPOs of bond dimension 2. For instance, to project out neighboring spin ups on sites 1 and 2, we use the MPO matrices

$$M_1^{\uparrow\downarrow} = M_1^{\downarrow\uparrow} = M_2^{\uparrow\downarrow} = M_2^{\downarrow\uparrow} = 0 \tag{5.35}$$

$$M_1^{\uparrow\uparrow} = \begin{pmatrix} 0 & 1 \end{pmatrix};\ M_1^{\downarrow\downarrow} = \begin{pmatrix} 1 & 1 \end{pmatrix}; M_2^{\uparrow\uparrow} = \begin{pmatrix} 1 \\ 0 \end{pmatrix};\ M_2^{\downarrow\downarrow} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}\ . \tag{5.36}$$

This projection increases the MPS bond dimension by a factor of two, which is then reduced back to its original value D using SVD compression. The method becomes exact in the limit $\epsilon \to 0$ and $D \to \infty$, and its convergence has been numerically checked by extrapolating in these parameters. Ground states are found by doing imaginary time evolution at fixed λ, also using tMPS. Note that, while we use the MPS ansatz described in Sec. 3.2.1, which is mainly intended for open boundary conditions, here we use it to simulate systems with periodic boundary conditions. While this introduces no additional errors, it does empirically require larger bond dimension than solving the same system with open boundary conditions.

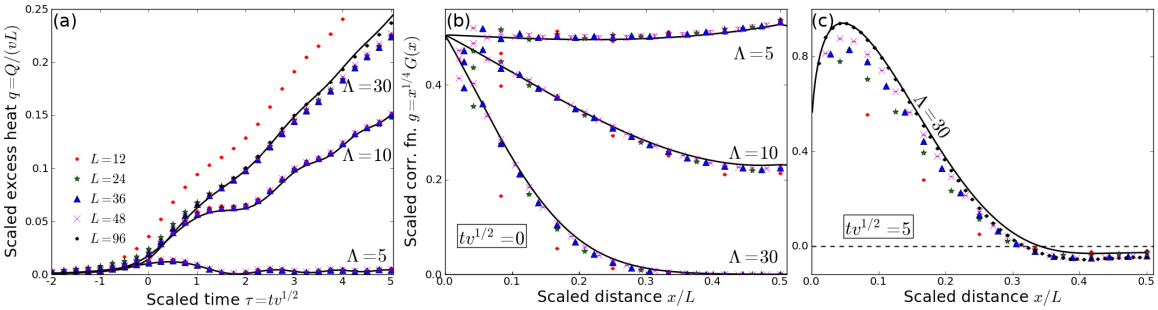

Figure 5.5: Results of tMPS simulations for ramping the MITP. The data show scaling collapse of MITP (colored dots) to the scaling limit of the TFI chain (black lines) as a function of scaled time $\tau = tv^{1/2}$ (a), scaled system size $\Lambda = Lv^{1/2}$ (a,b), and scaled length x/L (b,c). We have checked that data for $\Lambda = 30$ is near the $\Lambda = \infty$ limit for the TFI model (not shown), although finite size effects can still be seen in (c).

Fig. 5.5 shows calculated observables of the MITP model compared to the Ising scaling functions derived earlier.[5] We see clear scaling collapse in all three scaling directions: ramp time t, system size L, and distance x. As expected, simulations do not collapse as well for smaller system sizes. While the dynamic range of the simulations is limited, and the scaling collapse for finite systems is imperfect, we consider our results to be strong numerical evidence for the postulated universality of KZ scaling of systems in the Ising universality class.

We note that, while scaling collapse for the system sizes considered is not perfect at $tv^{1/2} = 5$, the correlation function goes negative as predicted from the TFI scaling function. Therefore, the athermal physics of the KZ ramp is robust against small system sizes and the breaking of integrability, and is a qualitative feature of our model that should be visible experimentally. A major open question is whether such athermal behavior will manifest in KZ ramps near non-integrable QCPs, where the relatively short-time dynamics of the KZ ramp will be in competition with the long-time expectation of eigenstate thermalization (cf. Chp. 7, Ref. 105).

[5]There is an overall non-universal scaling between the TFI and MITP models, given empirically by $v_{\mathrm{MITP}} = 4.84 v_{\mathrm{TFI}}$, $G_{\mathrm{MITP}} = 0.784 G_{\mathrm{TFI}}$, and $t_{\mathrm{MITP}} = 0.557 t_{\mathrm{TFI}}$. All values in Figs. 5.1 and 5.2 are those of the TFI chain, while Fig. 5.5 gives those of the MITP model.

5.3 Conclusions

In conclusion, we have solved for the full Kibble-Zurek dynamic scaling forms of the one-dimensional transverse field Ising chain at zero temperature. We provided numerical evidence for the universality of KZ scaling via tMPS simulations of the MITP model. The TFI scaling functions show a number of unexpected properties, including athermal negative correlations at a distance scale set by the ramp rate and the absence of dephasing in the long time limit for a cubic ramp protocol. As we will see in the next chapter, many of the properties discussed for the disorder-free TFI model with periodic boundary conditions are robust in the experimental case of open boundary conditions and disorder, opening up the exciting prospect of seeing these phenomena experimentally in the near future.

Chapter 6

Experimental considerations for the tilted Mott insulator

One important question that we have not yet addressed is whether the non-equilibrium phenomena discussed in the previous chapter will be visible experimentally. While we can estimate the relevant time scales within the TFI scaling theory, such an analysis neglects two experimentally-relevant effects: disorder and open boundary conditions. Open boundary conditions cannot matter in the thermodynamic limit, although they can affect finite-size scaling. However, disorder is a *relevant* perturbation of the TFI chain, so for a slow enough ramp, disorder must affect the dynamics.

In this chapter, we address the effects of disorder and open boundary conditions on the equilibrium and non-equilibrium properties of the MITP model, with an eye on experimental implications for the setup in Ref. 2. For the experimentally relevant case of open boundary conditions, we show how incipient long-range order leads to even/odd effects in finite size systems. We discuss what effect this has on ground state and non-equilibrium properties of the model, and in particular what aspects of the KZ scaling theory do or do not survive. Finally, we add disorder to the model with open

boundary conditions, and numerically establish the rough disorder strength necessary to disrupt the athermal negative correlations seen in the disorder-free model.

6.1 Ground state with open boundary conditions

One notable effect in the presence of open boundary conditions (o.b.c.) is a distinction between even and odd length systems; for a system with an even number of sites, it is possible to have a perfect dipole covering, but for an odd number of sites, it is not. In this section we will sketch the basic ground state properties of this even/odd effect, including some possible ways to detect it. We will mostly focus on deep in the AFM regime, where domain walls become the fundamental quasiparticles, deriving an effective low-energy theory that describes the spectrum of domain walls in this regime.

For comparison to experiment, whenever we estimate actual numbers, we use the approximate parameters from Ref. 2:

$$\begin{aligned} U &= 400\,\mathrm{Hz} \\ w &= 10\,\mathrm{Hz} \\ \mathcal{E}_f = 1.2U &= 480\,\mathrm{Hz} \\ \Delta_f = \mathcal{E}_f - U &= 80\,\mathrm{Hz}\ . \end{aligned} \tag{6.1}$$

In connection with the spin Hamiltonian, we refer to a singly-occupied site as a domain wall. In addition, for consistency with the remainder of this chapter, L will refer to the number of bonds in the boson chain, i.e., the number of sites in the effective spin model. Therefore, an even length boson chain will have odd L and vice versa.

As a starting point, let's look at the case of 15 and 16 sites with open boundary conditions ($L = 14$ and 15 respectively). The resulting spectra are shown in Fig. 6.1.

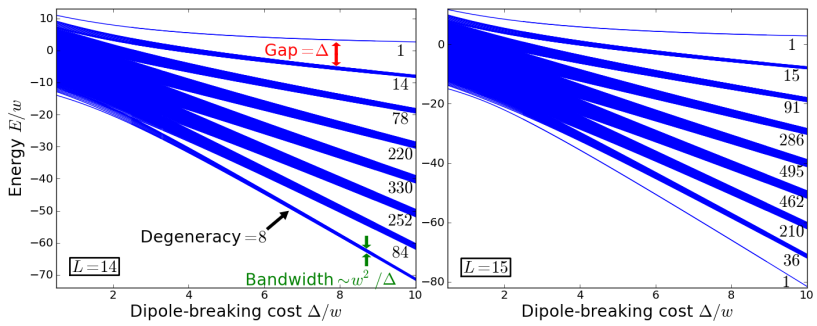

Figure 6.1: Spectra of the MITP model with open boundary conditions, on a lattice of $L = 14$ (a) and 15 (b) bonds. Degeneracies and energy gaps are labeled.

These spectra focus on the AFM side of the phase transition, where nominally there should be two nearly-degenerate ground states (fully degenerate in the thermodynamic limit). Note that the Ising $\mathbb{Z}_2$ symmetry here is translation symmetry, instead of a spin symmetry, and open boundary conditions do not respect this symmetry. Therefore, as a result of open b.c., there is no longer a nearly-degenerate ground state. In its place are ground states of two different types, depending on whether L is even or odd.

We can easily understand the ground states, and indeed the entire spectrum, in the limit $\Delta \gg t$. Consider first the extreme limit $w = 0$ (no hopping). For an even number of sites, there is a single isolated ground state, given in this limit by the perfect AFM, i.e., the state with a perfect dipole covering. For an odd number of sites, on the other hand, there is no perfect dipole covering. Therefore, there are $L/2+1$ degenerate ground states in the absence of hopping, given by a single domain wall on any of the odd-numbered sites.

Before introducing hopping, it's worth developing a clear notation for states with multiple domain walls. Denoting the domain wall as $|1\rangle$, since it is just an unpaired boson, and the dipole as $|d\rangle = |02\rangle$, an arbitrary state can be written as a string of unpaired bosons and domain walls, i.e., $|102021102\rangle = |1dd11d\rangle$. For a system with L bonds and W domains walls, there are $D = (L+1-W)/2$ dipoles filling in the gaps.

The degeneracy of the W domain wall sector is just the number of places to put the domain walls, namely the binomial coefficient $\begin{pmatrix} W+D \\ D \end{pmatrix}$. These degeneracies are indicated in Fig. 6.1. In order to add two domain walls to any given state, one has to break a dipole, which costs energy Δ.

We now want to introduce hopping perturbatively, $w \ll \Delta$. Consider first what it means to hop a domain wall. In the above language, this is a transition of the form $|1d\rangle \to |d1\rangle$ for a pair of sites. This can't happen at first order, but at second order there is one process to do this while remaining on the resonant manifold, namely

$$|102\rangle \to |111\rangle \to |021\rangle \ . \tag{6.2}$$

The intermediate state costs diagonal energy Δ, but gains two hopping terms. Therefore, the matrix element is $w' = 2w^2/\Delta$ (the factor of $2 = (\sqrt{2})^2$ comes from the bosonic commutation relations).

A second, trivial effect comes from the similar process

$$|102\rangle \to |111\rangle \to |102\rangle \ . \tag{6.3}$$

In other words, a dipole can gain energy w' by temporarily breaking and reforming. Therefore, there is a diagonal shift given by w' times the number of dipoles, D.

Finally, also with a matrix element w', there is an effective interaction induced between nearest neighbor domain walls, given by the exchange process

$$|11\rangle \to |02\rangle \to |11\rangle \ . \tag{6.4}$$

Since the intermediate step lowers the energy by Δ, this is a repulsive interaction.

To summarize, at second order in w/Δ, the effective Hamiltonian for the branch with W domain walls and $D = (L + 1 - W)/2$ dipoles is

$$H_{\text{eff}} = \Delta W - w'D - w' \sum_{j=1}^{D+W-1} \left[c_j^\dagger c_{j+1} + c_{j+1}^\dagger c_j - n_j n_{j+1} \right] , \tag{6.5}$$

where $c_j^\dagger$ creates a domain wall on "site" j and $n_j = c_j^\dagger c_j$ is the number operator. The notation c_j has been chosen to emphasize that, due to hardcore constraint (i.e., $n_j \leq 1$) and open boundary conditions, the domain walls can be treated as fermions. The "sites" j on which this effective Hamiltonian is defined are the lattice of the effective $\{|1\rangle, |d\rangle\}$ basis, i.e., a lattice with either a domain wall or a dipole on each site. Thus, the effective lattice size is $L_{\text{eff}} = D + W$.

While the above construction describes the entire spectrum in the limit $\Delta \gg w$, most current work is focused on the ground state properties [2]. Given this, we describe a few concrete experimental implications for the ground state.

First, boundary conditions should be visible in the parity of the number of domain walls. For a system with an even (odd) number of sites, it should only be possible to see an even (odd) number of domain walls. Assuming that this is fairly easily distinguishable, one could go a step farther and see how the histogram shifts as a function of Δ. For sufficiently large Δ, the signal should be dominated by one domain wall (L even) or zero domain walls (L odd). Furthermore, the correction, i.e., the occupation probability of two additional domain walls, should perturbatively go as $\sim (w/\Delta)^2$; for $\Delta_f/w = 8$, as estimated from experimental parameters, this should already by a pretty good approximation (see Fig. 6.2a,c).

Second, within the $W = 1$ manifold for an odd number of sites, Eq. 6.5 becomes a free particle Hamiltonian. So, to see the expected particle-in-a-box behavior of the ground state, one can take a series of measurements to map out the occupation of bosons and post-select measurements with a single domain wall. If the ramp was

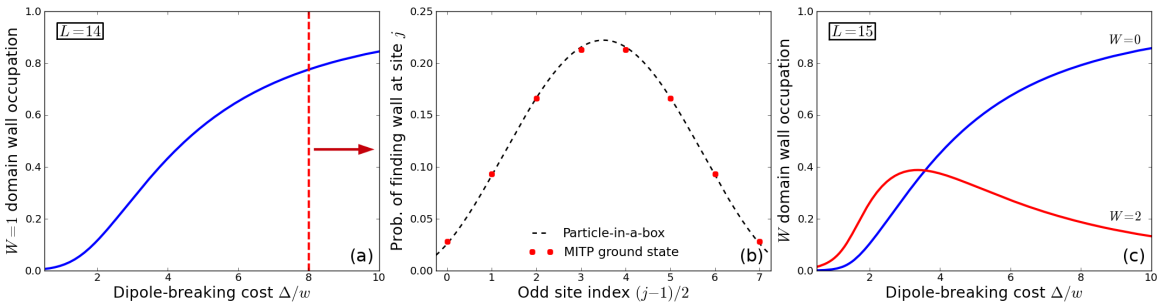

Figure 6.2: (a,b) Ground state of the $L = 14$ MITP model with o.b.c. (a) Probability of finding the ground state wavefunction in the single domain wall sector ($W = 1$) as a function of Δ/w. (b) For the states in this sector at the estimated experimental value $\Delta/w = 8$, a histogram of the domain wall locations (red dots). Black dashed line shows the comparable particle-in-a-box ground state. (c) For $L = 15$, probability of finding the ground state in the pure AFM ($W = 0$) and with one broken dipole ($W = 2$).

adiabatic and started from the Mott insulating ground state, then the domain wall should be in the ground state of the particle-in-a-box (see Fig. 6.2b).

Finally, there is a theoretical point that bears mentioning. As a result of the o.b.c. breaking the $\mathbb{Z}_2$ symmetry of this model, one thing that should not be found, even in the thermodynamic limit, are two degenerate spontaneous-symmetry-broken ground states, as would exist with p.b.c. or in the AFM Ising model (where $\mathbb{Z}_2$ is a spin symmetry). However, there are remnants of the long-range order associated with the "symmetry-broken" ground states. In o.b.c., while even/odd effects do show up in the AFM ground state, they do not show up in the PM ground state (the Mott insulator). This is related to the fact that the insulator is fundamentally a quantum disordered state, with a short correlation length, so no information can pass from one boundary to the other that would allow it to globally notice the boundary conditions. However, since the AFM is ordered, what happens at one boundary does affect what happens at the other boundary, so even/odd effects occur precisely because of long-range order. This is important because one could imagine getting similar long domains of dipoles by just using locally staggered effective magnetic fields to favor the configuration

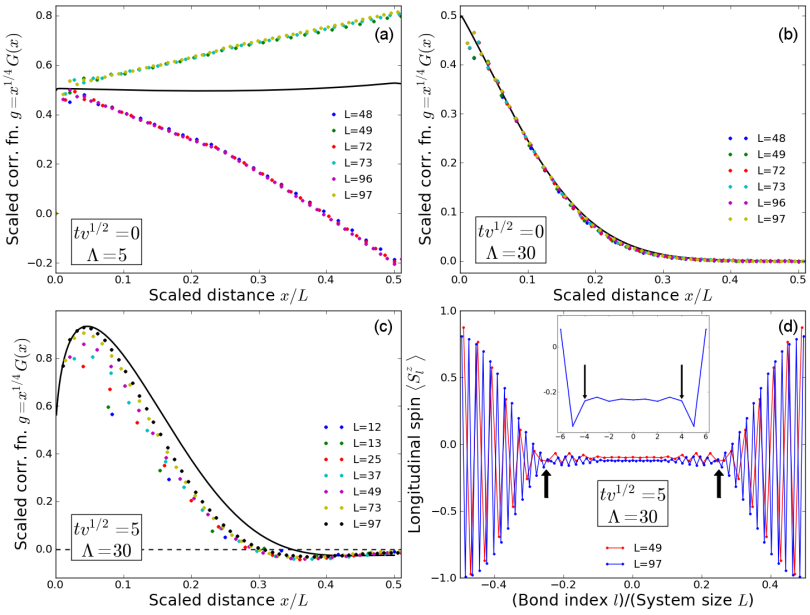

Figure 6.3: (a-c) Spin-spin correlation function of the MITP model with open boundary conditions, compared to the TFI scaling function with p.b.c. (black lines). (d) Expectation values of the z components of the individual spins, $\langle S_l^z \rangle$, for the same parameters as (c) with $L = 97$, 49, and 13 (inset) bonds, where l is the bond index. Arrows indicate the 180° phase shift (AFM domain wall) that yields the negative correlations in (c).

$|0202\ldots\rangle$. However, this state would not have a particle-in-a-box domain wall ground state with an odd length system; it would just put the one domain wall locally at whichever end the local fields were designed to place it at.

6.2 Kibble-Zurek with open boundary conditions

We now turn to the effect of open boundary condition on the Kibble-Zurek dynamics of the MITP model. Comparing the observables to those obtained from the finite size scaling theory of the TFI model with periodic boundary conditions (p.b.c.), we expect the the two to be identical in situations where the finite size effects vanish. In other words, boundary conditions should not matter in the KZ-TDL. However, once finite size effects become relevant, we should no longer expect the systems to match.

To compare against the TFI model, we numerically simulate a ramp of the MITP model with o.b.c. using the same methods as Sec. 5.2.2. The only difference is that, in projecting out neighboring spin ups, we do not apply the "end projector" that acted on sites 1 and L. We solve for the spin-spin correlation function $G(x)$ for the same criteria as in Fig. 5.5, the results of which are show in Fig. 6.3a-c. We see that large Λ matches the TFI model much better than small Λ, confirming our expectation that o.b.c. are irrelevant in the KZ-TDL. Note that Fig. 6.3 shows scaled correlations as a function of scaled length x/L, and we see that the solutions match better at small distances $x \ll L$, where boundary effects are not important, as opposed to $x/L \sim 0.5$, where boundary effects are non-trivial. Also notice that the even/odd effects show up quite strongly in situations where the boundaries matter, as expected.

An additional complication with o.b.c. is that, compared to the translationally invariant system with p.b.c., there is no unique way to define $G(x)$; due to the loss of translation invariance,

$$\langle S_j^z S_{j+x}^z \rangle - \langle S_j^z \rangle \langle S_{j+x}^z \rangle \neq \langle S_j^z S_{j+x}^z \rangle - \langle S^z \rangle^2 \ , \tag{6.6}$$

where $\langle S^z \rangle \equiv \frac{1}{L} \sum_j \langle S_j^z \rangle$. Choosing either side of Eq. 6.6 in defining $G(x)$ gives quantitatively different results. Similarly, including sites near the end of the chain in the definition of $G(x)$ introduces end effects whereas considering only spin correlations between sites in the middle third of the chain emphasizes the bulk behavior at the expense of throwing away data. The results in Fig. 6.3 use the definition

$$G(x) \equiv (-1)^x \frac{1}{L-x} \sum_{j=1}^{L-x} \left[\langle S_j^z S_{j+x}^z \rangle - \langle S^z \rangle^2 \right] \ . \tag{6.7}$$

Empirically (data not shown), the various definitions of $G(x)$ that we described above gave quantitatively different, but qualitatively similar, results. However, as expected, all of the definitions converge to the same result in the KZ-TDL.

Most importantly, the athermal negative correlations that we saw with p.b.c. remain in the case of o.b.c. (Fig. 6.3c), and therefore are robust with respect to both boundary conditions and small system size. In fact, o.b.c. provides an additional route to seeing negative correlations experimentally. As seen in Fig. 6.3d, due to the absence of the end projector, the system selects the AFM with higher spin-up probability at the ends. By breaking the degeneracy of the ground states, the domain walls associated with negative correlations can be seen simply by measuring the expectation value of S_j^z for each site j, rather than measuring correlation functions. As expected from Fig. 6.3c, for constant Λ these domain walls occur at a roughly constant value of j/L. This prediction, confirmed in Fig. 6.3d, implies that the position of the domain wall relative to the edge of the sample is tunable by the ramp rate v.

6.3 Kibble-Zurek with disorder

Experimentally, of course, disorder is omnipresent. In this section, we attempt to determine whether the non-equilibrium properties we have discussed are robust against the presence of weak-to-moderate disorder. It is well known that disorder is a relevant perturbation for the TFI model [106, 107], which has profound effects on its Kibble-Zurek scaling [108, 109]. Therefore, if one ramps slowly enough or probes correlations on a sufficiently large length scale, disorder effects become relevant and qualitatively change the scaling behavior. This topic is interesting in its own right, and deserves a separate discussion. However, our focus here will be on the robustness to disorder of the disorder-free KZ scaling theory in small, experimentally accessible systems.

The particular type of disorder that we consider is a random on-site longitudinal field h_l^z,

$$H_{\text{eff}} = \mathcal{P}\left\{-\sum_l \left[S_l^x + \left(\Delta + h_l^z\right)\left(\frac{S_l^z+1}{2}\right)\right]\right\}\mathcal{P}\ , \tag{6.8}$$

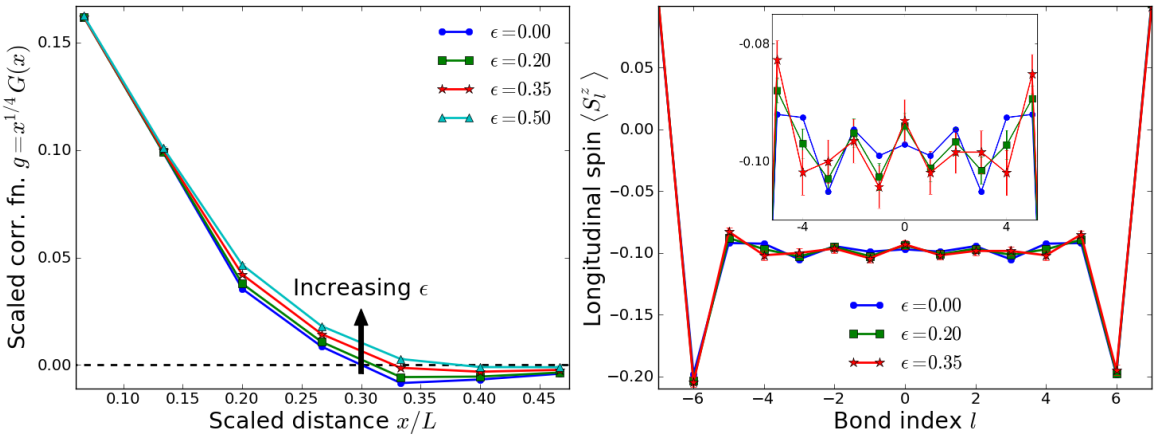

Figure 6.4: $L = 15$ site MITP model with open boundary conditions and disorder after a KZ ramp with the same parameters as Fig. 6.3d. (a) Spin-spin correlations and (b) bond-by-bond $\langle S^z \rangle$ are shown for a range of disorder strength, with average and standard error taken over 1000 disorder realizations. The inset to (b) zooms in on the middle of the chain.

using the same notation as Eq. 5.34 and with h_l^z sampled from a Gaussian of width ϵ. This is only an approximate model of the experimental disorder [110], but it should give a good initial feel for disorder effects.

We use tMPS to simulate the time evolution of this Hamiltonian for an $L = 15$ bond MITP chain with open boundary conditions and a wide range of disorder strength, the results of which are shown in Fig. 6.4. Clearly, strong enough disorder modifies the behavior of $\langle S^z \rangle$ and its correlations, decreasing the size of the qualitatively athermal negative correlations. However, a non-zero signal of negative correlations robustly remains up to the largest disorder strength that we simulate ($\epsilon = 0.5$). We estimate that the experimental conditions, with hopping energy 10 Hz and 5 Hz of disorder, correspond to $\epsilon \approx 0.35$. At this disorder strength, while the negative correlation signal is smaller than at zero disorder, it may still be measurable. Given the data in Fig. 6.4, we would recommend looking for these negative correlations via measurements of the full correlation function as opposed to looking

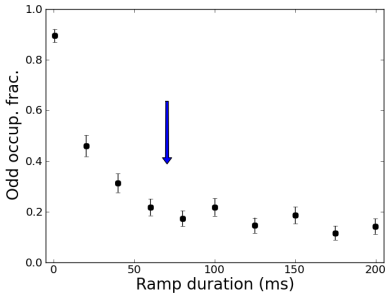

Figure 6.5: Estimated ramp rates necessary to see athermal negative correlations in a twelve-site bosonic chain. The experimental data, borrowed from Ref. 2, shows their measurement of the "paramagnetic fraction," namely the fraction of sites with odd bosonic occupation.

for domain walls in $\langle S^z \rangle$, as the correlation function signal is stronger than the domain wall signal near $\epsilon \approx 0.35$.

6.4 Conclusions

In this chapter we have investigated the properties of the MITP model under experimentally relevant conditions: with open boundary conditions and in the presence of disorder. We find that these experimental constraints can be a blessing in disguise, in some cases yielding phenomena not seen in the translation invariant case. To this end, we have made a number of concrete proposals for experimental measurements that should be within the reach of the current experimental setup [2]. In particular, we have shown that these systems have strong even/odd effects, which are inextricably tied to the long-range AFM order of the ground state. These effects show up in both the ground state and in the KZ dynamics, although for fast enough ramps the KZ length scale is much less than the system length, so boundary conditions become irrelevant. In this limit, open boundary conditions actually provide a new paradigm for measuring the athermal spin correlations, as the broken translation symmetry allows one to measure S^z directly instead of measuring its correlations. Finally, we

find that the athermal correlations are remarkably robust against disorder, so that it may be possible to see this effect even at current experimental levels of disorder.

Based on the robustness of the negative correlations against both boundary conditions and disorder, we can estimate the time scale necessary to see this phenomenon experimentally. Using the approximate experimental energy scales from Eq. 6.1 and Ref. 2, we simulate the MITP model with open boundary conditions for the unfrustrated case (L odd) and look for the presence of negative correlations, using the correlation function defined in Eq. 6.7. While the smallest system shown in Ref. 2, a six-site bosonic chain, does not exhibit negative correlations due to finite size effects, they occur for all larger systems with L odd. In Fig. 6.5, we show the estimated minimum ramp rate necessary to see negative correlations for a twelve-site bosonic chain. Clearly, the time scales necessary to see these non-equilibrium effects are well within the experimental range. [1]

[1] These estimates do not include the effects of disorder. While this would quantitatively modify the predicted ramp rates at current levels of experimental disorder, the numbers will change as the experiments continue to improve.

Chapter 7

Finite-size crossover of eigenstate thermalization

The behavior of isolated quantum systems are a topic of renewed recent interest [111], driven in large part by their experimental realization in ultracold atoms. One particularly interesting question involving the dynamics of isolated systems is their long-time behavior when started in some out-of-equilibrium state. After long enough times, do these systems equilibrate? If so, what are the properties of the equilibrium state, and how robust are these properties to changes in the initial state? Finally, how does the isolated quantum equilibrium compare to the non-isolated equilibrium given by quantum statistical mechanics?

As we will see, the answer to these questions is quite complicated. In particular, it is now well-understood that, while most initial conditions will eventually evolve to a long-time equilibrium, the properties of this equilibrium state depend on interactions and disorder in the Hamiltonian of interest. In this chapter, we will discuss what is understood about the long time relaxation of an isolated quantum system, and present our preliminary data on the role of interactions on long time equilibrium of finite size isolated systems.

7.1 Dephasing in many-body systems

Consider an arbitrary time-independent Hamiltonian H with energy eigenstates $|\alpha\rangle$, under which we evolve a wavefunction $|\psi(t)\rangle$[1] with initial state

$$|\psi(0)\rangle = \sum_{\alpha} C_{\alpha}|\alpha\rangle \ . \tag{7.1}$$

Then the wavefunction at time t is given by

$$|\psi(t)\rangle = \sum_{\alpha} C_{\alpha} e^{-iE_{\alpha}t}|\alpha\rangle \ , \tag{7.2}$$

where $|\alpha\rangle$ corresponds to energy eigenvalue E_{α} of the Hamiltonian H. Furthermore, we will assume that the energy levels are non-degenerate ($E_{\alpha} \neq E_{\beta}$ for all $\alpha \neq \beta$). We work with such Hamiltonians in a finite Hilbert space, so that expressing $|\psi\rangle$ as a discrete sum over non-degenerate levels is justified. One must take great care in approaching the thermodynamic limit, where the discrete energy spectrum approaches a continuum.

For this wavefunction, the expectation value of observable $\mathcal{O}$ is

$$\langle \mathcal{O}(t) \rangle = \sum_{\alpha\beta} C^{*}_{\beta} C_{\alpha} e^{-i(E_{\alpha}-E_{\beta})t} \mathcal{O}_{\beta\alpha} \ , \tag{7.3}$$

where $\mathcal{O}_{\beta\alpha} \equiv \langle \beta|\mathcal{O}|\alpha\rangle$. Clearly the time evolution of this observable is complicated, involving quasiperiodic oscillations dependent on both the eigenstate energies and their matrix elements with respect to the observable. However as we will see in the next section, its properties can be much better understood in the long time limit, where the off-diagonal terms in Eq. 7.3 tend to cancel, due to their quasi-random phases.

[1] We note that, as in Sec. 2.2.1, these arguments can be readily generalized to cases where the system starts in a mixed state density matrix.

7.1.1 Generic dephasing: the diagonal ensemble

Taking the long time average of Eq. 7.3, which we denote $\langle\langle \mathcal{O} \rangle\rangle$, all off-diagonal terms of the sum drop out, leaving

$$\langle\langle \mathcal{O} \rangle\rangle = \sum_{\alpha} |C_\alpha|^2 \, \mathcal{O}_{\alpha\alpha} \ . \tag{7.4}$$

Notice that this is the same as the expectation value of $\mathcal{O}$ for the density matrix

$$\rho_{\text{diag}} = \sum_{\alpha} |C_\alpha|^2 \, |\alpha\rangle\langle\alpha| \ , \tag{7.5}$$

known as the diagonal ensemble. The process in which pure state observables approach those of the diagonal ensemble, in which the off-diagonal phase information of the pure state is lost, is known as dephasing.

Before proceeding to analyze the diagonal ensemble for some specific cases, we note some of its general properties. First, while the wavefunction and its associated density matrix are always in a pure state, the diagonal ensemble can be a mixed state. This comes from the fact that phase information was lost, although not to environmental dephasing as in many mixed states, but instead to time averaging.

A second point regarding the diagonal ensemble is that, while phase information is not preserved, the weight $|C_\alpha|^2$ in each eigenstate is, meaning that whatever energy distribution is initially prepared, that distribution is stationary under time evolution. This is in contrast to an open system, where thermal equilibrium is reached by exchanging energy with the environment and settling into a new energy distribution given by the Gibbs ensemble.

7.1.2 Integrable dephasing: the generalized Gibbs ensemble

While integrability of a quantum system is primarily thought of as a nice property for finding analytical solutions of a quantum many-body problem, it is of much more fundamental importance in the context of dynamics. Before explaining this, we state our definition of integrability: by integrable, we mean a quantum system whose Hamiltonian is separable, i.e., $H = \sum_k H_k$ and $[H_k, H_{k'}] = 0$, with a number of conserved modes k that scales polynomially with the number of sites and/or particles in the system (logarithmically with the size of the Hilbert space). The prototypical example of this is free fermions or bosons, where k are then the momenta of the modes. While this class encompasses a wide range of models, including the transverse-field Ising chain, it is a fairly restrictive definition that is not universally agreed upon [111]. For example, this definition does not include models that can only be solved by the Bethe ansatz, such as the Lieb-Liniger gas [112, 113].

As a consequence of integrability, the number of independent conserved quantities in the diagonal ensemble can be significantly reduced. One can see this by noting that the energy eigenstates must now be separable: $|\psi\rangle = \prod_k |\psi_k\rangle$. As a result, if the system is initialized in a product state of these modes k, then the diagonal ensemble is also separable and, while it is still true that there are exponentially many conserved quantities given by the weight in each energy eigenstate $|\alpha\rangle$, these quantities are no longer independent of each other. For instance, if the (non-degenerate) Hamiltonian is that of free fermions, then it is formally true that the diagonal ensemble for a system initialized in such a product state can be written

$$\rho_{\mathrm{GGE}} = \prod_k e^{-\beta_k H_k} = e^{-\sum_k \beta_k H_k} \ , \tag{7.6}$$

which is known as the generalized Gibbs ensemble (GGE, see Ref. 102). Here β_k is a Lagrange multiplier, known as the inverse mode temperature, which is chosen to

give the correct occupation of mode k to match the initial conditions. This is not an exact equality for the case of a free boson Hamiltonian, as not just the number operator in mode k, $n_k = b_k^\dagger b_k$, but also its higher moments are conserved quantities. However, for physically reasonable initial conditions it is often found [102] that the GGE with only a single mode temperature set to match $\langle n_k \rangle$ gives a good quantitative description of the time-averaged observables in a system of free bosons.

As there are no interaction terms between modes, the GGE can be effectively used to describe integrable systems out of thermal equilibrium, particularly cases where the momentum distribution differs from thermal equilibrium. Indeed, the GGE was first postulated to describe such a phenomenon [102], which had been experimentally realized in a one-dimensional Bose gas [35]. Since that time, the GGE has been successfully applied to a number of theoretical and experimental situations [35, 37, 102], and its ability to correctly describe the non-equilibrium properties of isolated integrable systems is by now well accepted.

7.2 Thermalization in the generalized Ising chain

To summarize, the long time average of observables for an arbitrary Hamiltonian H and initial state $|\psi(0)\rangle$ are well described by the diagonal ensemble

$$\rho_{\text{diag}} = \sum_\alpha |C_\alpha|^2 \, |\alpha\rangle\langle\alpha| \ , \tag{7.7}$$

which for an integrable system initialized in a product state can simplify to the generalized Gibbs ensemble (GGE),

$$\rho_{\text{GGE}} = \prod_k e^{-\beta_k H_k} \ . \tag{7.8}$$

While the properties of these density matrices are not immediately clear, what is clear is that they will not, in general, be the same as the Gibbs ensemble $\rho_G = e^{-\beta H}$. One might therefore be skeptical, since the success of ultracold atoms relies on the ability to cool these isolated quantum gases to very low temperatures, and if temperature were ill-defined, one would need to seriously re-examine most ultracold atom experiments in the literature.

Fortunately, quantum mechanics provides a mechanism for effective thermalization in closed systems, loosely related to the classical statistical mechanics ideas of ergodicity and chaos. At the most basic level, the idea behind thermalization of closed systems is that, while the system as a whole can never be thought of as truly achieving a Gibbs distribution, the reduced density matrix of a small subsystem can approach the Gibbs ensemble. The remainder of the system acts as a heat bath for the subsystem, and thus any observable that is "reasonable" will have properties well-described by the Gibbs ensemble.[2]

To understand thermalization in closed systems, we examine the behavior within a microcanonical ensemble with narrow energy window ΔE. This approach is not fundamentally different than the canonical ensemble implicitly used above, harkening back to the arguments of classical statistical mechanics. Small systems may have a wide range of support for the distribution of weight over energy levels, $p(e)$, where $e = E/\mathcal{V}$ is the energy density. However, in the thermodynamic limit this distribution becomes sharply peaked, leading to identical expectation values in the microcanonical, canonical, and grand canonical ensembles.

[2] The term "reasonable" here is not particularly well-defined or understood. It is common to instead say that thermalization only works for *local* observables [105, 114], but even then non-local properties such as momentum distributions tend to also be well-described thermally. For the remainder of this chapter, we will leave the term loose, noting that it remains an interesting open question precisely what observables will notice differences from the thermal ensemble.

7.2.1 The eigenstate thermalization hypothesis

Consider a microcanonical ensemble Ω centered at energy E with small width ΔE, which contains N_Ω energy eigenstates. By the postulates of statistical mechanics, the energy of an arbitrary observable $\mathcal{O}$ must, in the long time average, match its microcanonical average:

$$\langle\langle \mathcal{O} \rangle\rangle = \langle \mathcal{O} \rangle_{\text{mc}} \implies \sum_{\alpha \in \Omega} |C_\alpha|^2 \, \mathcal{O}_{\alpha\alpha} = \frac{1}{N_\Omega} \sum_{\alpha \in \Omega} \mathcal{O}_{\alpha\alpha} \ , \tag{7.9}$$

where the C's are normalized microcanonically: $\sum_{\alpha \in \Omega} |C_\alpha|^2 = 1$. There are two main possibilities for how this can come about. One possibility is that, for all "reasonable" initial conditions, either the values of $|C_\alpha|^2$ are all very close to N_Ω^{-1} or the values of C_α and $\mathcal{O}_{\alpha\alpha}$ are randomly distributed and uncorrelated, such that their sum gives the microcanonical average. This option, known as canonical typicality, seems to work for a variety of observables and initial conditions [115, 116], usually argued for using an assumption of locality of observables and Hamiltonians. While canonical typicality can be useful in certain situations, such as understanding when integrable Hamiltonians give effectively thermal behavior [116], it can certainly be defeated by appropriately chosen initial conditions. This has been seen experimentally in various setups that are well described by a GGE, but not a thermal ensemble [35, 37].

A more robust possibility known as the eigenstate thermalization hypothesis (ETH) was originally proposed by Deutsch [117] and Srednicki [118] and has since been seen numerically [105] in certain systems using large-scale exact diagonalization. The eigenstate thermalization hypothesis postulates that, for each eigenstate $|\alpha\rangle$ in the microcanonical ensemble and in the limit of large system size, the expectation value of any "reasonable" observable $\mathcal{O}$ in $|\alpha\rangle$ approaches that of the microcanonical ensemble: $\mathcal{O}_{\alpha\alpha} \approx \langle \mathcal{O} \rangle_{\text{mc}}$. Its name derives from the fact that, if it holds, then each eigenstate "looks" thermal, since its observables can't be distinguished from

those of the equipartitioned microcanonical ensemble. One major advantage of the ETH is that, since thermalization occurs for each eigenstate, the time averaged value of the observable does not depend on initial conditions (within the microcanonical ensemble).

It is now widely accepted that, for a broad class of non-integrable systems, the properties of many observables obey the eigenstate thermalization hypothesis [105, 114]. However, it is also well understood that there are situations where eigenstate thermalization does not occur. One such situation which we have already discussed is the case of integrable systems, where two states close in energy can nevertheless look qualitatively different from the standpoint of observables, since they can be made from completely different modes which do not interact. Another important limit where eigenstate thermalization breaks is in the presence of strong disorder, even with interactions; this is known as many-body localization [119, 120]. Understanding the behavior in between perfectly integrable and/or localized systems and strongly interacting non-integrable systems with respect to eigenstate thermalization is a major open question. In Sec. 7.2.3, we will present some initial numerical data and analytical thoughts on the crossover between non-thermalizing integrable systems and strongly non-integrable systems that obey the ETH, using a generalization of the one-dimensional Ising model as a numerical test case.

7.2.2 The eigenstate thermalization crossover

Consider an integrable Hamiltonian H_0 to which a non-integrable perturbation γV is added. For a finite-size system, it seems fairly evident that eigenstate thermalization cannot occur at an arbitrarily small γ, since the distribution of eigenvalues and their observables must vary smoothly from the integrable $\gamma = 0$ case, where the ETH is not satisfied. Therefore, the onset of eigenstate thermalization in finite size systems must take the form of a crossover, the properties of which were first investigated by Rigol

[121] for 1D hardcore bosons with next nearest neighbor hopping and interactions to make them non-integrable. What remains unclear is what happens to this finite size crossover as the system size is increased toward the thermodynamic limit (TDL).

One possibility is that the crossover will asymptote to some sort of phase transition with the onset of eigenstate thermalization at a finite, non-zero value γ_c. There is numerical evidence for such a phase transition in the case of many-body localization [120], and it has been suggested [114] that such a phase transition could exist in the case of the integrability/non-integrability crossover by analogy to certain classical systems, where there can exist a non-zero amount of integrability-breaking γ – known as the Kolmogorov-Arnold-Moser (KAM) threshold [122–124] – below which the system does not exhibit classical chaos.

A second possibility is that the onset of thermalization in the TDL occurs at arbitrarily small γ. In the next section, we provide evidence that appears to support this possibility. While our numerical evidence and theoretical predictions are not yet sufficient to strongly argue for this hypothesis, we point out that at least one other group has produced similar evidence [125]. We therefore present the data and ideas that we have so far on this topic, with the hope that our on-going efforts will soon be able to more strongly support or reject this hypothesis.

7.2.3 Results for generalized Ising chain

The model that we consider is a generalization of the 1D Ising chain in which non-integrable interactions are added. In particular, we add a nearest-neighbor $s^x s^x$ interaction which, while irrelevant near the critical point in the low energy renormalization group sense, nevertheless appears to break the integrability. The Hamiltonian is

$$H = -\sum_{j=1}^{L} \left[(1-\lambda)\, s_j^x + s_j^z s_{j+1}^z + \gamma s_j^x s_{j+1}^x \right] \ , \tag{7.10}$$

where $\gamma = 0$ is the integrable point which we tune away from. Doing the same set of transformations as were used to solve the TFI model (see Chp. 5), the above Hamiltonian can be re-expressed in terms of Wigner fermions. The $s^x s^x$ term involves four fermion scattering,

$$\begin{aligned} H_{xx}/\gamma &= L - 4\sum_k n_k + \frac{4}{L}\sum_{k,k'} n_k n_{k'} \left[1 - \cos(k-k')\right] + \\ &\quad \frac{4}{L}\sum_{\substack{k_1<k_2 \\ k_1'<k_2'}} c^\dagger_{k_1'} c^\dagger_{k_2'} c_{k_2} c_{k_1} \left(e^{ik_2'} - e^{ik_1'}\right)\left(e^{-ik_2} - e^{-ik_1}\right)\delta_{k_1+k_2,k_1'+k_2'} \ . \quad (7.11) \end{aligned}$$

This term conserves fermion number modulo 2, so we continue to work in the even fermion parity sector. Also, noting that the total fermion momentum P is conserved, we work in the $P = 0$ sector, additionally restricting ourselves to the subspace with even spatial parity (an extra symmetry within the $P = 0$ sector).[3]

As the ETH is a statement regarding the properties of individual energy eigenstates, we numerically perform exact diagonalization on this Hamiltonian to gain access to these eigenstates. We then wish to look at the spread of observable expectation values for eigenstates within some microcanonical energy window ΔE, but first we must specify the energy E (or equivalently temperature) where this microcanonical ensemble is centered. Borrowing on the ideas of Pal, Oganesyan, and Huse [120, 126], we consider the onset of thermalization at infinite temperature, where the canonical ensemble is $\rho = \mathbb{1}$. Then, for a system in the TDL, a microcanonical ensemble corresponding to infinite temperature is defined as an energy window ΔE centered at the middle of the energy spectrum, that captures a vanishingly small energy density $\Delta E/L \to 0$ as $L \to \infty$. One convenient method for doing this is to include only the middle third of the energy spectrum, i.e., the eigenstates with energy $E_{N/3}$ through

[3] While all of the preceding arguments considered properties of a full Hamiltonian, they really only hold within a given symmetry sector. For example, as there are no interactions between sectors of different total momentum P in the generalized Ising model solved here, they cannot thermalize between one another, but rather thermalize within each sector.

$E_{2N/3}$ if the N eigenstates are sorted according to $E_1 < E_2 < \ldots < E_N$. Since the r.m.s. width of the energy distribution scales as $\sqrt{L}$ [120] by a simple argument,[4] the middle third of the spectrum incorporates energy densities $\Delta E/L \sim 1/\sqrt{L}$, which vanishes in the TDL as desired. However, its worth noting that for a finite size Hamiltonian, this particular choice of microcanonical ensemble includes some states at finite temperature, leading inevitably to finite size effects. For these eigenstates, we attempt to discover what properties characterize thermalization, and how they scale with system size.

Prior to seeing thermalization, we must first confirm that our Hamiltonian is indeed non-integrable. While there is no cut-and-dry analytical test for this, one successful technique which has been applied in many-body localization is to look at the many-body spectral statistics. For a non-interacting problem, different modes don't talk to one another, meaning that two many-body eigenstates can lie arbitrarily close in energy. Once integrability-breaking interactions are introduced, however, closely-spaced energy levels repel one another. Therefore, while the density of states looks very similar with or without interactions (Fig. 7.1a-b), the distribution of energy spacings between neighboring levels changes significantly as interactions are turned on, exhibiting a dip near zero due to level repulsion (Fig. 7.1c-d). We can encompass the level spacing in a single statistic [126]

$$r_n = \frac{\min\left(E_{n+1} - E_n, E_n - E_{n-1}\right)}{\max\left(E_{n+1} - E_n, E_n - E_{n-1}\right)} , \tag{7.12}$$

where the N energy levels are sorted as $E_1 < E_2 < \ldots < E_N$. The average value of this statistic should change from $\langle r \rangle \approx 0.386$ in the non-interacting limit, given by the statistics of an uncorrelated (Poissonian) distribution, to $\langle r \rangle \approx 0.530$ in the strongly interacting limit, given by the statistics of a Gaussian orthogonal ensemble

[4] Break H up into a sum of on-site Hamiltonians H_j. Then $\langle H^2 \rangle = \sum_{ij} \langle H_i H_j \rangle$. For our nearest neighbor Hamiltonian, $\langle H_i H_j \rangle = 0$ if $|i - j| > 1$ (for an appropriate breakup of H). Since $\langle H_i^2 \rangle \sim \langle H_i H_{i+1} \rangle \sim 1$ as the matrix elements are of order unity, we find that $\langle H^2 \rangle \sim L$.

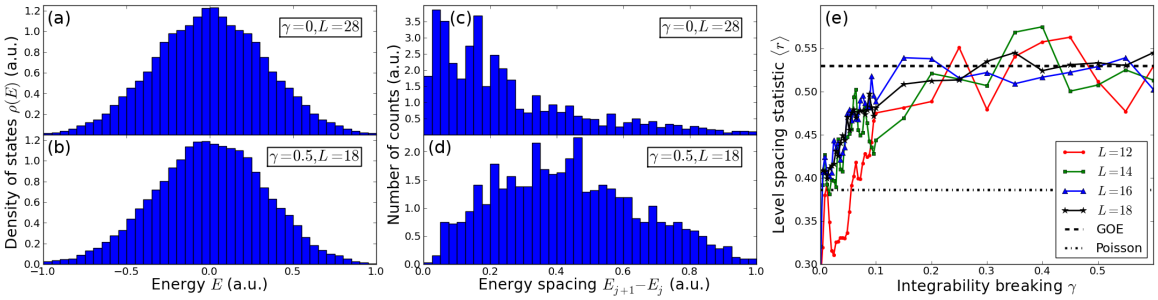

Figure 7.1: (a,b) Density of states for the generalized Ising model in Eq. 7.10, exhibiting no significant differences between the integrable case (a) and the non-integrable case (b). (c,d) Energy level spacing for the same values of the integrability-breaking parameter γ, demonstrating level repulsion for the non-integrable Hamiltonian. In (c), a large peak at $E_{j+1} - E_j = 0$, corresponding to degenerate levels in the integrable model, has been removed. (e) Finite size and γ dependence of the level spacing statistic r, defined in Eq. 7.12. The dashed lines shows the expected value of r for GOE and Poissonian level statistics [120]. All panels show data at $\lambda = \frac{1}{2}$, restricted to the total momentum $P = 0$ sector with even spatial parity. Panels (c-e) are for infinite temperature, i.e., only using the middle third of the energy spectrum.

(GOE) from random matrix theory [126]. While this general trend holds (Fig. 7.1e), confirming that the γ term breaks integrability, the data is much too noisy to establish finite-size scaling behavior. Therefore, we search for a better tool to investigate the finite size scaling of eigenstate thermalization.

According to the ETH, all reasonable observables should undergo a crossover from randomly distributed to microcanonically equivalent as interactions are increased. While we in principle have complete freedom to choose such an observable, it is important to step back before proceeding and ask what properties we wish this observable to satisfy. First, in order to be able to accumulate statistics over the entire middle third of the spectrum, we need the observable to be "microcanonically" the same over this wide range in energy. While a priori it seems unnecessary to look over such a large energy window, other studies which looked over much smaller energy windows to avoid seeing the energy-dependence of the observable (e.g., [121, 125]) had, as a result, very limited statistics on finite size scaling. Second, we wish to examine an observable that would be relatively straightforward to measure experimentally

and/or utilize in more complicated numerical techniques, such as DMRG or QMC methods. We will present some initial data for an observable which satisfies the first of these criteria but, as we will see, not the second. In this sense, this observable is a valuable tool for examining the finite-size crossover within the context of numerical exact diagonalization, but a more physical observable will need to be found before approaching the problem through more advanced methods or experiments.

The observable we consider (Fig. 7.2) is the number of broken momentum pairs,

$$\mathcal{B} = \sum_{k>0} \left(n_k + n_{-k} - 2 n_k n_{-k} \right) , \tag{7.13}$$

where $n_k = c_k^\dagger c_k$ is number of fermions at momentum k. This observable is suggested by the form of the unperturbed integrable Hamiltonian (Eq. 5.3), which commutes with $\mathcal{B}$ since all fermion operators occur in $(+k, -k)$ pairs. Therefore, the eigenstates with $\gamma = 0$ have $\mathcal{B}$ as a good quantum number (Fig. 7.2a). However, the scattering terms in Eq. 7.11 do not commute with $\mathcal{B}$, so as they are turned on, the eigenstate values of this observable mix. Eventually, for high enough γ, the spread of $\langle \mathcal{B} \rangle$ between eigenstates narrows and approaches a steady state (Fig. 7.2e), which has weak dependence on energy within the middle third of the spectrum (Fig. 7.2d). The spread of these $\mathcal{B}_{\alpha\alpha}$ values is characterized by their standard deviation

$$\sigma_{\mathcal{B}} = \sqrt{\sum_\alpha \left(\mathcal{B}_{\alpha\alpha} - \overline{\mathcal{B}_{\alpha\alpha}} \right)^2} , \tag{7.14}$$

where $\overline{\mathcal{B}_{\alpha\alpha}}$ is an average value over the α's. We see in Fig. 7.2f that $\sigma_{\mathcal{B}}$ has strong L and γ dependence, as expected.

To heuristically identify a crossover scale $\gamma_c(L)$, we take $\sigma_{\mathcal{B}}$ and do the following: 1) subtract the average value between $\gamma = 0.8$ and 1.0 to remove the offset, and 2) locate the value of γ where $\sigma_{\mathcal{B}}$ is at half of it's $\gamma = 0$ value by fitting to the four

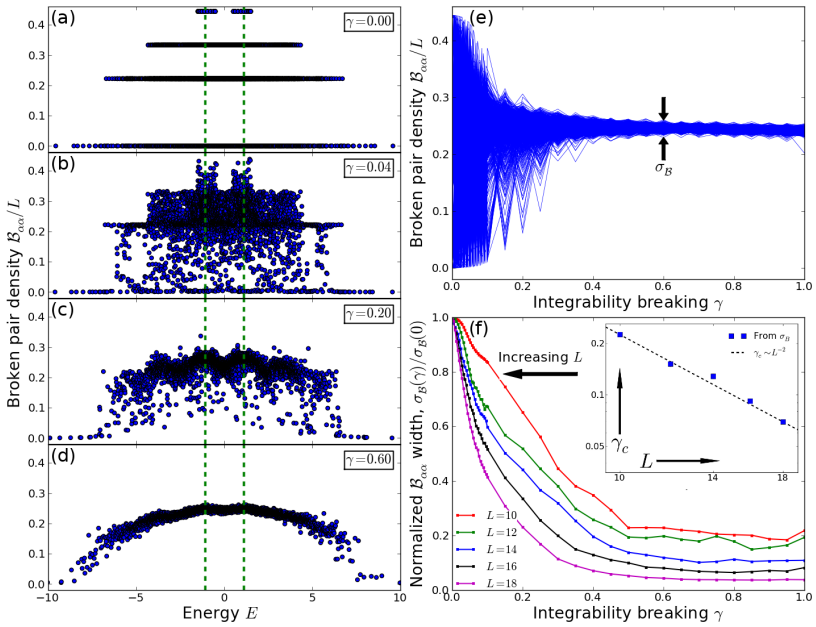

Figure 7.2: (a-d) Expectation values $\mathcal{B}_{\alpha\alpha}$ of the broken pair operator $\mathcal{B}$ (Eq. 7.13) in the energy eigenstates $|\alpha\rangle$, showing a crossover from integrable behavior ($\gamma = 0$) to eigenstate thermalization ($\gamma = 0.6$). Dashed green lines outline the middle third of the spectrum, used to find infinite temperature behavior in (e-f). (e) Evolution of the $\mathcal{B}_{\alpha\alpha}$ distribution as a function of γ. (f) Width of the $\mathcal{B}_{\alpha\alpha}$ distribution as a function of system size L and γ, normalized by the width at $\gamma = 0$. (inset) Heuristic estimate of the crossover scale from $\sigma_{\mathcal{B}}$, as described in the text. Dashed line shows our calculation for the breakdown of perturbation theory, $\gamma_c \sim L^{-2}$. As in Fig. 7.1, we use $\lambda = \frac{1}{2}$ and $P = 0$.

nearest points. This point, γ_c, is our estimate of the crossover scale. As seen in the inset to Fig. 7.2f, γ_c decreases strongly as a function of L. While consistent with the hypothesis that the phase transition occurs at infinitesimal integrability breaking in the TDL (i.e., $\gamma_c \to 0$ as $L \to \infty$), the range of system sizes we are able to solve is still not sufficient to draw any definitive conclusions.

Finally, we can place some limits on the finite-size scaling properties of eigenstate thermalization in this model by noting that, in order for the integrable Hamiltonian's eigenstates to mix sufficiently for their observables to satisfy the ETH, perturbation theory must break down.[5] If we consider an arbitrary eigenstate $|\mathbf{n}\rangle$ of the unper-

[5] After completing this analysis, we became aware of similar ideas in Ref. 125.

turbed TFI Hamiltonian, then at second order in perturbation theory, the energy of $|\mathbf{n}\rangle$ shifts by

$$\Delta E_{\mathbf{n}}^{(2)} = \sum_{\mathbf{m}\neq\mathbf{n}} \frac{|\langle\mathbf{m}|V|\mathbf{n}\rangle|^2}{E_{\mathbf{m}} - E_{\mathbf{n}}} \ . \tag{7.15}$$

Perturbation theory breaks down when the characteristic size of this term approaches the level spacing $E_{\mathbf{m}} - E_{\mathbf{n}}$ between coupled levels, i.e., levels where $\langle\mathbf{m}|V|\mathbf{n}\rangle \neq 0$.

To determine a naive bound on the breakdown of perturbation theory for the Hamiltonian in Eq. 7.10, we note that for each mode $k_m = 2\pi(m+1/2)/L$, there exists some Bogoliubov rotation such that the eigenstates of the unperturbed Hamiltonian are just number eigenstates of the Bogoliubov modes [78]. So $|\mathbf{n}\rangle$ can be represented as a binary string of length L, i.e., $|00111011\cdots\rangle$, where a 1 (0) at position m indicates that mode m is (not) occupied. The zeros and ones will be completely random in some arbitrary infinite temperature eigenstate.

Denote the Bogoliubov creation operators $a_k^\dagger = u_k c_k^\dagger + v_k c_{-k}$, where u_k and v_k are elements of the unitary Bogoliubov rotation matrix. In terms of these operators, the off-diagonal elements of the interaction matrix V do not conserve particle number, i.e., terms of the form $aaaa$, $a^\dagger aaa$, ..., $a^\dagger a^\dagger a^\dagger a^\dagger$ all exist. These terms have two things in common: 1) they enter with magnitudes proportional to γ/L and 2) they conserve the total momentum, modulo 2π. In counting up the total number of non-zero matrix elements of V, we are free to choose three different momenta at random, with the fourth fixed by momentum conservation. Therefore, the total number of terms in the sum in Eq. 7.15 scales as L^3, as there are L different momenta to choose from. By the first condition, these terms all enter with magnitude $|\langle\mathbf{m}|V|\mathbf{n}\rangle|^2 \sim \gamma^2/L^2$. Meanwhile, as the unperturbed levels come from a band with bandwidth of order unity, the characteristic scale of energy difference in the denominator of Eq. 7.15 will be ~ 1. This leads to an effective density of connected states $\rho_{\text{conn}} \sim L^3$, corresponding to a characteristic gap of $\Delta_{\text{conn}} \sim L^{-3}$ between connected levels. So we expect a scale γ_c

for the breakdown of perturbation theory given by

$$\left(\frac{\gamma_c}{L}\right)^2 L^3 \sim \Delta_{\text{conn}} \implies \gamma_c \sim L^{-2} \ . \tag{7.16}$$

While a bit heuristic, the above argument sets a limit on the fastest a system could thermalize as the non-integrable interaction $\gamma s^x s^x$ is introduced. For example, it appears to rule out the possibility of exponential approach to eigenstate thermalization for large systems, i.e., $\gamma_c \not\sim e^{-L}$. However, we are not yet able to show that this breakdown of perturbation theory is sufficient for eigenstate thermalization, so we have not proven that $\gamma_c \to 0$ as $L \to \infty$. While the strong decrease in observable variance as a function of system size shown in the inset to Fig. 7.2f provides evidence in favor of this possibility, the numerics have not yet been pushed to sufficiently large systems to establish that γ_c does not asymptote to a finite value in the TDL.

7.3 Conclusions

In this chapter, we investigated the finite size scaling of eigenstate thermalization as integrability is broken, using a variant of the Ising model as a numerical test case. Compared to the other chapters in this dissertation, this chapter contains relatively raw research, and therefore we are not able to draw many definitive conclusions at this time. We showed some preliminary evidence that, for a fairly aphysical observable $\mathcal{B}$, the crossover scale for eigenstate thermalization decreases sharply with system size. From a perturbative calculation, we set an upper limit on the rate at which the crossover scale can approach zero as the system size increases. However, we are not yet able to determine the asymptotic behavior of this crossover in the thermodynamic limit. To do so in a definitive manner, we will likely need a combination of more sophisticated numerics – such as MPS-based techniques – to allow investigation of larger system sizes, a more predictive theoretical approach – in particular one which

suggests some finite size scaling form which we can attempt to match to the data – and a more physically relevant observable to measure, such as the spin-spin correlation function. We have been working on all of these fronts, so hopefully soon we will be able to reach a more concrete conclusion.

Chapter 8

Open questions in closed systems

In this dissertation, we have discussed a number of topics in closed quantum systems. We have developed and analyzed the FCIQMC algorithm, both understanding and improving its ability to mitigate the sign problem, which is crucial to solving problems in fermionic systems. Turning to dynamics, we have solved the scaling properties for one of the simplest quantum critical systems, the integrable transverse-field Ising chain. We saw that qualitatively athermal spin correlations that arise in the dynamics are not only universal, but robust against experimental considerations such as boundary conditions and disorder. Finally, we turned to the long time behavior of the TFI chain with an integrability-breaking term added, presenting some initial data on the finite size scaling of eigenstate thermalization, a crucial ingredient for utilizing these closed systems to simulate condensed matter systems in thermal equilibrium.

This work, however, only begins to touch on the many open questions in the field of isolated quantum systems. We have seen that even in some of the simplest cases, such as the Kibble-Zurek scaling theory of the transverse-field Ising model, one finds interesting, unexpected, and universal non-equilibrium behavior. However, breaking the integrability renders the problem much more difficult, and it is in these non-integrable systems that the majority of open questions reside.

Within the context of scaling theory, many questions about the dynamics remain. The TFI chain discussed in Chp. 5 is one of the simplest examples of a critical point available: an integrable system with an integrable critical point. We saw numerically that breaking the integrability did not change the scaling behavior, but the MITP model still has an integrable QCP, meaning that the low energy properties remain integrable. One important goal is to find scaling functions for a non-integrable QCP, where even the low energy degrees of freedom are strongly correlated. In this case, one may address the interplay between eigenstate thermalization, which occurs in the long time limit, and the relatively short time behavior of critical scaling dynamics. As a precursor, one could instead examine the dynamics during a ramp of an RG-relevant non-integrable perturbation near an integrable QCP; for instance, adding a longitudinal magnetic field to the TFI chain.

Another class of models with interesting critical scaling properties are those with $z = \infty$ QCPs [108, 109], notably the infinite randomness critical point of the Ising and Heisenberg models [107, 127], as well as disorder-free models exhibiting a Berezinsky-Kosterlitz-Thouless (BKT) phase transition such as the quantum spin-1/2 XXZ model [128] and the spin-1 XXZ model with uniaxial single-spin anisotropy [129]. The model of Ref. 129 has particular applicability as a low-energy theory of the one-dimensional Bose-Hubbard model, which has much closer connection to current experiments than these spin models [35, 37]. In addition, with these one-dimensional bosonic systems there is some question as to the existence of anomalous dynamic exponents [104, 129], which may be addressable within a full BKT scaling theory.

At higher energies, our understanding of thermalization remains incomplete. For instance, most work on thermalization has focused on one-dimensional systems for the simple reason that they are the most numerically tractable. However, in higher dimensions, cold gases empirically thermalize well enough for mapping out equations

of state [130]. This dimensional dependence is still not fully understood, nor is, for example, the effect of an optical lattice [131, 132].

Experimentally, the controllable onset of thermalization remains to be seen. Since the first observation of athermal behavior in an integrable system [35], there have been numerous attempts to controllably add integrability-breaking terms to the Hamiltonian in order to measure the onset of thermalization. However, as of yet the time scales necessary to see thermalization in these systems do not appear to be accessible; the systems heat up due to environmental interactions before characteristics of eigenstate thermalization manifest [36, 37].

In order to theoretically address these questions, new numerical techniques will be needed. DMRG-based methods have been quite successful in performing time evolution and finding ground states of one-dimensional systems. However, their applicability for finding highly excited states is currently quite limited [133]. In addition, while time evolution works well for one-dimensional systems, it does not scale well to higher dimensions. One option for using DMRG-like techniques in higher dimensions are tensor product states [134, 135], but at this point that technology is still at too early a stage to be competitive with DMRG in performing time evolution.

Time evolution and excited state targeting with QMC suffer from a sign problem similar to that of the ground state of fermionic Hamiltonians, even for Hamiltonians where the ground state algorithm has no sign problem. This can be understood physically as the fact that, for most bases, excited and/or time-evolved states have nodes and phases, even when the ground state does not. For example, if we want to do real-time evolution via applying the operator $U = e^{-i\tau H} \approx 1 - i\tau H$, we can break U up into four positive matrices $\{\mathbf{U}^{\pm}, \mathbf{U}^{\pm I}\}$, such that

$$U = \mathbf{U}^{+} + i\,\mathbf{U}^{+I} - \mathbf{U}^{-} - i\,\mathbf{U}^{-I}\ . \tag{8.1}$$

Then, working in a configuration space consisting of a basis element and a phase (i.e., $\{\pm 1, \pm i\}$), the time evolution stochastically simulates

$$\begin{pmatrix} \mathbf{c}_{n+1}^{+} \\ \mathbf{c}_{n+1}^{+I} \\ \mathbf{c}_{n+1}^{-} \\ \mathbf{c}_{n+1}^{-I} \end{pmatrix} = \begin{pmatrix} \mathbf{U}^{+} & \mathbf{U}^{+I} & \mathbf{U}^{-} & \mathbf{U}^{-I} \\ \mathbf{U}^{-I} & \mathbf{U}^{+} & \mathbf{U}^{+I} & \mathbf{U}^{-} \\ \mathbf{U}^{-} & \mathbf{U}^{-I} & \mathbf{U}^{+} & \mathbf{U}^{+I} \\ \mathbf{U}^{+I} & \mathbf{U}^{-} & \mathbf{U}^{-I} & \mathbf{U}^{+} \end{pmatrix} \begin{pmatrix} \mathbf{c}_{n}^{+} \\ \mathbf{c}_{n}^{+I} \\ \mathbf{c}_{n}^{-} \\ \mathbf{c}_{n}^{-I} \end{pmatrix} \tag{8.2}$$

As in the analysis of Sec. 4.1.2, the dominant instability here is to the ground state of $\mathbf{U}^{+} + \mathbf{U}^{+I} + \mathbf{U}^{-} + \mathbf{U}^{-I}$. An interesting open question is how this "phase problem" is affected by the presence of annihilation. Based on our experience with FCIQMC, it seems as if solving such a problem would be difficult, but not necessarily impossible.

In conclusion, the dynamics of closed quantum systems remain a difficult but exciting frontier, becoming increasing relevant due to continued improvements in ultracold atom experiments. Indeed, it is often simpler to see non-equilibrium effects on experimental time scales than it is to find the ground state, so as theorists we must continue working to turn this into an advantage rather than a burden. The difficulty lies not just in numerically or analytically solving the Schrödinger equation, but in attempting to place results within a broader unifying context; as a fairly new field, research on closed system dynamics still seems relatively disjointed. While in this dissertation we have examined some particular limits of the dynamics – namely their long time behavior and critical scaling – the majority of phase space remains untouched. We hope and expect that, as this area is explored, it will continue to yield exciting and unexpected new phenomena.

Appendix A

Proof of quantization equivalence conditions for FCIQMC

We will now show that a) the conditions described in Sec. 4.1.3 uniquely characterize when the first and second quantized sign problem differ and b) if they differ, the second quantized sign problem is better. We introduce the notation

$$\widetilde{\mathbf{U}}_\alpha = \mathbf{U}_\alpha^+ + \mathbf{U}_\alpha^- \ , \tag{A.1}$$

where $\alpha \in \{F, D, P\}$, $\mathbf{U}_\alpha^{+(-)}$ consists of strictly non-negative elements, and the actual matrix corresponding to U is $\mathbf{U}_\alpha = \mathbf{U}_\alpha^+ - \mathbf{U}_\alpha^-$.

Let's start by considering under what conditions the first and second quantized algorithms have the same sign problem. From the arguments in Sec. 4.1.3, it is clear this will be the case if

$$\langle \mathbf{P_i} | \left(\mathbf{U}_F^+ + \mathbf{U}_F^- \right) | \mathbf{P_j} \rangle = \langle \mathbf{D_i} | \left(\mathbf{U}_D^+ + \mathbf{U}_D^- \right) | \mathbf{D_j} \rangle \ . \tag{A.2}$$

Expanding out the determinants and permanents, this can be rewritten

$$\frac{1}{N!}\sum_{\substack{\mathbf{i}\in\mathbf{D_i}\\ \mathbf{j}\in\mathbf{D_j}}}|\langle\mathbf{i}|U|\mathbf{j}\rangle| = \left|\sum_{\substack{\mathbf{i}\in\mathbf{D_i}\\ \mathbf{j}\in\mathbf{D_j}}}\langle\mathbf{i}|\mathbf{D_i}\rangle\langle\mathbf{j}|\mathbf{D_j}\rangle\langle\mathbf{i}|U|\mathbf{j}\rangle\right| . \tag{A.3}$$

By the triangle inequality, the left side is greater than or equal to the right side, with equality iff all terms within the absolute values on the right side are of the same sign. Therefore, the equality implies that, for all $\mathbf{i}\in\mathbf{D_i}$ and $\mathbf{j}\in\mathbf{D_j}$,

$$\begin{aligned} s_{\mathbf{D_iD_j}}\,|\langle\mathbf{i}|U|\mathbf{j}\rangle| &= \mathrm{sgn}\langle\mathbf{i}|\mathbf{D_i}\rangle\mathrm{sgn}\langle\mathbf{j}|\mathbf{D}_j\rangle\langle\mathbf{i}|U|\mathbf{j}\rangle \iff \\ s_{\mathbf{D_iD_j}} &= \mathrm{sgn}\langle\mathbf{i}|\mathbf{D_i}\rangle\mathrm{sgn}\langle\mathbf{j}|\mathbf{D_j}\rangle\mathrm{sgn}\langle\mathbf{i}|U|\mathbf{j}\rangle \\ &\text{or} \\ \langle\mathbf{i}|U|\mathbf{j}\rangle &= 0\ , \end{aligned} \tag{A.4}$$

where $s_{\mathbf{D}_i\mathbf{D}_j}=\pm1$ is an arbitrary sign that is the same for all $\mathbf{i}\in\mathbf{D_i}$ and $\mathbf{j}\in\mathbf{D_j}$. If we consider two distinct FQ basis elements $\mathbf{j},\mathbf{j}'\in\mathbf{D_j}$, then either

- $\langle\mathbf{i}|U|\mathbf{j}\rangle = 0$ or
- $\langle\mathbf{i}|U|\mathbf{j}'\rangle = 0$ or
- $\mathrm{sgn}\langle\mathbf{j}|\mathbf{D_j}\rangle\mathrm{sgn}\langle\mathbf{i}|U|\mathbf{j}\rangle = \mathrm{sgn}\langle\mathbf{j}'|\mathbf{D_j}\rangle\mathrm{sgn}\langle\mathbf{i}|U|\mathbf{j}'\rangle$. (A.5)

These conditions, which we will see uniquely characterize when first and second quantization differ, are precisely the conditions we assumed were met by the "filtered" psip distribution in Sec. 4.1.3. Since the conditions in Eq. A.5 are equivalent to those in Eq. A.2, then *if Eq. A.5 holds, the first and second quantized sign problems are identical.*

Second, we show that *if the conditions in Eq. A.5 are not met, then the sign problem in second quantization is strictly better.* Let $|\widetilde{\varphi}_D\rangle$ be the highest energy

eigenstate of $\widetilde{\mathbf{U}}_D$ with eigenvalue $\widetilde{E}_D$. Construct the (normalized) FQ state $|\varphi_F\rangle$ such that

$$\langle \mathbf{i}|\varphi_F\rangle = \frac{1}{\sqrt{N!}}\langle \mathbf{D_i}|\widetilde{\varphi}_D\rangle \ , \tag{A.6}$$

for all $|\mathbf{i}\rangle \in |\mathbf{D_i}\rangle$. Since $\widetilde{\mathbf{U}}_D$ is a positive matrix, $|\widetilde{\varphi}_D\rangle$, and thus $|\varphi_F\rangle$, is positive valued.

Now consider the variational energy of $|\varphi_F\rangle$ with respect to the matrix $\widetilde{\mathbf{U}}_F$:

$$\widetilde{E}'_F \equiv \langle \varphi_F|\widetilde{\mathbf{U}}_F|\varphi_F\rangle = \sum_{\mathbf{ij}} \langle \mathbf{i}|\varphi_F\rangle\langle \mathbf{j}|\varphi_F\rangle\langle \mathbf{i}|\widetilde{\mathbf{U}}_F|\mathbf{j}\rangle \ . \tag{A.7}$$

By definition, $\langle \mathbf{i}|\widetilde{\mathbf{U}}_F|\mathbf{j}\rangle = |\langle \mathbf{i}|U|\mathbf{j}\rangle|$. Utilizing this,

$$\begin{aligned}\widetilde{E}'_F &= \frac{1}{N!}\sum_{\mathbf{ij}} \langle \mathbf{D_i}|\widetilde{\varphi}_D\rangle\langle \mathbf{D_j}|\widetilde{\varphi}_D\rangle \, |\langle \mathbf{i}|U|\mathbf{j}\rangle| \\ &= \frac{1}{N!}\sum_{\mathbf{D_i D_j}} \langle \mathbf{D_i}|\widetilde{\varphi}_D\rangle\langle \mathbf{D_j}|\widetilde{\varphi}_D\rangle \sum_{\mathbf{i}\in\mathbf{D_i}}\sum_{\mathbf{j}\in\mathbf{D_j}} |\langle \mathbf{i}|U|\mathbf{j}\rangle| \, ,\end{aligned} \tag{A.8}$$

where in the last step we broke the sum up into determinant sectors.

From this point onward, assume that the conditions in Eq. A.5 are not met. Then there must exist some $\mathbf{j}, \mathbf{j}' \in \mathbf{D}_j$ such that

- $\langle \mathbf{i}|U|\mathbf{j}\rangle \neq 0$ and
- $\langle \mathbf{i}|U|\mathbf{j}'\rangle \neq 0$ and
- $\mathrm{sgn}\langle \mathbf{j}|\mathbf{D_j}\rangle \mathrm{sgn}\langle \mathbf{i}|U|\mathbf{j}\rangle \neq \mathrm{sgn}\langle \mathbf{j}'|\mathbf{D_j}\rangle \mathrm{sgn}\langle \mathbf{i}|U|\mathbf{j}'\rangle$. (A.9)

Consider the sign of these two terms in the sum $\sum_{\mathbf{j}\in\mathbf{D_j}}$ from Eq. A.8. It is clear that

$$|\langle \mathbf{i}|U|\mathbf{j}\rangle| = |\langle \mathbf{i}|U|\mathbf{j}\rangle \mathrm{sgn}\langle \mathbf{j}|\mathbf{D_j}\rangle| \tag{A.10}$$

for all $\mathbf{j} \in \mathbf{D_j}$. However, by Eq. A.9, $\mathbf{j}$ and $\mathbf{j}'$ have different signs for the term inside the absolute value on the right hand side of Eq. A.10. Therefore, by the triangle inequality,

$$\widetilde{E}'_F > \frac{1}{N!} \sum_{\mathbf{D_i} \mathbf{D_j}} \langle \mathbf{D_i} | \widetilde{\varphi}_D \rangle \langle \mathbf{D_j} | \widetilde{\varphi}_D \rangle \sum_{\mathbf{i} \in \mathbf{D_i}} \left| \sum_{\mathbf{j} \in \mathbf{D_j}} \langle \mathbf{i} | U | \mathbf{j} \rangle \mathrm{sgn} \langle \mathbf{j} | \mathbf{D_j} \rangle \right| \ . \tag{A.11}$$

We are now in a position to rewrite the right hand side in terms of determinants, since

$$\sum_{\mathbf{j} \in \mathbf{D_j}} |\mathbf{j}\rangle \mathrm{sgn} \langle \mathbf{j} | \mathbf{D_j} \rangle = \sqrt{N!} |\mathbf{D_j}\rangle \ . \tag{A.12}$$

Thus,

$$\begin{aligned} \sum_{\mathbf{j} \in \mathbf{D_j}} \langle \mathbf{i} | U | \mathbf{j} \rangle \mathrm{sgn} \langle \mathbf{j} | \mathbf{D_j} \rangle &= \sqrt{N!} \langle \mathbf{i} | U | \mathbf{D_j} \rangle \\ &= \sqrt{N!} \langle \mathbf{i} | U \mathcal{A} | \mathbf{D_j} \rangle \\ &= \sqrt{N!} \langle \mathbf{i} | \mathcal{A} U | \mathbf{D_j} \rangle \\ &= \mathrm{sgn} \langle \mathbf{i} | \mathbf{D_i} \rangle \langle \mathbf{D_i} | U | \mathbf{D_j} \rangle \ . \end{aligned} \tag{A.13}$$

Now that everything is in terms of $\mathbf{D_i}$, we can bring out a factor of $N!$ corresponding to the sum $\sum_{\mathbf{i} \in \mathbf{D_i}}$. Then

$$\begin{aligned} \widetilde{E}'_F &> \sum_{\mathbf{D_i} \mathbf{D_j}} \langle \mathbf{D_i} | \widetilde{\varphi}_D \rangle \langle \mathbf{D_j} | \widetilde{\varphi}_D \rangle \left| \langle \mathbf{D_i} | U | \mathbf{D_j} \rangle \right| \\ &= \sum_{\mathbf{D_i} \mathbf{D_j}} \langle \mathbf{D_i} | \widetilde{\varphi}_D \rangle \langle \mathbf{D_j} | \widetilde{\varphi}_D \rangle \langle \mathbf{D_i} | \widetilde{\mathbf{U}}_D | \mathbf{D_j} \rangle = \widetilde{E}_D \ . \end{aligned} \tag{A.14}$$

Therefore, by the variational principle, $\widetilde{E}_F > \widetilde{E}_D$, where $\widetilde{E}_F$ is the maximum eigenvalue of $\widetilde{H}_F$.

In conclusion, we have shown that $\widetilde{E}_F = \widetilde{E}_D$ if and only if the conditions in Eq. A.5 hold. If they do not hold, then $\widetilde{E}_F > \widetilde{E}_D$. Therefore, the sign problem in a second

quantized formalism is at least as good as in the first quantized case with same set of single-particle orbitals, becoming strictly better precisely when the conditions in Eq. A.5 break.

Bibliography

[1] A. Schirotzek, C.-H. Wu, A. Sommer, and M. W. Zwierlein. *Phys. Rev. Lett.*, 102:230402, 2009.

[2] J. Simon, W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner. *Nature*, 472:307, 2011.

[3] J. S. Spencer, M. Kolodrubetz, B. K. Clark, and W. M. Foulkes. In preparation.

[4] M. Kolodrubetz and B. K. Clark. arXiv:1204.1490, to be published in Phys. Rev. B, 2012.

[5] M. Kolodrubetz, D. Pekker, B. K. Clark, and K. Sengupta. *Phys. Rev. B*, 85:100505, 2012.

[6] M. Kolodrubetz, B. K. Clark, and D. A. Huse. *Phys. Rev. Lett.*, 109:015701, 2012.

[7] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. *Phys. Rev. Lett.*, 75:1687, 1995.

[8] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. *Phys. Rev. Lett.*, 75:3969, 1995.

[9] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. *Science*, 269:198, 1995.

[10] W. Ketterle and M. W. Zwierlein. In M. Inguscio, W. Ketterle, and C. Salomon, editors, *Ultra-cold Fermi Gases, Proceedings of the International School of Physics Enrico Fermi, Course CLXIV*, pages 95–287. IOS Press, Amsterdam, 2007.

[11] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch. *Nature*, 415:39, 2002.

[12] Y.-J. Lin, R. L. Compton, K. Jimenez-Garcia, J. V. Porto, and I. B. Spielman. *Nature*, 462:628, 2009.

[13] J. Zhang, F. M. Cucchietti, C. M. Chandrashekar, M. Laforest, C. A. Ryan, M. Ditty, A. Hubbard, J. K. Gamble, and R. Laflamme. *Phys. Rev. A*, 79:012305, 2009.

[14] K. Kim, S. Korenblit, R. Islam, E. E. Edwards, M.-S. Chang, C. Noh, H. Carmichael, G.-D. Lin, L.-M. Duan, C. C. Joseph Wang, J. K. Freericks, and C. Monroe. *New Journal of Physics*, 13:105003, 2011.

[15] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. In M. Inguscio, S. Stringari, and C. Wieman, editors, *Bose-Einstein Condensation in Atomic Gases, Proceedings of the International School of Physics Enrico Fermi, Varenna, 7-17 July 1998, Course CXL*, pages 67–176. IOS Press, Amsterdam, 1999.

[16] L. Dunoyer. *J. du Radium*, 8, 1911.

[17] J. V. Prodan, W. D. Phillips, and H. Metcalf. *Phys. Rev. Lett.*, 49:1149, 1982.

[18] R. Grimm, M. Weidemüller, and Y. Ovchinnikov. *Advances in Atomic, Molecular, and Optical Physics*, 42:95–170, 2000.

[19] M. D. Swallows, M. Bishof, Y. Lin, S. Blatt, M. J. Martin, A. M. Rey, and J. Ye. *Science*, 331:1043, 2011.

[20] S. Ospelkaus, K.-K. Ni, D. Wang, M. H. G. de Miranda, B. Neyenhuis, G. Quemner, P. S. Julienne, J. L. Bohn, D. S. Jin, and J. Ye. *Science*, 327:853, 2010.

[21] A. V. Gorshkov, A. M. Rey, A. J. Daley, M. M. Boyd, J. Ye, P. Zoller, and M. D. Lukin. *Phys. Rev. Lett.*, 102:110503, 2009.

[22] W. S. Bakr, J. I. Gillen, A. Peng, S. Folling, and M. Greiner. *Nature*, 462:74, 2009.

[23] L. Landau. *Physics of the Soviet Union*, 2:46–51, 1932.

[24] C. Zener. *Proceedings of the Royal Society of London A*, 137:696–702, 1932.

[25] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga. *Rev. Mod. Phys.*, 82:1225, 2010.

[26] H. Feshbach. *Annals of Physics*, 5:357, 1958.

[27] M. Greiner, C. A. Regal, and D. S. Jin. *Nature*, 426:537, 2003.

[28] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, A. J. Kerman, and W. Ketterle. *Phys. Rev. Lett.*, 92:120403, 2004.

[29] M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, C. Chin, J. H. Denschlag, and R. Grimm. *Phys. Rev. Lett.*, 92:120401, 2004.

[30] M. W. Zwierlein, A. Schirotzek, C. H. Schunck, and W. Ketterle. *Science*, 311:492, 2006.

[31] G. B. Partridge, W. Li, Y. A. Liao, R. G. Hulet, M. Haque, and H. T. C. Stoof. *Phys. Rev. Lett.*, 97:190407, 2006.

[32] Y.-i. Shin, C. H. Schunck, A. Schirotzek, and W. Ketterle. *Nature*, 451:689, 2008.

[33] K. Van Houcke, F. Werner, E. Kozik, N. Prokof'ev, B. Svistunov, M. J. H. Ku, A. T. Sommer, L. W. Cheuk, A. Schirotzek, and M. W. Zwierlein. *Nat. Phys.*, 8:366, 2012.

[34] S. Trotzky, P. Cheinet, S. Fölling, M. Feld, U. Schnorrberger, A. M. Rey, A. Polkovnikov, E. A. Demler, M. D. Lukin, and I. Bloch. *Science*, 319:295, 2008.

[35] T. Kinoshita, T. Wenger, and D. S. Weiss. *Nature*, 440:900, 2006.

[36] J. Schmiedmayer. Private communication.

[37] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D. A. Smith, E. Demler, and J. Schmiedmayer. arXiv:1112.0013, 2011.

[38] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 2007.

[39] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.

[40] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. *J. Chem. Phys.*, 21:1087–1092, 1953.

[41] S. Sorella. *Phys. Rev. B*, 71:241103, 2005.

[42] J. Toulouse and C. J. Umrigar. *J. Chem. Phys.*, 126:084102, 2007.

[43] G. H. Booth, A. J. W. Thom, and A. Alavi. *J. Chem. Phys.*, 131:054106, 2009.

[44] S. R. White. *Phys. Rev. Lett.*, 69:2863, 1992.

[45] S. R. White. *Phys. Rev. B*, 48:10345, 1993.

[46] K. G. Wilson. *Rev. Mod. Phys.*, 47:773, 1975.

[47] U. Schollwöck. *Annals of Physics*, 326:96, 2011.

[48] U. Schollwöck. *Rev. Mod. Phys.*, 77:259, 2005.

[49] K. A. Hallberg. *Advances in Physics*, 55:477, 2006.

[50] M. Q. Weng, D. N. Sheng, Z. Y. Weng, and R. J. Bursill. *Phys. Rev. B*, 74:012407, 2006.

[51] S. Yan, D. A. Huse, and S. R. White. *Science*, 332:1173, 2011.

[52] P. Pippan, S. R. White, and H. G. Evertz. *Phys. Rev. B*, 81:081103, 2010.

[53] H. Li and F. D. M. Haldane. *Phys. Rev. Lett.*, 101:010504, 2008.

[54] G. Vidal. *Phys. Rev. Lett.*, 91:147902, 2003.

[55] G. Vidal. *Phys. Rev. Lett.*, 93:040502, 2004.

[56] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal. *Journal of Statistical Mechanics: Theory and Experiment*, 2004:P04005, 2004.

[57] S. R. White and A. E. Feiguin. *Phys. Rev. Lett.*, 93:076401, 2004.

[58] F. Verstraete, J. J. Garcia-Ripoll, and J. I. Cirac. *Phys. Rev. Lett.*, 93:207204, 2004.

[59] N. Hatano and M. Suzuki. In A. Das and B.K. Chakrabarti, editors, *Quantum Annealing and Other Optimization Methods*, pages 37–68. Springer, 2005.

[60] H. Takasaki, T. Hikihara, and T. Nishino. *J. Phys. Soc. Jpn.*, 68:1537, 1999.

[61] S. R. White. *Phys. Rev. B*, 72:180403, 2005.

[62] R. C. Grimm and R. G. Storer. *J. Comp. Phys.*, 7:134, 1971.

[63] J. S. Spencer, N. S. Blunt, and W. M.C. Foulkes. *J. Chem. Phys.*, 136:054110, 2012.

[64] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi. *Phys. Rev. B*, 85:081103, 2012.

[65] M. H. Kalos and F. Pederiva. *Phys. Rev. Lett.*, 85:3547, 2000.

[66] D. Cleland, G. H. Booth, and A. Alavi. *J. Chem. Phys.*, 132:041103, 2010.

[67] R. Assaraf, M. Caffarel, and A. Khelif. *J. Phys. A: Math. and Theor.*, 40:1181, 2007.

[68] A. Szabo and N. S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Publications, 1996.

[69] A. S. Mishchenko, N. V. Prokofev, A. Sakamoto, and B. V. Svistunov. *Phys. Rev. B*, 62:6317, 2000.

[70] H. Fröhlich, H. Pelzer, and S. Zienau. *Philos. Mag.*, 41:221, 1950.

[71] F. Chevy. *Phys. Rev. A*, 74:063628, 2006.

[72] C. J. M. Mathy, M. M. Parish, and D. A. Huse. *Phys. Rev. Lett.*, 106:166404, 2011.

[73] R. Combescot and S. Giraud. *Phys. Rev. Lett.*, 101:050404, 2008.

[74] N. V. Prokof'ev and B. V. Svistunov. *Phys. Rev. B*, 77:125101, 2008.

[75] D. M. Ceperley and M. H. Kalos. In K. Binder, editor, *Monte Carlo Methods in Statistical Physics*. Springer-Verlag, 1979.

[76] H. J. M. van Bemmel, D. F. B. ten Haaf, W. van Saarloos, J. M. J. van Leeuwen, and G. An. *Phys. Rev. Lett.*, 72:2442, 1994.

[77] D. M. Ceperley and B. J. Adler. *J. Chem. Phys.*, 81:5833, 1984.

[78] S. Sachdev. *Quantum Phase Transitions*. Cambridge University Press, 1999.

[79] L. Onsager. *Phys. Rev.*, 65:117, 1944.

[80] T. W. B. Kibble. *J. Phys. A: Math. Gen.*, 9:1387, 1976.

[81] W. H. Zurek. *Nature*, 317:505, 1985.

[82] B. Damski. *Phys. Rev. Lett.*, 95:035701, 2005.

[83] A. del Campo, G. De Chiara, G. Morigi, M. B. Plenio, and A. Retzker. *Phys. Rev. Lett.*, 105:075701, 2010.

[84] C. De Grandi, A. Polkovnikov, and A. W. Sandvik. *Phys. Rev. B*, 84:224303, 2011.

[85] C. De Grandi, V. Gritsev, and A. Polkovnikov. *Phys. Rev. B*, 81:012303, 2010.

[86] A. Chandran, A. Erez, S. S. Gubser, and S. L. Sondhi. arXiv:1202.5277, 2012.

[87] S. Deng, G. Ortiz, and L. Viola. *Europhys. Lett.*, 84:67008, 2008.

[88] S. Sachdev, K. Sengupta, and S. M. Girvin. *Phys. Rev. B*, 66:075128, 2002.

[89] J. Dziarmaga. *Phys. Rev. Lett.*, 95:245701, 2005.

[90] C. N. Yang. *Phys. Rev.*, 85:808, 1952.

[91] B. Damski and W. H. Zurek. *Phys. Rev. A*, 73:063405, 2006.

[92] M. M. Rams and B. Damski. *Phys. Rev. Lett.*, 106:055701, 2011.

[93] J. I. Latorre and R. Orus. *Phys. Rev. A*, 69:062302, 2004.

[94] R. Schützhold and G. Schaller. *Phys. Rev. A*, 74:060304, 2006.

[95] H.-Q. Zhou and J. P. Barjaktarevic. *J. Phys. A*, 41:412001, 2008.

[96] E. Barouch and B. M. McCoy. *Phys. Rev. A*, 3:786, 1971.

[97] T. T. Wu. *Phys. Rev.*, 149:380, 1966.

[98] A. Polkovnikov and V. Gritsev. *Nat. Phys.*, 4:477, 2008.

[99] L. Cincio, Ja. Dziarmaga, M. M. Rams, and W. H. Zurek. *Phys. Rev. A*, 75:052321, 2007.

[100] K. Sengupta, S. Powell, and S. Sachdev. *Phys. Rev. A*, 69:053616, 2004.

[101] R. W. Cherng and L. S. Levitov. *Phys. Rev. A*, 73:043614, 2006.

[102] M. Rigol, V. Dunjko, V. Yurovsky, and M. Olshanii. *Phys. Rev. Lett.*, 98:050405, 2007.

[103] F. Pollmann, S. Mukerjee, A. G. Green, and J. E. Moore. *Phys. Rev. E*, 81:020101, 2010.

[104] J.-S. Bernier, G. Roux, and C. Kollath. *Phys. Rev. Lett.*, 106:200601, 2011.

[105] M. Rigol, V. Dunjko, and M. Olshanii. *Nature*, 452:854, 2008.

[106] D. S. Fisher. *Phys. Rev. Lett.*, 69:534, 1992.

[107] D. S. Fisher. *Phys. Rev. B*, 51:6411, 1995.

[108] J. Dziarmaga. *Phys. Rev. B*, 74:064416, 2006.

[109] T. Caneva, R. Fazio, and G. E. Santoro. *Phys. Rev. B*, 76:144427, 2007.

[110] R. Ma. Private communication.

[111] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore. *Rev. Mod. Phys.*, 83:863, 2011.

[112] E. H. Lieb and W. Liniger. *Phys. Rev.*, 130:1605, 1963.

[113] E. H. Lieb. *Phys. Rev.*, 130:1616, 1963.

[114] A. Polkovnikov. *Annals of Physics*, 326:486, 2011.

[115] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghi. *Phys. Rev. Lett.*, 96:050403, 2006.

[116] S. Popescu, A. J. Short, and A. Winter. *Nat. Phys.*, 2:754, 2006.

[117] J. M. Deutsch. *Phys. Rev. A*, 43:2046, 1991.

[118] M. Srednicki. *Phys. Rev. E*, 50:888, 1994.

[119] D. M. Basko, I. L. Aleiner, and B. L. Altshuler. *Annals of Physics*, 321:1126, 2006.

[120] A. Pal and D. A. Huse. *Phys. Rev. B*, 82:174411, 2010.

[121] M. Rigol. *Phys. Rev. Lett.*, 103:100403, 2009.

[122] A. N. Kolmogorov. *Dokl. Akad. Nauk SSSR*, 98:527–530, 1954.

[123] V. I. Arnold. *Uspehi Mat. Nauk*, 18:13–40, 1963.

[124] J. Moser. *Nachr. Akad. Wiss. Göttingen Math.-Phys. Kl. II*, pages 1–20, 1962.

[125] C. Neuenhahn and F. Marquardt. arXiv:1007.5306, 2010.

[126] V. Oganesyan and D. A. Huse. *Phys. Rev. B*, 75:155111, 2007.

[127] Daniel S. Fisher. *Phys. Rev. B*, 50:3799, 1994.

[128] H. Bethe. *Zeitschrift fr Physik A Hadrons and Nuclei*, 71:205, 1931.

[129] E. Canovi, D. Rossini, R. Fazio, and G. E. Santoro. *J. Stat. Mech.*, page P03038, 2009.

[130] N. Navon, S. Nascimbéne, F. Chevy, and C. Salomon. *Science*, 328(5979):729, May 2010.

[131] S. S. Natu, K. R. A. Hazzard, and E. J. Mueller. *Phys. Rev. Lett.*, 106:125301, 2011.

[132] W. S. Bakr, P. M. Preiss, M. E. Tai, R. Ma, J. Simon, and M. Greiner. *Nature*, 480:500, 2011.

[133] J. J. Dorando, J. Hachmann, and G. K.-L. Chan. *J. Chem. Phys.*, 127:084109, 2007.

[134] F. Verstraete and J. I. Cirac. *Phys. Rev. A*, 70:060302, 2004.

[135] G. Vidal. *Phys. Rev. Lett.*, 99:220405, 2007.