The Interplay of Localization and Interactions in Quantum Many-Body Systems

Thesis by

Shankar Iyer

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

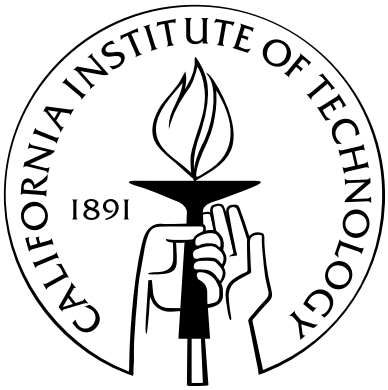

California Institute of Technology

Pasadena, California

2013

(Defended May 29, 2013)

For Sowmya, who inspired me to finish this.

Acknowledgements

My five years at Caltech have been among the most wonderful and transformative of my life. This would not have been possible without the kindness and guidance of my advisor, Gil Refael. Early on, he gave me the freedom to pursue the problems that interested me, and throughout my graduate studies, he remained a constant source of support, encouragement, and enthusiasm. I learned a lot of physics from him, but I also learned so much more: how to talk to other scientists, how to communicate my work, even how to organize my notebooks. His personal and scientific influence is all over this thesis.

Many others enriched my time in the Condensed Matter Group. I particularly want to recognize David Pekker for being such a patient guide through the intricacies of dirty bosons. I also want to thank Lesik Motrunich, Michael Cross, and Oskar Painter for serving on my thesis committee, and Loly Ekmekjian for all her help over the years. It was very enjoyable to work and learn alongside my fellow graduate students, among them Paraj Bhattacharjee, Scott Geraedts, Kun Woo Kim, Tony Lee, Shu-Ping Lee, Hsin-Hua Lai, Debaleena Nandi, and Karthik Seetharam. I cannot thank Liyan Qiao and Zhaorui Li enough for their friendship and the incredible honor of choosing Sowmya and me as the godparents of their daughter Amy. For two years,

Liyan, Zhaorui, and Amy transformed our office into a second home and made coming to work an absolute joy.

Outside of our research group, I am indebted to Hemanth Siriki for putting up with me for four years. I also want to thank Bill Fefferman, Karmina Caragan, Senthil Radhakrishnan, and several others for sharing weekend evenings with me. I loved getting to know Pasadena and the greater Los Angeles area with all of them. Finally, I deeply appreciate the hospitality that the Nadadur family showed me on so many occasions during these years. It was great to have such a welcoming home to retreat to on weekends when I needed an escape from Caltech.

Before beginning graduate school, I spent one year as a Fulbright Fellow at Leiden University in The Netherlands. I want to thank the Netherland-America Foundation, Jan Zaanen, and the Instituut-Lorentz for their roles in making that year possible. I also want to express my gratitude to the Kauffman-Abel family for friendship and hospitality during that time and during my later visits to Leiden.

From 2003-2007, I was an undergraduate at Princeton University. There, I began a senior thesis project that eventually expanded and evolved into Chapter 4 of this thesis. I thank my senior thesis advisor, David Huse, for introducing me to research in condensed matter theory and for helping me to revive our project when I met him again at the Boulder Summer School in 2010. I am also very grateful to Vadim Oganesyan and Arijeet Pal for their insight into the problem of many-body localization. During my Princeton years, I made many lifelong friends who still enrich my life: among them are Meha Jain, Yijia Lu, Anuj Nadadur, Ben Pollak, and Leizhi

Sun, but there are too many others to name.

My interest in physics developed during high school in East Brunswick, and I want to thank all the excellent teachers who nurtured my curiosity in physics and other subjects. Moreover, I cannot overestimate the influence of my classmates during this period. Had I not developed a passion for physics alongside them, I would not now be completing this thesis.

Before all of this, I had the great fortune of being born into a large, loving family. Whatever I have achieved would have been impossible without my family's love and support, especially in the trying times. Most of all, I want to express my deep gratitude to my sister (Jayasree Iyer), my parents (Latha and Mohan Iyer), and my grandparents (A. Chellammal, S. Padmanabha Iyer, N. Jayalakshmi, and H. Sankara Iyer). However, I also want to thank the rest of my extended family in New Jersey, India, and elsewhere around the world.

Finally, I want to thank Sowmya and her family for welcoming me into their lives over the past couple of years. My life has become immeasurably better for having Sowmya in it, and I am **so** excited for our future together...

Abstract

Disorder and interactions both play crucial roles in quantum transport. Decades ago, Mott showed that electron-electron interactions can lead to insulating behavior in materials that conventional band theory predicts to be conducting [105]. Soon thereafter, Anderson demonstrated that disorder can localize a quantum particle through the wave interference phenomenon of Anderson localization [11]. Although interactions and disorder both separately induce insulating behavior, the interplay of these two ingredients is subtle and often leads to surprising behavior at the periphery of our current understanding. Modern experiments probe these phenomena in a variety of contexts (e.g. disordered superconductors, cold atoms, photonic waveguides, etc.); thus, theoretical and numerical advancements are urgently needed. In this thesis, we report progress on understanding two contexts in which the interplay of disorder and interactions is especially important.

The first is the so-called "dirty" or random boson problem. In the past decade, a strong-disorder renormalization group (SDRG) treatment by Altman, Kafri, Polkovnikov, and Refael has raised the possibility of a new unstable fixed point governing the superfluid-insulator transition in the one-dimensional dirty boson problem [7, 8, 9]. This new critical behavior may take over from the weak-disorder criticality of Giamarchi and Schulz when disorder is sufficiently strong [58]. We analytically determine the scaling of the superfluid susceptibility at the strong-disorder fixed point and connect our analysis to recent Monte

Carlo simulations by Hrahsheh and Vojta [70]. We then shift our attention to two dimensions and use a numerical implementation of the SDRG to locate the fixed point governing the superfluid-insulator transition there. We identify several universal properties of this transition, which are fully independent of the microscopic features of the disorder.

The second focus of this thesis is the interplay of localization and interactions in systems with high energy density (i.e., far from the usual low energy limit of condensed matter physics). Recent theoretical and numerical work indicates that localization can survive in this regime, provided that interactions are sufficiently weak. Stronger interactions can destroy localization, leading to a so-called many-body localization transition [17, 109, 111]. This dynamical phase transition is relevant to questions of thermalization in isolated quantum systems: it separates a many-body localized phase, in which localization prevents transport and thermalization, from a conducting ("ergodic") phase in which the usual assumptions of quantum statistical mechanics hold. Here, we present evidence that many-body localization also occurs in quasiperiodic systems that lack true disorder.

Contents

Chapter 1

Introduction

1.1 Noninteracting Particles in Perfect Media

The most basic theories of many-body and solid-state physics treat noninteracting particles propagating in perfectly periodic environments (e.g., electrons in crystal lattices). This thesis concerns situations in which this picture fails completely. More precisely, it examines physical systems where particle-particle interactions and aperiodicity (e.g., disorder) combine to lead to novel phenomena. Before delving into the effects of these two ingredients, it is useful to revisit the simplified theories that ignore them and recall their successes and inadequacies[1].

1.1.1 Electrons in a Metal as an Ideal Fermi Gas

In 1900, Drude formulated the first model of conduction in metals by considering a classical gas of conduction electrons moving in a static background of positive ions. He assumed that electrons propagate freely between collisions with the ionic back-

[1]Readers who are familiar with the important roles that disorder and interactions individually play in many-body physics may wish to skip to 1.4.2, where we begin discussing the specific systems that we study in Chapters 2 and 3.

ground, that each electron's collisions are typically separated by a mean scattering time τ, and that the electrons emerge from these scattering events with velocities obeying the classical Maxwell-Boltzmann distribution. Using these assumptions, he was able to make predictions for certain quantities that roughly match experimental observations. Nevertheless, other measurements deviated strongly from Drude's predictions. Most notably, experimental values of the specific heat c_v did not show the classically expected value of $\frac{3}{2}k_B$ per conduction electron. This particular deficiency was remedied in the early years of quantum mechanics by Sommerfeld, who replaced the Maxwell-Boltzmann distribution in Drude theory with the Fermi-Dirac distribution for the occupation of a single-particle state of energy E:

$$f_{\mathrm{FD}}(E) = \frac{1}{e^{\frac{E-\mu}{k_B T}} + 1} \tag{1.1}$$

Nevertheless, several quantitative and qualitative mysteries remained. Most fundamentally, the theories of Drude and Sommerfeld failed to explain why some solid-state materials are nonmetallic [12].

To begin to classify materials as metals and insulators, it was necessary to treat the positive ions of Drude and Sommerfeld's theories more carefully. X-ray diffraction experiments suggested that these ions are often arranged in periodic lattices, motivating the study of the quantum mechanics of a particle in a perfectly periodic potential:

$$-\frac{\hbar^2}{2m}\nabla^2\Psi(\vec{r}) + U(\vec{r})\Psi(\vec{r}) = E\Psi(\vec{r}) \tag{1.2}$$

Here, $U(\vec{r}+\vec{R}) = U(\vec{r})$, where $\vec{R}$ is any of the elementary lattice vectors that translate the lattice into itself. The eigenstates of this problem satisfy Bloch's theorem:

$$\psi_{n,k}(\vec{r}+\vec{R}) = e^{i\vec{k}\cdot\vec{R}}\psi_{n,k}(\vec{r}) \tag{1.3}$$

where $\vec{k}$ is a vector within the unit cell of the reciprocal lattice defined by $e^{i\vec{K}\cdot\vec{R}} = 1$ (known as the *first Brillouin zone*). Bloch's theorem allows us to solve the Schrödinger equation within a single unit cell, because up to a phase, the eigenstates are periodic between unit cells. The eigenenergies $E_n(\vec{k})$ typically segregate into a series of *bands*, labeled by the index n, separated by energetic band gaps. In the case of the Schrödinger equation (1.2), there are an infinite number of bands, corresponding to the discrete, but infinite, solutions within a unit cell. An example for a particles in a weak periodic potential is shown in Figure 1.1.

Often, the problem is simplified further by making the *tight-binding approximation* that there are only a fixed number of orbitals that a particle can occupy within each unit cell:

$$\hat{H}_{\text{tb}} = J\sum_{j=1}^{L}(\hat{c}_{j+1}^{\dagger}\hat{c}_j + \hat{c}_j^{\dagger}\hat{c}_{j+1}) - \sum_{j=1}^{L}\mu_j\hat{n}_j \tag{1.4}$$

If the potential μ_j repeats every n lattice sites, then there is an n-site unit cell, and correspondingly, there will be n bands. In this tight-binding approximation, Bloch's theorem guarantees that the eigenstates will be periodic (again, up to a phase) each n sites.

Each band can hold two electrons per unit cell, corresponding to the two possible

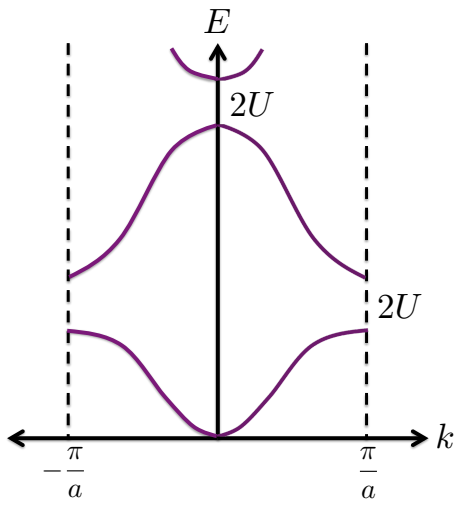

Figure 1.1: The energy E vs. quasimomentum k for particles in 1D in a weak periodic potential of period a and strength U. The quadratic dispersion for a free particle is folded into the first Brillouin zone, and degenerate states whose quasimomenta differ by an integer number of reciprocal lattice vectors get split by the potential. This induces band gaps of size $2U$.

values of electronic spin. Suppose there is an odd number of conduction electrons per unit cell. Then, in the ground state, the highest occupied band will only be half filled. In the limit of an infinitely large system, there will be no energetic barrier to adding another electron. In other words, the Fermi energy E_{F} lies in the conduction band, and the material conducts at low temperatures. Alternatively, suppose there is an even number of conduction electrons per unit cell. Then, the highest occupied band is fully filled, and it is necessary to pay the energetic cost of the band gap to add another electron. Therefore, E_F lies in a band gap, and the system is an insulator at low temperatures. This forms the most basic band classification of solid-state materials into metals and insulators [12].

1.1.2 Bose-Einstein Condensation

In parallel with these studies of the properties of solid-state materials, the observation of viscosity-free flow in liquid helium at low temperatures motivated the study of systems composed of many bosons. The occupation of single-particle energy levels by bosons is described the Bose-Einstein distribution:

$$f_{\mathrm{BE}}(E) = \frac{1}{e^{\frac{E-\mu}{k_B T}} - 1} \tag{1.5}$$

Meanwhile, the density of states of an ideal gas scales like $\rho(E) \sim E^{\frac{d-2}{2}}$. Einstein noted that, in three dimensions and at low enough temperature, there is a bound to the number of particles that can be accommodated in excited single-particle states. Hence, there will be a macroscopic accumulation of bosons in the single-particle ground state, a phenomenon now known as Bose-Einstein condensation [48]. The ground state condensate can be thought of as a single coherent quantum state with an overall number and phase.

Important experimental properties of helium-4 deviate from the predictions for the ideal Bose gas. For instance, helium-4 has a divergence in the specific heat at the transition to the so-called superfluid state instead of a cusp at the Bose-Einstein condensation temperature [62]. Furthermore, superfluidity can occur in two dimensions [83, 40]. Nevertheless, the ideal Bose gas treatment gives a rough estimate of the transition temperature in helium-4 and was recognized as being the conceptual foundation for the understanding of superfluid helium-4 [95].

1.2 The Role of Interactions

1.2.1 From Fermi gases to Fermi liquids

At least in the case of electrons in metals, it may seem like a miracle that the band theory picture is ever useful. After all, electrons interact strongly through Coulomb forces. Indeed, the typical Coulomb energy per particle in a good metal is typically three to five times the Fermi energy [42]. Why then is an independent electron picture often applicable?

Landau answered this question by arguing that there are indeed entities that can be treated as independent in a metal, but that these are not the bare electrons. Instead, in Landau's Fermi Liquid Theory, the bare electrons effectively get "dressed" by weak interactions, giving rise to emergent *quasiparticles*. These quasiparticles are in one-to-one correspondence with the bare electrons and are therefore also described by Fermi-Dirac statistics. However, these quasiparticles are effectively free (i.e., noninteracting). Landau showed that these emergent particles are stable to quasiparticle-quasiparticle scattering, due to the difficulty of constructing momentum-conserving scattering events with particles near the effective Fermi surface. Thus, the difference between the noninteracting and interacting electron systems could be encapsulated in a few renormalized parameters describing the quasiparticles (e.g., an effective quasiparticle mass m^* that differs from the bare electron mass m_e) [12]. Figure 1.2 shows the ground state momentum occupation function in a Fermi Liquid. In terms of the bare electrons, interactions induce some filling of momenta beyond the Fermi mo-

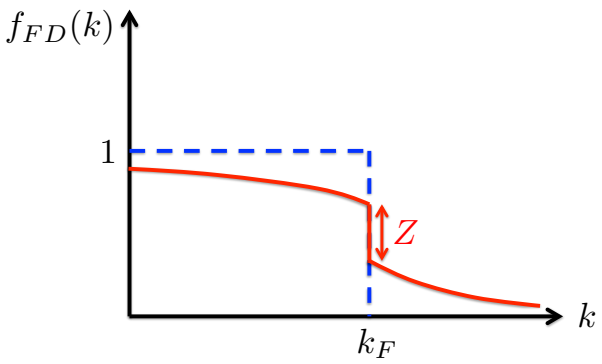

Figure 1.2: The ground state momentum occupation for noninteracting fermions (blue dashed line) and for a Fermi liquid (red solid line). In terms of the bare electrons, there is some filling of momenta beyond the Fermi wavenumber k_F, but there is a relic of the step function occupation in the noninteracting system (the "quasiparticle residue," here called Z). This feature is a reflection of the fact that the momentum occupation in terms of quasiparticles *is* indeed that of a noninteracting system.

mentum k_F of the non-interacting problem; however, in terms of *quasiparticles*, only momenta up to k_F are filled.

1.2.2 From Bose-Einstein Condensates to Superfluidity

To explain the discrepancies between experimental observations on helium-4 and Bose-Einstein condensation, it was also necessary to study the role of interactions in many-boson systems. Bogoliubov introduced weak interactions to a Bose gas and showed that excitations on top of the Bose-Einstein condensate obey a linear dispersion $\omega_{\vec{k}} \sim |\vec{k}|$ at small wavevector $|\vec{k}|$ [26]. Then, it was possible to invoke an argument by Landau that, to preserve energy and momentum conservation, the condensate cannot exchange energy with its environment through low-energy processes if it flows at velocities less than a critical value v_c [62, 6]. Thus, the condensate will flow without dissipation.

Eventually, the interacting Bose gas picture was also recognized as being intimately related to superconductivity and to superfluidity in helium-3. Electrons and helium-3 atoms are both fermions; thus, for a superfluid to form, a mechanism is required to "pair" these fermions into effective bosons. In ordinary superconductors, this is due to effective attractive electron-electron interactions mediated by phononic modes of the lattice [33, 16]. In helium-3, the pairing mechanism instead involves magnetic interactions between atoms of the same spin [62, 91].

1.2.3 Fundamental Failures of Noninteracting Models

We have seen above that, for sometimes rather subtle reasons, *noninteracting* models can provide useful starting points for the study of weakly *interacting* systems of bosons and fermions. There are circumstances, however, where noninteracting pictures fail completely. We will now discuss two such situations.

1.2.3.1 The Tomonaga-Luttinger Liquid in One Dimension

For instance, consider the low temperature properties of fermions in one spatial dimension. Here, no finite interaction strength can truly be considered weak. The Fermi surface in one dimension consists of only two points at $\pm k_F$. These degenerate states will be split by even weak interactions, destroying the entire Fermi surface and rendering the Fermi Liquid Theory inapplicable [139].

Tomonaga realized that the emergent "particles" in one dimension are not the quasiparticles of Fermi Liquid Theory but rather collective bosonic modes. He used the process of *bosonization* to map the system to a harmonic chain that describes

these modes:

$$\hat{H}_{\mathrm{b}} = \frac{u}{2\pi} \int dx \left[\frac{1}{K} (\pi \Pi_\phi)^2 + K (\partial_x \phi)^2 \right] \tag{1.6}$$

The Hamiltonian (1.6) represents the case of spinless fermions, and the bosonic operators $\hat{\phi}$ and $\hat{\Pi}_\phi$ obey $[\hat{\phi}(x), \hat{\Pi}_\phi(x')] = i\delta(x - x')$. The parameter K depends on the microscopic interactions: $K < 1$ corresponds to attractive interactions, $K = 1$ corresponds to noninteracting fermions, and $K > 1$ corresponds to repulsive interactions [136]. Luttinger calculated the ground state momentum distribution and showed that it is a power law that depends on K [97]. This is in contrast to the occupation function for a Fermi Liquid in Figure 1.2.

Later, Haldane realized that the model (1.6) is much more generally applicable to one-dimensional systems. In particular, it can also be applied to repulsively interacting bosons, but the relationship of K to the microscopic interactions is different: $K = \infty$ corresponds to the noninteracting limit, and K decreases with interaction strength [64].

1.2.3.2 The Mott Insulator

There is another fundamental limitation of noninteracting models that applies in arbitrary dimensions: the Fermi and Tomonaga-Luttinger Liquid Theories can break down at nonperturbative interaction strengths. Landau's argument, for example, relies on an adiabatic continuity between the bare electrons of the truly noninteracting model and the quasiparticles of the weakly interacting model [12]. It is important to question when this assumption of adiabatic introduction of interactions holds.

In fact, when Fermi Liquid Theory was proposed, experiments had already shown that there are insulating materials that band theory incorrectly classifies as conductors. Classic examples are the transition metal oxides with an odd number of electrons per unit cell (e.g., V_2O_3). Since these materials often exhibit antiferromagnetism at low temperatures, Slater proposed that the unit cell gets "doubled" due to the spin structure. Then, the insulating behavior would still follow from band theory since there would now be an even number of particles per unit cell. There was, however, a problem of scales: the antiferromagnetic ordering does not set in until $T \sim 10^2 K$ while the insulating behavior can persist until much higher temperatures of the order $T \sim 10^3 - 10^4 K$ [42]. This motivated Mott to look outside the band theoretic paradigm for an explanation for the insulating behavior. He reasoned that electrons cannot be mobile if they do not have enough kinetic energy to overcome strong electron-electron Coulomb repulsion. Instead, the electrons get localized in a particular unit cell, forming an insulating state known as the *Mott insulator* [105]. Each localized electron still carries a spin degree of freedom, and these spins can order at low temperatures, giving rise to the antiferromagnetic ground state. Thus, the insulating state sets the stage for the spin ordering and not vice versa [42].

Although motivated by the inadequacies of band theory in describing certain electronic systems, Mott's argument is clearly not restricted to fermions. Many-boson systems can also transition from the superfluid state to a Mott insulating state at strong interaction strengths, and we will now proceed to discuss models that can capture this physics in both fermionic and bosonic systems.

1.2.4 The Fermi and Bose Hubbard Models

A minimal model for studying the role of interactions in electronic systems is the *Fermi-Hubbard model*:

$$\hat{H}_{\mathrm{FH}} = J \sum_{\langle j,k\rangle,\sigma} (\hat{c}^\dagger_{j,\sigma}\hat{c}_{k,\sigma} + \hat{c}^\dagger_{k,\sigma}\hat{c}_{j,\sigma}) + U \sum_j \hat{n}_{j,\uparrow}\hat{n}_{j,\downarrow} - \mu \sum_{j,\sigma} \hat{n}_{j,\sigma} \tag{1.7}$$

Here, we introduce on-site interactions between electrons to a tight-binding model (1.4) with a single-site unit cell. The operators $\hat{c}_j$ and $\hat{c}^\dagger_j$ obey the fermionic anti-commutation relation $\{\hat{c}_j, \hat{c}^\dagger_k\} = \delta_{j,k}$ and the number operator $\hat{n}_j = \hat{c}^\dagger_j\hat{c}_j$. The ground state of the one-dimensional model at half-filling (i.e., one particle per unit cell) was obtained by Lieb and Wu, who showed that there is no conducting state for $U > 0$ [93]. In higher dimensions, no such exact solution has been found, and the Fermi-Hubbard model has inspired decades of numerical and theoretical work. The relevance of this model to the high-temperature superconductors and heavy fermion materials has especially motivated efforts to understand its phase diagram.

A somewhat simpler[2] variant of the Fermi-Hubbard model is the *Bose-Hubbard model*, where the particles hopping on the lattice are (typically spinless) bosons:

$$\hat{H}_{\mathrm{BH}} = J \sum_{\langle j,k\rangle} (\hat{b}^\dagger_j\hat{b}_k + \hat{b}^\dagger_k\hat{b}_j) + \frac{U}{2} \sum_j \hat{n}_j(\hat{n}_j - 1) - \mu \sum_j \hat{n}_j \tag{1.8}$$

Here, the operators $\hat{b}_j$ and $\hat{b}^\dagger_j$ obey the bosonic commutation relation $[\hat{b}_j, \hat{b}^\dagger_k] = \delta_{j,k}$

[2]The Bose-Hubbard model is simpler because it is amenable to quantum Monte Carlo, while fermionic models often suffer from the so-called "sign" problem that precludes efficient numerical simulation [94].

and the number operator $\hat{n}_j = \hat{b}_j^\dagger \hat{b}_j$. The ground state phase structure of this model was studied in the late 1980s by Fisher, Weichman, Grinstein, and Fisher and is shown in Figure 1.3 [55]. These authors began by examining the model at zero hopping strength, $J = 0$. Here, the model is trivially insulating with a fixed density of bosons per site that is determined by minimizing $\frac{U}{2} n(n-1) - \mu n$. As the hopping is raised, isolated boson-hole pairs are added on top of this inert background, but the overall density stays fixed. There remains a finite gap to the addition of single bosons or holes, and the system remains a Mott insulator. The transition to superfluity generically occurs when the gap to the addition of single bosons or holes closes, and then, any additional particles added to the system propagate freely on top of the inert background. This determines the general phase structure seen in Figure 1.3, with "Mott lobes" at small hopping. The superfluid-insulator transition takes on a special character at the tip of each lobe, where there is an emergent particle-hole symmetry. For this reason, the overall density of bosons stays fixed even after exiting the Mott lobe, and the ground state instead becomes a superposition of different ways to redistribute the same density around the system [55].

1.3 Aperiodicity and Localization

We have seen above that strong interactions can drive otherwise metallic or even superfluid systems into a Mott insulating state. There is another distinct route to insulating behavior that does not require interactions. To understand this mechanism, we need to relax the assumption of a perfectly periodic medium.

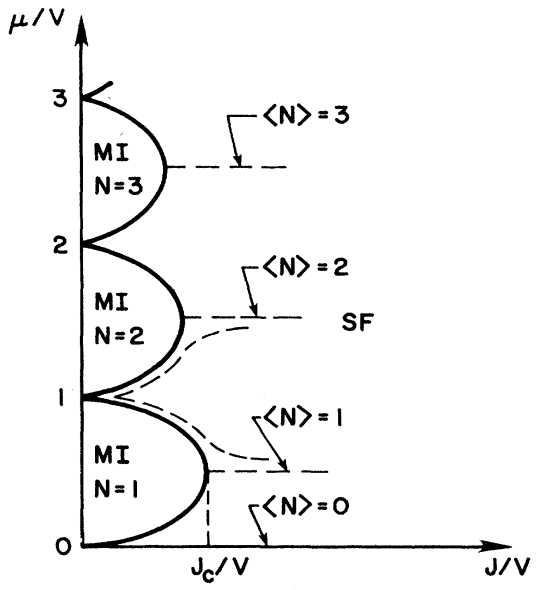

Figure 1.3: The general phase structure of the Bose-Hubbard model (1.8), as worked out by Fisher, Weichman, Grinstein, and Fisher [55]. Note that V in their phase diagram is the interaction strength that we call U in our definition of the Bose-Hubbard model (1.8).

1.3.1 Anderson Localization

In 1958, Anderson was motivated by Feher's experiments on spin diffusion in silicon to examine the role of disorder in quantum transport [11]. He studied a tight-binding single-particle model of the form (1.4) where μ_j varies randomly from site to site. Anderson sampled the μ_j randomly from a box distributions of width W. This serves as a toy model of the situation in real materials where the environment (e.g., a crystalline lattice) is never perfect: impurities and structural imperfections are inevitable. Anderson showed that, below some ratio of the hopping to the disorder strength, the single-particle eigenstates will all be localized. This means that their amplitude decays exponentially in space with a characteristic length ξ, the *localization length* [11]. This is in contrast to the *extended* states (1.3) that form in periodic

systems where Bloch's theorem holds. Soon thereafter, Mott and Twose argued that all eigenstates will be localized in one spatial dimension for arbitrarily weak disorder in the potential [107]. This reveals an important aspect of Anderson localization: it occurs even when $J >> W$, meaning that the kinetic energy is great enough to surmount any barriers imposed by the potential. Anderson localization is thus not due to any form of classical trapping and is instead a wave interference phenomenon.

In 1979, the nature of Anderson localization in higher dimensions was clarified by Abrahams, Anderson, Licciardello, and Ramakrishnan [2], who built on earlier ideas of Thouless and collaborators [45, 135, 92]. These authors formulated a scaling theory of localization by considering how the conductance g varies with the linear system size L in dimension d. At small disorder strength, the conductance should be ohmic:

$$g(L) \sim L^{d-2} \tag{1.9}$$

Meanwhile, at strong disorder, localization implies that the conductance decays exponentially with L:

$$g(L) \sim \exp\left(-\frac{L}{\xi}\right) \tag{1.10}$$

Abrahams et al. examined the function:

$$\beta(g) \equiv \frac{d \ln g}{d \ln L} \tag{1.11}$$

which describes how the conductance varies as the system size is changed. The goal is then to connect the ohmic (1.9) and localized (1.10) behaviors, from which it is

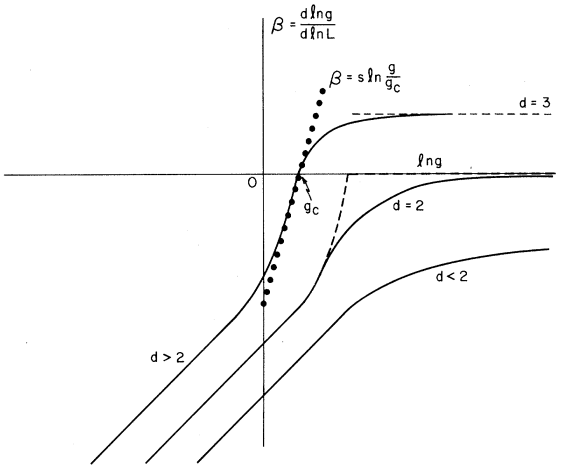

Figure 1.4: The beta function vs. conductance g, as derived in the paper by Abrahams et al. [2].

possible to extract a picture like the one in Figure 1.4. In one dimension, the β function (1.11) is always negative, indicating that the system flows towards lower and lower conductances as the sample size is raised. This is consistent with the notion that that one-dimensional disordered systems are always localized. In three dimensions, the β function crosses zero at some finite value of the conductance g. This indicates that the presence of an *Anderson transition* between localized and ohmic behavior. The case of two dimensions is marginal, because the level of analysis in equation (1.9) suggests that the ohmic conductance is independent of system size. Thus, the sign of the β function is determined by logarithmic corrections to equation (1.9). Abrahams et al. showed that this correction is typically negative, meaning that the two-dimensional model is always localized[3] [2].

Let us now focus on the nature of the Anderson transition in three dimensions. Figure 1.5 shows a caricature of the density of states of the three-dimensional An-

[3]The correction can become positive in models with spin-orbit coupling, leading to a transition [69].

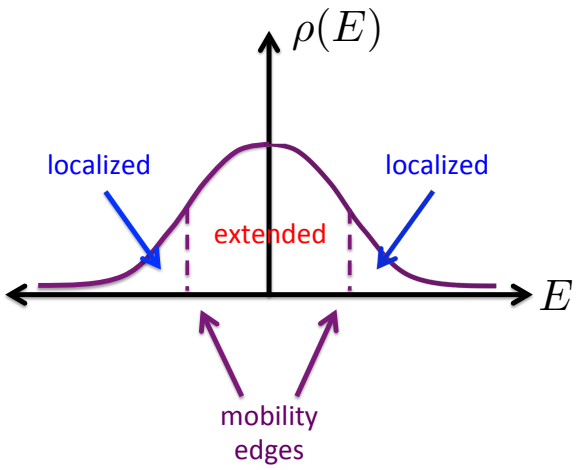

Figure 1.5: Sketch of the density of states of the three-dimensional Anderson model. Mobility edges separate the localized and extended states in the spectrum.

derson model. Whereas in one and two dimensions the entire spectrum would be localized, here extended and localized states can coexist. However, they generally cannot coexist at the same energy: as argued by Mott, if there were degenerate extended and localized states, they would be strongly mixed by small perturbations. This mixing would render both states extended [108]. Consequently, localized and extended states are typically separated by so-called *mobility edges*, as shown in Figure 1.5[4]. For a system of electrons, the conductance properties will be determined by the location of the Fermi energy (i.e., whether E_{F} lies in a region of localized or extended states). The Anderson transition can be tuned by changing the disorder strength, thus moving the mobility edge past E_{F}. Alternatively, the transition can be tuned by changing the number of electrons in the system at fixed disorder strength, thus tuning E_{F} past the mobility edge.

The scaling theory (as presented in Figure 1.4), suggests that the conductance

[4]At sufficiently strong disorder strength, the entire spectrum can be localized even in three dimensions, and in this case, there would be no mobility edge.

will go to zero as we approach the thermodynamic limit in a localized regime. In real solid-state systems, it is important to remember that the electrons interact with a phononic bath. By absorbing and emitting phonons, electrons can hop between distant localized orbitals. The probability of such an event depends upon the splitting in energy between the orbitals (i.e., hopping is more favored between nearly degenerate states) and on the distance separating the orbitals, because the overlap of the orbital wavefunctions decays exponentially in space. Mott determined how to balance these effects and thereby obtained the *variable-range hopping conductivity*:

$$\sigma \propto \exp\left[-\left(\frac{T_0}{T}\right)^{\frac{1}{d+1}}\right] \tag{1.12}$$

in d dimensions. Here, T is the temperature of the phononic bath, and T_0 is a characteristic temperature related to the typical energy splitting of orbitals within a localization volume ξ^d [106]. The variable-range hopping process is ultimately responsible for bringing the electrons into thermal equilibrium with the phononic bath. The problem of many-body localization, which we will discuss in Section 1.4.2.2 below, is intimately tied to the question of whether electron-electron interactions can induce variable-range hopping in the absence of coupling to phonons [17].

1.3.2 Other Models of Localization

1.3.2.1 Localization with Hopping Disorder

Localization is a general wave phenomenon that occurs in many disordered quantum systems; it is not special to the model studied by Anderson. For instance, localization also occurs with "off-diagonal" disorder in the hopping term:

$$\hat{H}_{\mathrm{rh}} = -\sum_{j=1}^{L} J_j(\hat{c}_{j+1}^{\dagger}\hat{c}_j + \hat{c}_j^{\dagger}\hat{c}_{j+1}) \tag{1.13}$$

This one-dimensional random hopping model has a particle-hole symmetry that is missing in the Anderson model. The related classical problem of the random harmonic chain was studied in the early 1950s by Dyson [43], and the quantum model received attention in the following decades [133, 46]. More recently, higher-dimensional versions of this model have been studied [103]. Furthermore, the model has been realized experimentally in photonic waveguides [89], and there is also a proposal to engineer it in atomic lattices [150].

Perhaps the most intriguing feature of the one-dimensional model is that all eigenstates are localized *except one*: the $E = 0$ state. This can be shown by setting $\hat{H}_{\mathrm{rh}}|\psi\rangle = 0$, which yields the equation:

$$-J_j\psi_{j+1} - J_{j-1}\psi_{j-1} = 0 \Rightarrow \frac{\psi_{j+1}}{\psi_{j-1}} = -\frac{J_{j-1}}{J_j} \tag{1.14}$$

This in turn implies that the localization length ξ near $E = 0$ obeys:

$$\begin{aligned}\frac{1}{\xi} &= -\lim_{n\to\infty}\frac{1}{2n}\ln\left|\frac{\psi_{2n+1}}{\psi_1}\right| \\ &= -\lim_{n\to\infty}\frac{1}{2n}[(\ln|J_{2n-1}| + \ln|J_{2n-3}| + \cdots + \ln|J_1|) \\ &\qquad -(\ln|J_{2n}| + \ln|J_{2n-2}| + \cdots + \ln|J_2|)] \end{aligned} \tag{1.15}$$

For uncorrelated disorder, the limit goes to zero. This means that ξ diverges, indicating the presence of an extended zero energy state. As a function of E, the localization length is known to diverge as $\xi(E) \sim -\ln|E|$ and this divergence is accompanied by a *Dyson singularity* $\rho(E) \sim \frac{1}{|E\ln^3 E|}$ in the density of states [43, 133, 149].

1.3.2.2 Aubry-André Localization

A couple of decades after Anderson's original work on localization, Aubry and André demonstrated that localization can even occur without "true" disorder. They studied a tight-binding model of the form:

$$\hat{H}_{AA} = J\sum_j(\hat{c}^\dagger_{j+1}\hat{c}_j + \hat{c}^\dagger_j\hat{c}_{j+1}) + h\sum_j\cos(2\pi kj + \delta)\hat{n}_j \tag{1.16}$$

with irrational wavevector k. This makes the potential *incommensurate* with the lattice, which in turn means that the potential pattern is deterministic but nonrepeating [13]. Harper had studied a similar model much earlier, but he had focused on a special ratio of hopping to potential strength [65]. Aubry and André showed that this point actually lies at a localization transition. It separates a weak potential phase, where

all single-particle eigenstates are extended, from a strong potential phase, where all eigenstates are localized.

The easiest way of demonstrating that the transition exists is by noting that the model (1.16) is self-dual when $\delta = 0$. On a finite size lattice, it is conventional to enforce periodic boundary conditions by choosing $k = \frac{\ell}{L}$ such that ℓ and L are mutually prime. The self-duality can be realized by switching to Fourier space ($c_j = \frac{1}{\sqrt{L}} \sum_q e^{iqj} c_q$) and then performing a rearrangement of the wavenumbers q such that the real-space potential term looks like a nearest-neighbor hopping in Fourier space and vice versa. The duality construction reveals that, if the AA model has a transition, it must occur at $\frac{J}{h} = \frac{1}{2}$. In the thermodynamic limit, there is indeed a transition at this value for nearly all irrational wavenumbers k [134]. When $\frac{J}{h} > \frac{1}{2}$, all single-particle eigenstates are spatially extended, and by duality, localized in momentum space; when $\frac{J}{h} < \frac{1}{2}$, all single-particle eigenstates are spatially localized, and by duality, extended in momentum space. Exactly at $\frac{J}{h} = \frac{1}{2}$, the eigenstates are multifractal [127, 67]. The spatially extended phase of the AA model is characterized by ballistic, not diffusive, transport [13]. Recently, Albert and Leboeuf have argued that localization in the AA model is a fundamentally more classical phenomenon than disorder-induced Anderson localization, and that the AA transition at $g = \frac{1}{2}$ is most simply viewed as the classical trapping that occurs when the maximum eigenvalue of the kinetic (or hopping) term crosses the amplitude of the incommensurate potential [4].

In the 1980s and 1990s, physicists continued to study this *quasiperiodic* local-

ization transition for its own peculiarities and because it mimics the situation in disordered systems in $d \geq 3$, where there is also a single-particle localization transition[5] [134, 30, 49, 1, 47, 127, 67, 138]. It is important to keep in mind that, in contrast to the three-dimensional Anderson model, the Aubry-André model (1.16) lacks a mobility edge. However, a mobility edge is introduced by many perturbations to the model (1.16) [116].

1.4 Disorder, Interactions, and Phase Transitions

We have seen above that interactions and disorder can both localize particles, leading to Mott and Anderson insulators respectively. We might, consequently, be led to believe that the combination of both ingredients would favor insulating behavior even more strongly. This naive assumption underestimates the considerable richness of the disordered many-body problem. Interactions can destroy disorder-induced localization, and disorder can transform Mott insulators into exotic glassy states. We will now discuss these fascinating phenomena. We will first address some general aspects of the role of disorder at and near phase transitions; then, we will concretize these ideas to two issues of special relevance to this thesis.

[5]The AA model was also actively investigated in the mathematical physics literature, because it involves a Schrödinger operator (i.e., the "almost Mathieu" operator) with particularly rich spectral properties. The contributions of mathematical physicists put the initial work on Aubry and André on more rigorous footing [128, 20, 77, 76].

1.4.1 Disorder at Phase Transitions

1.4.1.1 The Harris Criterion

Suppose we introduce disorder into a system that exhibits a continuous phase transition in the clean limit. Two questions immediately arise: does a sharp transition survive the presence of disorder, and if so, how do the universal properties of the transition change? In 1974, Harris formulated a general criterion that can guide the effort to address these questions. Harris originally formulated his argument in the context of thermal phase transitions tuned by the temperature. Here, we use the thermodynamic language for convenience, but Harris' argument applies equally well to a quantum phase transition tuned by some other parameter. At the clean transition, the correlation length ξ diverges as $|T - T_c|^{-\nu}$. Meanwhile, in a disordered system, there is actually a "local" critical temperature $T_{c,\text{loc}}$ that varies spatially, because of local variations in the couplings. To maintain the integrity of the clean critical point, the typical fluctuation of $T_{c,\text{loc}}$ in a correlation volume ξ^d ought to be smaller than the distance to criticality $|T - T_c|$. The central limit theorem suggests that fluctuations of the mean of $T_{c,\text{loc}}$ in a localization volume decay as $\xi^{-\frac{d}{2}}$. Thus, we require:

$$\xi^{-\frac{d}{2}} = |T - T_c|^{-\frac{d\nu}{2}} < |T - T_c| \tag{1.17}$$

This yields an inequality, the *Harris criterion*, that must be obeyed by the critical point of any disordered system:

$$\nu d > 2 \tag{1.18}$$

If a clean critical point obeys the condition (1.18), then the universal properties will be unaffected by introducing weak disorder. Alternatively, if the clean critical point violates the condition, there are two possibilities: weak disorder may completely smear the transition away, or a critical point with new universality may emerge [66, 120].

1.4.1.2 Griffiths Phases and Rare Events

Let us consider the case where a transition survives in the disordered system. Irrespective of whether or not the universality changes, there will be consequences for the phases that the transition separates. We label these abstract phases A and B. Due to the spatial variation in the couplings, when the bulk of the system is in phase B, there can still be regions that are locally in phase A, as shown in Figure 1.6. These regions are known as *Griffiths regions*, and the phase in which they appear is known as a *Griffiths phase*. If the disorder is spatially uncorrelated, the Griffiths regions get exponentially rare as their size increases. Nevertheless, arbitrarily large Griffiths regions can appear [61, 140].

Pollet et al. have recently argued that Griffiths regions exclude the possibility of standard phase transitions between energetically gapped and gapless phases in disordered systems. As we approach the transition from the side that is gapped in the clean system, there will appear arbitrarily large regions of the competing gapless phase, driving the system gapless [113]. According to Pollet et al., any transition between gapped and gapless phases in a disordered system must be of the "Griffiths type." This means that the transition occurs when the disorder distribution first allows the possibility of a region of the competing phase. To concretize this further,

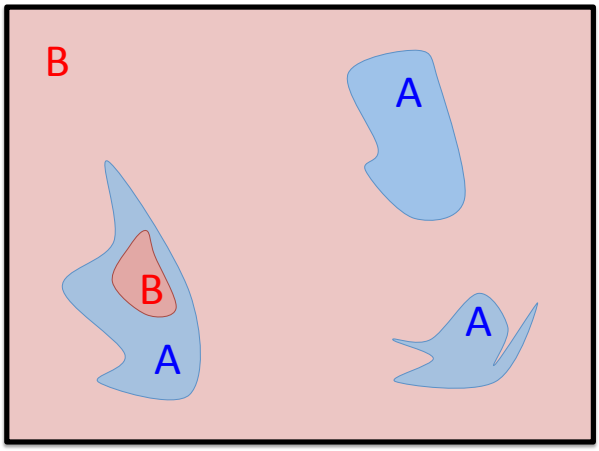

Figure 1.6: Schematic depiction of a Griffiths phase in a disordered system. The bulk of the system is in phase B, a gapped phase of the clean system. However, there are regions that are locally in the phase A, and these regions can get arbitrarily large in the thermodynamic limit. Since phase A is gapless, the Griffiths phase will also be gapless due to the presence of these rare regions.

suppose phases A and B are gapless and gapped phases of the clean system respectively, as shown in panel (a) of Figure 1.7. In the disordered system, two Griffiths phases will flank the "true" phase transition. This situation is shown in panel (b). Both Griffiths phases will be gapless, so the only transition between gapped and gapless phases is the transition out of the pure phase B, which is indeed of Griffiths type. Note that, if the disorder distributions are unbounded, it is always possible to find regions of the competing phase, so the gapped phase B will, presumably, be completely destroyed, as in panel (c) of Figure 1.7 [63].

Griffiths regions are examples of rare events in disordered systems, and these rare events can lead to very unusual behavior: average values of physical properties can behave very differently from their typical values. An instructive example was found by Fisher, who studied the paramagnet-ferromagnet phase transition of the random transverse field Ising model. He examined the two-point spin-spin correlation

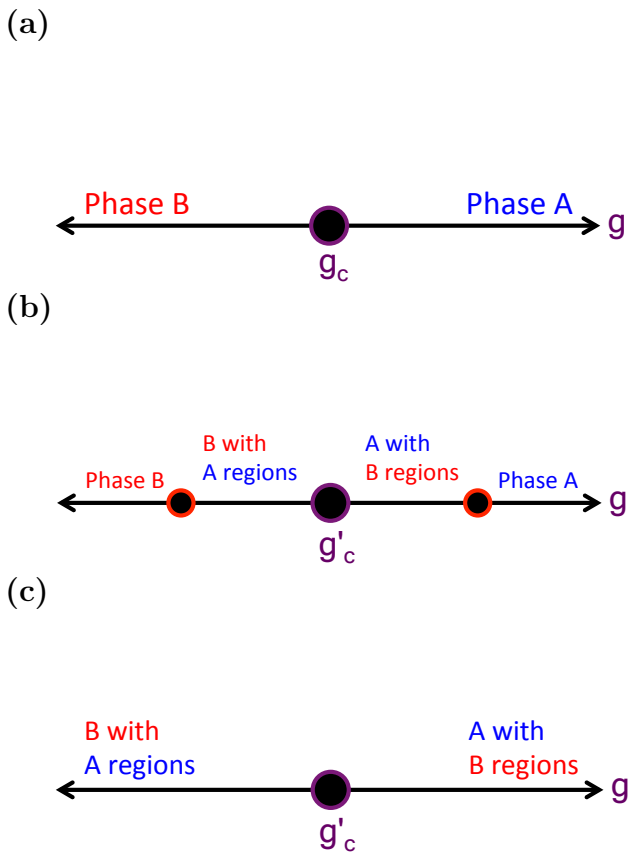

Figure 1.7: Schematic depictions of the impact of disorder around a phase diagram. In panel (a), we show a clean system with a transition between phase B and phase A at the tuning parameter $g = g_c$. With disorder, the tuning parameter varies randomly with mean g. When the bulk of the system is in phase B, there can be rare regions of phase A if the maximum value of the tuning parameter is on the phase A side of the clean phase transition. This situation, with Griffiths phases flanking the transition, is shown in panel (b). With unbounded disorder distributions, there can always be rare regions of the competing phase, as shown in panel (c).

function $\langle \hat{\sigma}_j^z \hat{\sigma}_{j+r}^z \rangle$ in the paramagnetic phase, and showed that the typical and mean correlations both decay exponentially. However, the correlation lengths over which these decay occur are very different, because the mean correlation at large distances is dominated by exponentially rare pairs that are perfectly correlated [54].

1.4.1.3 Infinite-Disorder Fixed Points and the Strong-Disorder Renormalization Group

The most extreme impact of disorder upon phase transitions occurs when a system exhibits a so-called *infinite-disorder fixed point*. These special fixed points are often identified through a powerful technique for studying strongly disordered systems, the *strong-disorder renormalization group* (SDRG). This method was originally developed by Dasgupta and Ma [37] and Bhatt and Lee [22]. Later, it was used extensively by Fisher to study a variety of one-dimensional quantum spin models [52, 54].

One way to motivate the procedure is to note that strong disorder makes the problem of finding the quantum ground state of a model more local. Having identified the strongest of all the disordered couplings in the Hamiltonian, we can then use the assumption of strong disorder to argue that other couplings in the real-space vicinity of this coupling are likely to be much weaker. This means that the ground state can locally be approximated by satisfying just this dominant coupling. Other terms in the Hamiltonian can then be incorporated as corrections[6]. This procedure generates new effective couplings in the model and thereby yields a new effective Hamiltonian. Since

[6]Quite often, these other terms are treated by means of perturbation theory, but this need not always be the case.

part of the ground state is specified in this step, some degrees of freedom of the system are decimated away. By repeating the procedure, we can iteratively specify the entire ground state [37, 22]. Moreover, we can examine the way in which the probability distributions of the disordered couplings flow as the renormalization proceeds. One possibility is that the model looks more and more disordered at larger length scales near criticality, thus flowing towards "infinite disorder." If this occurs, the strong disorder renormalization group becomes *asymptotically exact* near criticality, and it is sometimes even possible to predict quantities for which the corresponding behavior in the clean model is unknown [52, 53].

Fisher used the SDRG to show that several one-dimensional quantum spin models exhibit this type of fixed point. These include the random transverse-field Ising chain and various random antiferromagnets [52, 54]. Numerical work by Motrunich et al. subsequently provided evidence that infinite-disorder criticality also characterizes the two-dimensional random transverse-field Ising model [104]. The discrepancy between typical and mean values is most dramatic at infinite-disorder criticality. Again, a striking example was found by Fisher: in the random transverse-field Ising chain, the typical two-point correlations at criticality obey $\ln\langle\hat{\sigma}^z_j\hat{\sigma}^z_{j+r}\rangle \sim -\sqrt{x}$ while the mean two-point correlations obey $\langle\hat{\sigma}^z_j\hat{\sigma}^z_{j+r}\rangle \sim \frac{1}{x^{2-\phi}}$ where $\phi = \frac{\sqrt{5}-1}{2}$ is the golden mean. Since the latter is what is measured in neutron-scattering experiments, it is crucial to keep this distinction between mean and typical values in mind when confronted with infinite-disorder criticality [52, 54].

1.4.2 The Disordered, Interacting Many-Body Problem

We will now focus on two challenging problems of current theoretical and experimental interest: the dirty boson and many-body localization problems. In both cases, the interplay of disorder and interactions is absolutely crucial for characterizing the phases and phase transitions, but the inherent difficulty of each problem has left a full understanding elusive.

1.4.2.1 Superfluid-Insulator Transition for Disordered Bosons

In the 1980s, seminal experiments on helium adsorbed in Vycor first attracted the attention of theorists to the "dirty" or random boson problem [35, 115]. These experiments seemed to show a superfluid-insulator transition with some striking similarities to ideal Bose gas behavior[7]. Thus, in the decade before the realization of Bose-Einstein condensation in cold atomic gases, the helium in Vycor system was actually proposed as a possible realization of this elusive phenomenon. While studies of disordered bosons did not ultimately lead to the observation of Bose-Einstein condensation, the considerable richness of the dirty boson problem continued to stimulate theoretical interest.

Both this richness and the difficulty of the dirty boson problem originate in the volatility of the combination of disorder, interactions, and Bose statistics at low temperatures. The single-particle ground state in a disordered potential is generically localized [11]. The protocol for constructing the ground state of many noninteracting

[7]The most important similarity was the vanishing of the superfluid density $\rho_s \sim |T - T_c|^v$ near criticality. In particular, a crossover from the interacting exponent of $v = \frac{2}{3}$ to the ideal Bose gas value of $v = 1$ was observed as the density of helium was reduced [146].

bosons is to simply deposit all of them in this localized state, but such a configuration is *intrinsically unstable* to weak interactions. Thus, the noninteracting, disordered limit is pathological and cannot be the starting point for a perturbative analysis.

Instead, Giamarchi and Schulz pioneered the study of dirty bosons in one dimension by introducing weak disorder to a strongly interacting many-boson system [58]. The starting point of their analysis was a Hamiltonian of the form (1.6). Giamarchi and Schulz added a perturbative term describing a disordered potential and decomposed it into terms describing forward and backward scattering by impurities. They showed that the backward scattering leads to renormalization of the Luttinger parameter K and disorder strength D. Their proposed renormalization group flow diagram is shown in Figure 1.8. The superfluid-insulator transition is described by a fixed point at a universal value of the Luttinger parameter and vanishing disorder strength. For generic systems, the universal Luttinger parameter at criticality is $K = \frac{3}{2}$, and for particle-hole symmetric systems, it is $K = 2$. In either case, the transition is of Kosterlitz-Thouless type, meaning that superfluidity is destroyed by the proliferation of phase slips in the superfluid order parameter [83]. Disorder is dangerously irrelevant at the transition, because on the insulating side, the disorder strength grows under renormalization. Unfortunately, when this occurs, the perturbative approach of Giamarchi and Schulz breaks down and the insulating phase cannot be characterized [58].

The task of identifying the insulating phase was taken up by Fisher, Weichman, Grinstein, and Fisher, who studied the disordered Bose-Hubbard model in arbitrary

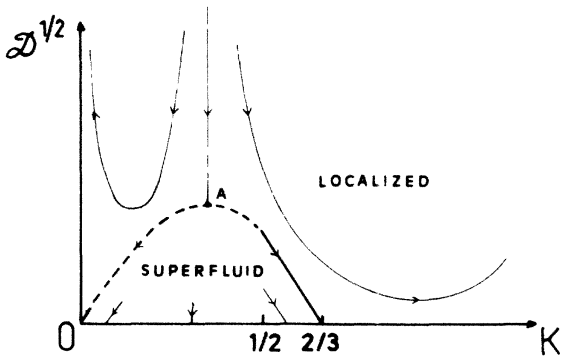

Figure 1.8: Renormalization group flows for weakly disordered bosons in one dimension, as derived by Giamarchi and Schulz [58]. The superfluid-insulator transition is described by a fixed point at vanishing disorder strength D. This transition separates a nearly uniform superfluid from an insulator in which the disorder strength grows with length scale, rendering the approach of Giamarchi and Schulz invalid.

dimension:

$$\hat{H}_{\mathrm{BH}} = J \sum_{\langle j,k \rangle} (\hat{b}_j^\dagger \hat{b}_k + \hat{b}_k^\dagger \hat{b}_j) + \frac{U}{2} \sum_j \hat{n}_j (\hat{n}_j - 1) - \sum_j \mu_j \hat{n}_j \tag{1.19}$$

Their proposed phase diagram, as presented in a review by Weichman, is shown in Figure 1.9 [146]. A Griffiths phase emerges and separates the Mott insulator and superfluid in the phase diagram. This phase is globally insulating but contains arbitrarily large regions of superfluid ordering. Thus, the phase is gapless. A finite fraction of the regions sit at the threshold for adding an additional boson; hence, in contrast to the Mott insulator, the phase has a finite compressibility:

$$\kappa \equiv \frac{1}{L^d} \sum_j \frac{\partial \langle \hat{n}_j \rangle}{\partial \mu} \tag{1.20}$$

Furthermore, in the thermodynamic limit, the largest regions of superfluid ordering

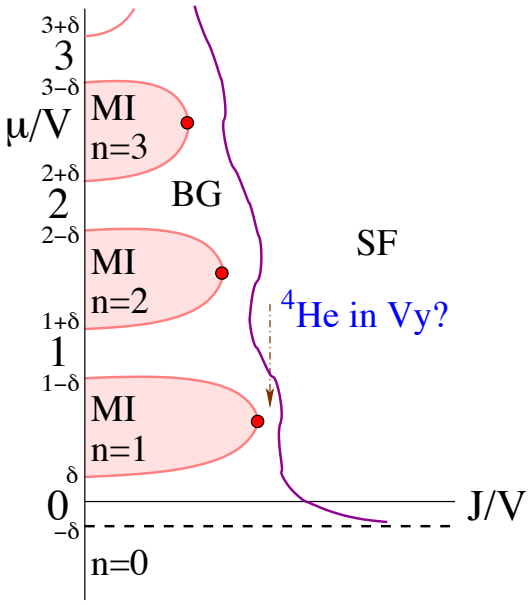

Figure 1.9: The phase diagram proposed by Fisher et al. for the disordered Bose-Hubbard model with $\mu_j \in (\mu - \delta, \mu + \delta)$. The Mott insulator and the superfluid are always separated by a new Griffiths phase, the Bose glass. This phase diagram is taken from a review by Weichman and also includes an estimate of where the helium-4 in Vycor system might fall in the phase diagram [146].

give the phase an *infinite* susceptibility:

$$\chi \equiv \frac{1}{L^d} \sum_j \frac{\partial \langle \cos\left(\hat{\phi}_j\right) \rangle}{\partial h} \tag{1.21}$$

when perturbed by an ordering field that couples to the phase of the superfluid order parameter[8]:

$$\hat{H}' = -h \sum_j \cos\left(\hat{\phi}_j\right) \tag{1.22}$$

This gapless, compressible, infinitely susceptible insulator is known as the *Bose glass*[9] [55]. Fisher et al. considered the possibility of a direct Mott insulator to superfluid transition in the presence of disorder, but this possibility appears to be excluded by the arguments of Pollet et al. [113].

There has been substantial progress on the dirty boson problem in the two and a half decades since the work of Giamarchi and Schulz and Fisher et al. For instance, the Bose-Hubbard model is amenable to quantum Monte Carlo, and in 2009, a numerical phase diagram for the three-dimensional model at unit filling was determined by Gurarie et al. This phase diagram is reproduced in Figure 1.10. It exhibits a highly nontrivial phase boundary and strong robustness of the superfluid phase to disorder [63]. Despite this progress, the dirty boson problem still presents many challenges. Most importantly, the universal properties of the superfluid-insulator transitions in

[8] As we will note later in the thesis, the boson creation and annihilation operators can be reexpressed in a phase-number representation $\hat{b}_j = \sqrt{\hat{n}_j} e^{i\hat{\phi}_j}$.

[9] In principle, there is a Griffiths phase also on the superfluid side of the transition, with large rare regions of insulator embedded in a superfluid bulk. However, because this phase maintains the long-range correlation properties of the superfluid phase, it is conventionally just referred to as a superfluid.

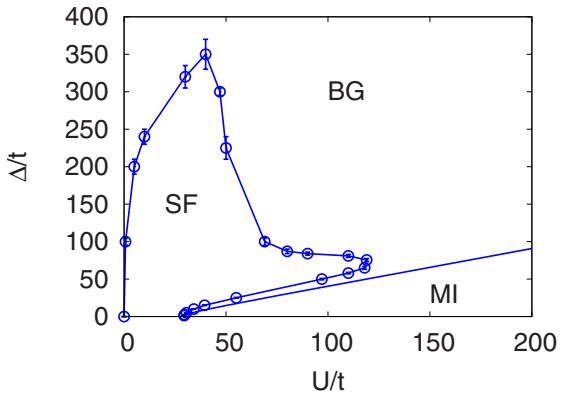

Figure 1.10: The phase diagram in the disorder $\frac{\Delta}{t}$ vs. interaction $\frac{U}{t}$ plane, as determined numerically by Gurarie et al. [63].

various dimensions are still contentious.

1.4.2.2 Thermalization and the Problem of Many-Body Localization

One of the goals of Anderson's original work on localization was to show that an isolated, disordered quantum system can fail to serve as its own heat bath. In a noninteracting localized phase, the system can get trapped near a nonequilibrium initial condition. Consequently, local quantities simply cannot equilibrate with their surroundings and never reach a thermal distribution [11]. Anderson's point is particularly crucial if it is also true in interacting systems. To determine if it is, it is necessary to understand the impact of interactions upon the full spectrum of a disordered system.

This issue has received renewed interest in recent years due to theoretical work by Basko, Aleiner, and Altshuler. In 2005, these authors studied a model in which all single-particle eigenstates are localized and considered the effect of introducing interactions. Suppose the operators $\hat{c}_\alpha^\dagger$ and $\hat{c}_\alpha$ are fermionic creation and annihilation

operators in the single-particle localized state indexed by α, and E_α is the energy of this single-particle state. Basko et al. proposed studying a model of the following form:

$$\hat{H}_{\mathrm{BAA}} = \sum_{\alpha} E_\alpha \hat{c}_\alpha^\dagger \hat{c}_\alpha + \frac{1}{2} \sum_{\alpha,\beta,\gamma,\delta} V_{\alpha\beta\gamma\delta} \hat{c}_\alpha^\dagger \hat{c}_\beta^\dagger \hat{c}_\gamma \hat{c}_\delta \tag{1.23}$$

Here, the interaction terms $V_{\alpha\beta\gamma\delta}$, which have characteristic strength λ, induce hopping on a high-dimensional lattice whose vertices represent different occupations of the single-particle localized states. We can consider the dynamics of the system beginning on one these vertices. In the absence of interactions, the system clearly remains localized, irrespective of the energy content. Basko et al. found that this picture changes qualitatively in the interacting problem. Figure 1.11 schematically depicts their surprising theoretical prediction for the behavior of the conductivity as a function of T. At low T[10], the system remains a *perfect* insulator, with $\sigma = 0$. Thermalization is impossible in this *many-body localized* phase. At larger T, interactions induce a finite σ, giving rise to a metallic or *many-body ergodic* phase in which thermalization is restored. The dividing point between these two regimes is now known as the *many-body localization (MBL) transition* [17].

This proposed metal-insulator transition is rather unique for several reasons. It is a quantum phase transition, as it involves a qualitative change in the nature of the many-body eigenstates; however, it occurs at finite energy density above the ground state. As such, it falls outside the usual paradigm of $T = 0$ quantum phase transitions [120]. Consequently, it *also* falls outside the usual paradigm of metal-insulator tran-

[10]The temperature T is more naturally interpreted as the energy density of the system.

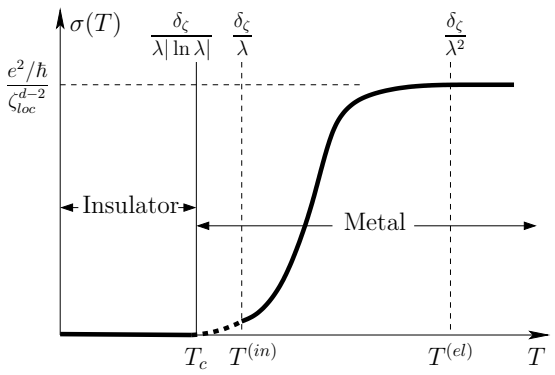

Figure 1.11: The DC conductivity σ vs. "temperature" T for an interacting system in which all single-particle eigenstates are localized, as depicted by Basko, Aleiner, and Altshuler [17]. The single particle localization length is $\zeta_{\rm loc}$ and the interaction strength is λ.

sitions, which are smeared into crossovers at any finite temperature [42]. Meanwhile, the many-body localization transition evades the arguments that prohibit $T \neq 0$ phase transitions in 1D, because it is precisely equilibrium statistical mechanics that breaks down when the transition occurs [109, 111]. This fact points to a final intriguing aspect of this phenomenon: work on the general problem of thermalization in quantum systems has led to the formulation of the so-called *eigenstate thermalization hypothesis* (ETH), which states that individual eigenstates of a generic quantum system should already encode thermal distributions of local quantities [129, 41]. The ETH is generically expected to hold in interacting quantum systems [117], but the many-body localized phase is an example of a situation in which this hypothesis fails.

Since 2005, the work of Basko et al. has inspired many numerical studies to confirm their predictions [109, 72, 78, 110, 111, 101, 21, 29, 15, 39, 79, 24]. These efforts are often confined to methods that scale in the same manner as exact diagonalization, because the problem demands information about highly excited states of the system.

Exact diagonalization studies *have* been fruitful: Oganesyan and Huse were able to identify a crossover in level statistics between localized and ergodic regimes, and Pal and Huse measured several other quantities that even provided estimates of the location of the transition in a spin model [109, 111]. Nevertheless, so far, it has been impossible to treat large systems and definitively characterize the transition. As such, many fundamental questions regarding the existence and nature of MBL remain unanswered and even unexplored.

1.5 Experimental Situation

This is a particularly exciting time to study localization and interaction effects in many-body systems, because these phenomena are being actively probed in the laboratory. We will now briefly discuss three particularly important experimental frontiers: solid-state systems, cold atoms, and photonic waveguides[11].

1.5.1 Solid-State Systems

We begin with the traditional domain of application of condensed-matter physics: the solid-state material. Here, the dirty boson problem, for instance, could be relevant to the low temperature behavior of disordered superconductors. A long-standing issue in understanding these materials has been to determine if the superconductor-insulator transition is "bosonic" in nature: in other words, do Cooper pairs survive on

[11]Note that the discussion in this section is definitely not representative of the state of the field as a whole. Instead, it is definitely biased towards experiments of possible relevance to the work reported in this thesis.

both sides of the transition, or alternatively, is the transition driven by pair breaking? In 2007, Crane et al. experimentally measured a finite-frequency superfluid stiffness on the insulating side of the superconductor-insulator transition in highly disordered indium oxide thin films. This indicates that Cooper pairs survive the transition and, consequently, that the transition may be understood in terms of the two-dimensional dirty boson problem [34]. This possibility has also been explored by Frydman et al. and Sherman et al. [56, 125].

Relatively recently, experimentalists have also realized superconductivity in a graphene sheet through proximity coupling to inhomogeneous puddles of tin deposited on its surface. A transmission electron microscope image of a sample of this "tin-decorated graphene" is shown in Figure 1.12. The resulting material behaves like a granular superconductor, with clean graphene separating the proximity-induced superconducting puddles. It is possible to probe the superconductor-insulator transition in this material by varying a gate voltage. Allain, Han, and Bouchiat have done so and found evidence for a transition that is bosonic in nature (i.e., it is driven by localization of Cooper pairs). Intriguingly, this transition exhibits properties similar to those we find in our analysis of the two-dimensional dirty boson problem in Chapter 2 below.

An entirely new experimental paradigm for the exploration of dirty boson physics has been developed by Yu et al., who built on theoretical work by Roscilde and Haas [148, 119]. Their work exploits an analogy between dirty bosons and certain nickel-based spin-one antiferromagnets in an applied magnetic field. Yu et al. have reported

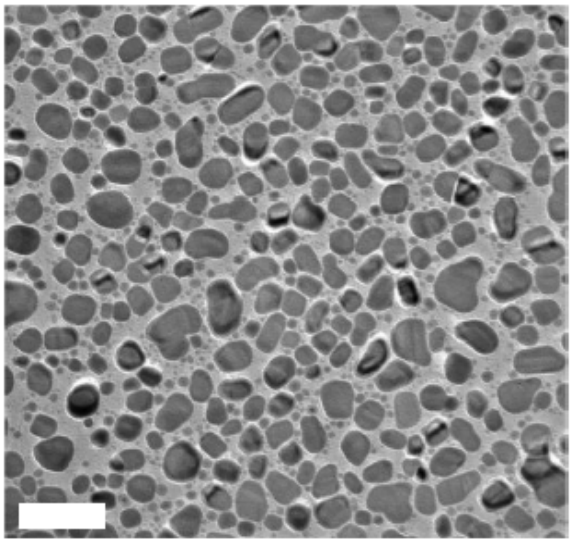

Figure 1.12: A transmission electron microscope image of tin-decorated graphene, as prepared in an experiment by Allain, Han, and Bouchiat [5]. The image shows puddles of tin on a graphene sheet. The scale bar is 200nm long.

evidence for a Bose glass in the spin system, and they have even reported evidence for an exotic incompressible Mott glass in the zero field limit[12]. Their finding is consistent with the relationship of exotic glassy phases to special symmetries of the dirty boson problem, a link that we explore later in this thesis.

There have even been proposals to study the physics of many-body localization in solid-state systems. Here, the interacting system of electrons is inevitably coupled to an external heat bath of phonons. Basko, Aleiner, and Altshuler have argued that, with this coupling, the transition gets smeared into a sharp crossover that nevertheless retains interesting signatures of the MBL phenomena. Specifically, the many-body localized phase survives as a regime in which conduction occurs *only* by phonon-induced hopping. Meanwhile, the many-body ergodic phase is characterized

[12]In the spin language, the compressibility (1.20) is actually a spin susceptibility to a longitudinal field, while the superfluid susceptibility (1.21) is a spin susceptibility to a transverse field.

by metallic transport. Basko et al. have suggested silicon doped with yttrium as a possible experimental system where these effects may be observable [18].

Finally, more peripheral to the focuses of this thesis, disorder effects have also been identified as crucial to certain technologically promising effects in complex oxides, for instance colossal magnetoresistance (CMR) in manganites. Experiments have shown that disorder leads to nanoscale phase separation in these materials, with different types of order coexisting precariously in different spatial regions. Then, small perturbations can tip the balance in favor of one type of order, inducing dramatic changes in physical quantities such as resistance. The CMR phenomenon is potentially very useful for constructing more powerful computer memories, so understanding disorder effects in these materials could be of significant practical relevance [36, 3].

1.5.2 Cold Atoms

In the mid-1990s, the first true Bose-Einstein condensates were realized in cold atomic gases by the groups of Wieman and Cornell at the University of Colorado, Boulder and Ketterle at MIT [10, 38]. Cold atomic gases have since developed into versatile experimental simulators of quantum condensed matter systems. These systems offer unprecedented control over parameters that cannot be easily tuned in conventional solid-state systems (e.g., particle-particle interaction strengths). In 2002, Greiner et al. were able to directly realize the Bose-Hubbard model (1.8) in a gas of rubidium-87 atoms and observe the three-dimensional Mott insulator to superfluid transition [60]. Following this milestone, several groups have introduced disorder into

their experimental setups in the hopes of observing Anderson localization and the Bose glass phase.

There are two primary methods of introducing disorder into cold atomic setups. The first is the method of optical speckle. Here, a primary lattice is generated by counterpropagating laser beams with wavelengths on the order of several hundred nanometers. Then, a random potential is superimposed on this lattice by passing another laser beam through a diffuser, which randomly scatters the light. This leads to random intensities of light in various spatial locations, generating a potential energy shift for the atoms which is proportional to the local intensity. Using this method, it is possible to generate disorder which varies on the scale of the lattice spacing. Beginning with the Inguscio group in 2005 [98], experimentalists have used speckle potentials to observe noninteracting Anderson localization [23, 82] and have also introduced strong interactions into the setup [147]. They have probed the superfluid-insulator transition [31] and have reported evidence of the superfluid to Bose glass transition [112].

The other commonly used method of introducing "disorder" into optical lattices is to superimpose an incommensurate modulating lattice on the primary lattice. This was proposed by Schulte et al. as an experimentally promising way of realizing localization physics [122], but it should be noted that this approach realizes the Aubry-André model (1.16) and not the Anderson model. In recent years, this experimental approach has been actively pursued by Inguscio and collaborators, who have observed the noninteracting Aubry-André transition [118] and have reported evidence

for a Bose glass [50]. They have additionally explored nonequilibrium physics within this setup by observing nonlinearity-induced subdiffusive expansion of a cloud of potassium atoms in the regime in which the noninteracting model is localized.

1.5.3 Photonic Waveguides

In the past decade, nonlinear light has emerged as another experimental tool to explore the physics of disorder. The essential idea is to fabricate single-mode photonic waveguides and then arrange them in an array such that their modes overlap. This realizes a tight-binding model with hopping terms between adjacent waveguides. Next, a potential can be introduced by varying the refractive index within each waveguide. Nonlinear interactions can also be incorporated by fabricating the waveguides from media that exhibit the optical Kerr effect, where the refractive index gets modified by the local intensity of light. An exciting feature of these systems is that the temporal propagation of light through the waveguide array is coupled to spatial propagation along the axis parallel to the waveguides. This allows for a direct imaging of dynamics [32]. Figure 1.13 provides a diagram of a typical experimental setup from a paper by Lahini et al. [87].

In 2007, the Segev group at the Technion reported the direct observation of Anderson localization in a two-dimensional lattice of waveguides [123]. Since then, the Silberberg group at the Weizmann Institute has made several important advancements. Figure 1.13 includes some images from their experiment showing Anderson localization in disordered waveguides [87]. This group has also experimentally realized

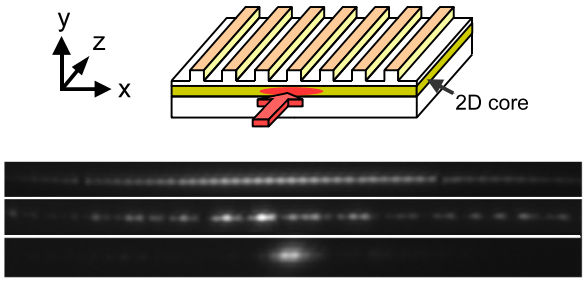

Figure 1.13: On top, a drawing of the experimental photonic waveguide setup from a paper by Lahini et al. [87]. The waveguides are arranged along the x-axis, and the light propagates along the z-axis. The red arrow shows the direction of the input beam. The three images below show experimental output patterns. The first is for periodic waveguides. The bottom two show localization in disordered waveguides; the two images correspond to cases where the input beam is aimed at weakly and strongly localized sites respectively.

the Aubry-André transition and explored the effects of nonlinearities upon Anderson and Aubry-André localized phases [87, 90]. Furthermore, they have been able to study differences in density-density correlation functions between light traveling in waveguides with diagonal and off-diagonal disorder[13] [89]. Finally, very recently, they have pioneered the study of topological edge states in quasiperiodic photonic waveguide lattices [85, 137]. These developments have opened an exciting new frontier in localization physics.

1.6 Overview of Thesis

The remainder of this thesis focuses on our own progress in understanding the dirty boson and many-body localization problems. Chapters 2 and 3 are devoted to the dirty boson problem in one and two dimensions respectively. First, in Chapter

[13]Off-diagonal (or hopping) disorder is realized by randomizing the spacing between adjacent waveguides[89].

2, we discuss the application of the strong-disorder renormalization group to the dirty boson problem by Altman, Kafri, Polkovnikov, and Refael [7, 9]. This work has raised the possibility that, in one dimension, the weak-disorder transition of Giamarchi and Schulz [58] could give way to a different critical behavior beyond a certain disorder strength. We analytically calculate the superfluid susceptibility (1.21) at the strong-disorder fixed point of Altman et al. Then, in Chapter 3, we numerically extend the approach of Altman et al. to higher dimensions, identify a new fixed point that governs the two-dimensional superfluid-insulator transition, and extract various universal properties of this transition. In Chapter 4, we turn our attention to the many-body localization problem and numerically investigate the real-time dynamics of a many-body quasiperiodic system. We provide evidence that the noninteracting Aubry-André transition becomes a many-body localization transition upon the introduction of interactions. We additionally develop toy models of the many-body localized and ergodic regimes and use our data to extract estimates for the phase boundary in the interactions versus hopping plane.

This thesis interpolates material from three papers by the author [74, 75, 73]. Chapter 2 uses material from References [74] and [75], both coauthored with David Pekker and Gil Refael. Meanwhile, Chapter 3 is based on Reference [75]. Finally, Chapter 4 is based on Reference [73], coauthored with Vadim Oganesyan, Gil Refael, and David Huse. Some material from each of these papers has also been incorporated into this introductory Chapter.

Chapter 2

Critical Susceptibility for Strongly-Disordered Bosons in One Dimension

2.1 Introduction

2.1.1 Motivation

In their seminal work on one-dimensional dirty bosons, Giamarchi and Schulz showed that perturbative disorder is irrelevant at the clean transition between the Mott insulating and superfluid phases [58]. Their analysis remains hugely influential in our understanding of the role of disorder in one-dimensional bosonic systems. Nevertheless, it is possible that the picture developed by Giamarchi and Schulz is not complete. Since their perturbative method cannot handle strong disorder, it cannot tell us whether the nature of the superfluid-insulator transition changes in this regime.

This possibility motivated Altman, Kafri, Polkovnikov, and Refael to formulate a strong-disorder renormalization group (SDRG) for the one-dimensional dirty boson problem. Using this approach, Altman et al. were able to identify a novel fixed point,

which may describe the superfluid-insulator transition in the strong-disorder regime. The transition proposed by Altman et al., while still being of Kosterlitz-Thouless type, occurs at a nonuniveral value of the Luttinger parameter. This difference from the transition of Giamarchi and Schulz is physically very significant: it suggests that, in the strong-disorder regime, superfluidity is destroyed by a process other than phase slips [7, 9].

The strong-disorder fixed point of Altman et al. remains controversial. It is not of the infinite-randomness class, and the SDRG procedure does not become asymptotically exact near criticality. Therefore, it is important to check the conclusions of the SDRG using other methods. The initial Monte Carlo simulations that followed the work of Altman et al. did not yield evidence in favor of strong-disorder criticality [14], motivating some to conclude that the transition is *always* of the type found by Giamarchi and Schulz. More recently however, Hrahsheh and Vojta performed a new Monte Carlo analysis, focusing on a stronger-disorder regime than was explored in the earlier simulations [70]. Their measurements of two quantities, the Luttinger parameter and the superfluid susceptibility, appeared to show the first independent evidence of strong-disorder criticality.

As we mentioned above, the nonuniversal critical Luttinger parameter was calculated by Altman et al. and was, in one sense, the key prediction of their work [9]. On the other hand, these authors did not calculate the critical superfluid susceptibility. To make closer contact between the SDRG and the Monte Carlo results of Hrahsheh and Vojta, we undertake a calculation of the susceptibility in this chapter.

2.1.2 Preview of the Results

Our principal result is that, at the strong-disorder transition, the divergence of the superfluid susceptibility is characterized by an anomalous dimension:

$$\lim_{L\to\infty} \frac{d\ln\chi}{d\ln L} = 2 - \eta_{\rm sd} \tag{2.1}$$

Here, the exponent:

$$\eta_{\rm sd} \approx \frac{1}{2\pi}\sqrt{y_i + y_i^2} \tag{2.2}$$

depends upon the bare disorder strength, parametrized by the quantity y_i. We plot $\eta_{\rm sd}$ as a function of y_i in Figure 2.1. The meaning of the parameter y_i will become more clear when we introduce the attractor distributions of the SDRG flows in Section 2.3 below. For now, note that $y_i = 0$ corresponds to a flat distribution of bare Josephson couplings. As y_i increases, the bare Josephson coupling distribution becomes progressively more strongly peaked near the RG scale, effectively reducing the disorder strength. Thus, the anomalous dimension monotonically increases as the disorder strength decreases.

The nonuniversal anomalous exponent (2.2) differs from the value at the transition of Giamarchi and Schulz, which is the Kosterlitz-Thouless value of $\eta = \frac{1}{4}$ [58]. As y_i is increased, eventually our estimate (2.2) will reach the Kosterlitz-Thouless value, presumably indicating the crossover to the weak-disorder regime. This crossover was indeed seen in the numerics of Hrahsheh and Vojta [70]. Regardless of the disorder strength, Hrahsheh and Vojta also observed that the anomalous dimension is related

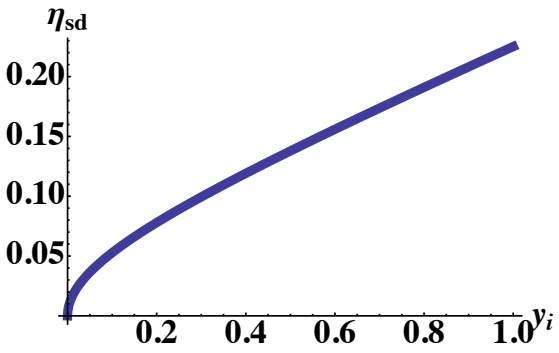

Figure 2.1: The anomalous exponent $\eta_{\rm sd}(y_i)$, as approximated in the small y_i regime by (2.2).

to the Luttinger parameter as $\eta \approx \frac{1}{2K}$. Our estimate (2.2) approximately matches the Luttinger parameter of Altman et al. via this relation [9].

However, our work also points to the need for some caution in interpreting numerical results. When tuning the transition near $y_i = 0$, the wideness of the initial distribution of Josephson couplings prohibits efficient self averaging of the anomalous exponent. Cleanly observing the "true" exponent $\eta_{\rm sd}$ requires exploring very large lattices, or equivalently, doing a very large amount of statistical sampling on smaller lattices.

2.1.3 Organization of the Chapter

We begin in Section 2.2 by introducing our model of dirty bosons, the disordered rotor model[1], and its relationship to the more familiar Bose-Hubbard model (1.19). Next, in Section 2.3, we describe the application of the strong-disorder renormalization group to this model. We present the renormalization group flows that were

[1]This is also the model that we will study in two dimensions in Chapter 3.

derived analytically by Altman et al. and also comment on the important relationship that these authors found between symmetry properties of the rotor Hamiltonian and the nature of the insulating phase.

We proceed to our calculation of the superfluid susceptibility in Section 2.4. We outline the general structure of our calculation, derive some intermediate results, and then piece together the calculation of χ and the anomalous dimension η. At key stages in the calculation, we compare to a numerical implementation of the SDRG. Then, in Section 2.5, we discuss the relationship of our results with the work of Altman et al. and Hrahsheh and Vojta. Finally, in Section 2.6, we summarize and examine the implications of our work for the superfluid-insulator transition.

2.2 Model

The model that we will study in this chapter and the next is the so-called *disordered rotor model*. We can motivate this model by beginning with a disordered Bose-Hubbard Hamiltonian that includes randomness in the interaction and hopping along with the usual chemical potential disorder[2]:

$$\begin{aligned} \hat{H}_{\mathrm{BH}} &= -\sum_{\langle jk \rangle} t_{jk}(\hat{b}_j^\dagger \hat{b}_k + \hat{b}_k^\dagger \hat{b}_j) + \sum_j U_j \hat{b}_j^\dagger \hat{b}_j (\hat{b}_j^\dagger \hat{b}_j - 1) \\ &\quad -\sum_j \mu_j \hat{b}_j^\dagger \hat{b}_j \end{aligned} \tag{2.3}$$

[2]We leave the dimension unspecified for the moment, because we will eventually study this model in both 1D and 2D.

Here, the creation and annihilation operators satisfy bosonic commutation relations:

$$[\hat{b}_j, \hat{b}_k^\dagger] = \delta_{jk} \tag{2.4}$$

and the hopping is between all nearest-neighbor sites on the lattice. An alternative representation of this model is given by the number and phase operators:

$$\hat{b}_j = e^{i\hat{\phi}_j}\sqrt{\hat{N}_j} \tag{2.5}$$

$$[\hat{\phi}_j, \hat{N}_k] = i\delta_{jk} \tag{2.6}$$

In terms of these operators:

$$\begin{aligned} \hat{H}_{\mathrm{BH}} &= -\sum_{\langle jk\rangle} t_{jk}\left(\sqrt{\hat{N}_j}e^{-i\hat{\phi}_j}e^{i\hat{\phi}_k}\sqrt{\hat{N}_k}\right. \\ &\qquad \left. +\sqrt{\hat{N}_k}e^{-i\hat{\phi}_k}e^{i\hat{\phi}_j}\sqrt{\hat{N}_j}\right) \\ &\quad +\sum_j U_j(\hat{N}_j-1)\hat{N}_j - \sum_j \mu_j\hat{N}_j \end{aligned} \tag{2.7}$$

To obtain a large, commensurate boson filling N_0, the chemical potential is tuned such that the on-site interaction and chemical potential terms are minimized for $\hat{N}_j = N_0$. This consideration fixes $\mu_j = (2N_0 - 1)U_j$. Then, if we expand the number operators

around this large filling as $\hat{N}_j = N_0 + \hat{n}_j$, the Hamiltonian becomes:

$$\begin{aligned}\hat{H}_{\mathrm{BH}} &= -\sum_{\langle jk\rangle} t_{jk} N_0 \left(\sqrt{1+\frac{\hat{n}_j}{N_0}} e^{-i\hat{\phi}_j} e^{i\hat{\phi}_k} \sqrt{1+\frac{\hat{n}_k}{N_0}} \right. \\ &\qquad \left. + \sqrt{1+\frac{\hat{n}_k}{N_0}} e^{-i\hat{\phi}_k} e^{i\hat{\phi}_j} \sqrt{1+\frac{\hat{n}_j}{N_0}} \right) \\ &\quad + \sum_j U_j \hat{n}_j^2 + (\text{const.}) \end{aligned} \tag{2.8}$$

The operators $\hat{n}_j$ now correspond to the particle number deviations from the large filling N_0. As such, n_j can take on any integer value from $-N_0$ to ∞, but we assume that N_0 is so large that we can let n_j run from $-\infty$ to ∞. The same approximation allows us to drop subleading terms in the hopping. We finally define the couplings $J_{jk} = 2t_{jk}N_0$ to arrive at the quantum rotor Hamiltonian:

$$\hat{H}_{\mathrm{rot}} = -\sum_{\langle jk\rangle} J_{jk} \cos(\hat{\phi}_j - \hat{\phi}_k) + \sum_j U_j \hat{n}_j^2 \tag{2.9}$$

This model, constructed as the large filling limit of a Bose-Hubbard Hamiltonian, can also be viewed as a description of an array of superconducting islands connected by Josephson junctions [7, 9].

When the Josephson couplings J_{jk} and charging energies U_j are uniform, the rotor model (2.9) exhibits a quantum phase transition between superfluid and Mott insulating phases at zero temperature. This transition is in the universality class of the classical XY transition in one higher dimension [146, 55, 28]. Consequently, the transition of the clean, one-dimensional rotor model is of Kosterlitz-Thouless type [83]. The critical exponent ν at a Kosterlitz-Thouless transition is effectively

infinite, meaning that the Harris criterion (1.18) is satisfied. From this perspective, it is perhaps unsurprising that weak disorder is irrelevant at this transition [58]. Nevertheless, this does not exclude the possibility that new physics may emerge at strong disorder, as we shall see in Section 2.3.

Irrespective of the disorder strength, the Hamiltonian (2.9) respects two important symmetries. First, there is the global U(1) phase rotation symmetry:

$$\hat{\phi}_j \to \hat{\phi}_j + \varphi \tag{2.10}$$

This means that the Hamiltonian conserves total particle number:

$$\hat{n}_{\text{tot}} = \sum_j \hat{n}_j \tag{2.11}$$

The model is also globally particle-hole symmetric:

$$\begin{aligned} \hat{n}_j &\to -\hat{n}_j \\ \hat{\phi}_j &\to -\hat{\phi}_j \end{aligned} \tag{2.12}$$

The particle-hole symmetry exists because the chemical potential coupling to the true particle number $\hat{N}_j$ has been tuned precisely to the value that enforces the large, commensurate filling. If this chemical potential is allowed to deviate from this value, then it would manifest in a disordered potential coupling to the particle number fluctuation $\hat{n}_j$, or equivalently, in offsets to the large filling. For simplicity, we do not

consider random offsets in our work. However, in this chapter and the next, we will occasionally comment on the likely influence of reintroducing these offsets.

2.3 Methodology: SDRG for the Disordered Rotor Model

In Chapter 1, we conceptually introduced the basic idea of the strong-disorder renormalization group (SDRG). Now, we will concretely describe the application of the method to the disordered rotor model (2.9). Our discussion follows the approach of Altman, Kafri, Polkovnikov, and Refael [7, 8, 9].

2.3.1 The Basic RG Steps

There are two types of disordered couplings in the rotor model: charging energies and Josephson couplings. In each step of the renormalization, we identify the maximum of all of these couplings, which defines the RG scale:

$$\Omega = \max\left[\{U_j\}, \{J_{jk}\}\right] \tag{2.13}$$

How we then proceed depends upon which type of coupling is dominant.

2.3.1.1 Site Decimation

Consider the case where the charging energy on site X is dominant. We define a local Hamiltonian in which this charging energy term is chosen to be the unperturbed

piece. All Josephson couplings entering the corresponding site are considered to be perturbations[3]:

$$\hat{H}_X = U_X \hat{n}_X^2 - \sum_k J_{Xk} \cos\left(\hat{\phi}_X - \hat{\phi}_k\right) \tag{2.14}$$

Satisfying the dominant coupling means setting $n_X = 0$ to leading order. This defines a degenerate manifold of local ground states: $|0, \{n_k\}\rangle$. In these kets, the first term corresponds to zero number fluctuation on site X and the second specifies the number fluctuations on all sites connected to X by a Josephson coupling. The degenerate space is infinitely large, corresponding to all possible choices of $\{n_k\}$. However, all matrix elements of the perturbative Josephson couplings in this ground state manifold vanish. The leading corrections then come from second order degenerate perturbation theory, in which we calculate corrections coming from excited states:

$$\begin{aligned} |0, \{n_k\}\rangle' &\approx |0, \{n_k\}\rangle \\ &\quad + \sum_{m \in k} \frac{J_{Xm}}{2U_X} (|1, n_m - 1\rangle + |-1, n_m + 1\rangle) \end{aligned} \tag{2.15}$$

In the terms giving the perturbative corrections, we assume that the number fluctuations on all neighboring sites except m remain unmodified from their values in $\{n_k\}$. We next consider the matrix elements of these states in $\hat{H}_X$. Up to a constant term, these matrix elements are identically those that would result from an effective Josephson coupling:

$$\tilde{J}_{jk} = \frac{J_{jX} J_{Xk}}{U_X} \tag{2.16}$$

[3]In 1D, there always only two Josephson couplings penetrating this site, and this property is strictly preserved by the RG.

Figure 2.2: The site decimation RG step in 1D.

between each two sites that were coupled to site X before the decimation step [7]. This process of site decimation is illustrated in Figure 2.2.

2.3.1.2 Link Decimation

Now, suppose that a Josephson coupling sets the RG scale. In this case, the local Hamiltonian is:

$$\hat{H}_{jk} = U_j \hat{n}_j^2 + U_k \hat{n}_k^2 - J_{jk} \cos\left(\hat{\phi}_j - \hat{\phi}_k\right) \tag{2.17}$$

The local approximation to be made here is that, to lowest order, the phases on these adjacent sites should be locked together. In other words, the degree of freedom to be specified is the relative phase $\phi_j - \phi_k$. This motivates the introduction of new operators:

$$\begin{aligned}
\hat{n}_C &= \hat{n}_j + \hat{n}_k \\
\hat{\phi}_C &= \frac{U_k \hat{\phi}_j + U_j \hat{\phi}_k}{U_j + U_k} \\
\hat{n}_R &= \frac{U_j \hat{n}_j - U_k \hat{n}_k}{U_j + U_k} \\
\hat{\phi}_R &= \hat{\phi}_j - \hat{\phi}_k
\end{aligned} \tag{2.18}$$

These operators satisfy the commutation relations:

$$[\hat{\phi}_C, \hat{n}_C] = i$$

Figure 2.3: The link decimation RG step in 1D.

$$[\hat{\phi}_R, \hat{n}_R] = i \tag{2.19}$$

with all other commutators vanishing. Thus, the transformation preserves the algebra of number and phase operators. A subtlety arises for the relative coordinate operators $\hat{n}_R$ and $\hat{\phi}_R$ because, as defined above, n_R need not be an integer and $\phi_R \in [-2\pi, 2\pi)$ as opposed to $\phi_R \in [0, 2\pi)$. To deal with this difficulty, we may make the additional approximation of treating ϕ_R as a noncompact variable. This makes n_R continuous instead of discrete. Then, in terms of the new cluster and relative coordinate operators, the local Hamiltonian (2.17) reads:

$$\hat{H}_{jk} = \frac{U_j U_k}{U_j + U_k} \hat{n}_C^2 + (U_j + U_k)\hat{n}_R^2 - J_{jk} \cos(\hat{\phi}_R) \tag{2.20}$$

To lowest order, we set $\phi_R = 0$. This decimation of the relative phase leaves the cluster phase ϕ_C unspecified; thus, two phase degrees of freedom have been reduced to one. The first term in $\hat{H}_{jk}$ shows that the inverse charging energies add like the capacitances of capacitors in parallel to give the charging energy for the cluster:

$$\tilde{U}_C = \frac{1}{\frac{1}{U_j} + \frac{1}{U_k}} = \frac{U_j U_k}{U_j + U_k} \tag{2.21}$$

Figure 2.3 depicts this process of link decimation [7].

2.3.2 Flow Equations, Their Solutions, and the Superfluid-Insulator Transition

In one dimension, it turns out that the SDRG can be implemented analytically. In other words, it is possible to translate the rules (2.16) and (2.21) into equations describing how the *distributions* of charging energies and Josephson couplings flow under the RG. We can then identify solutions to these *flow equations* which reveal universal features of the phases and phase transitions of the model (2.9) [7, 8, 9]. We will now describe this analytical approach, in preparation for applying it to our calculation of the susceptibility in Section 2.4.

Before formulating the flow equations, some preparatory work is necessary. First, it is helpful to define an RG "time":

$$\Gamma = \ln\left(\frac{\Omega_i}{\Omega}\right) \tag{2.22}$$

Here, Ω_i is the RG scale at the beginning of the procedure. In one dimension, the RG scale Ω monotonically decreases as the RG proceeds, and Γ monotonically increases from zero. It is also convenient to define the quantities:

$$\zeta = \frac{\Omega}{U} - 1 \tag{2.23}$$

and

$$\beta = \ln\left(\frac{\Omega}{J}\right) \tag{2.24}$$

In terms of the quantities ζ and β, the RG rules become *additive*. For instance, the site decimation rule (2.16) can be written as:

$$\beta_{jk} = \beta_{jX} + \beta_{Xk} \tag{2.25}$$

Similarly, the link decimation rule (2.21) becomes:

$$\zeta_C = \zeta_j + \zeta_k + 1; \tag{2.26}$$

The goal is then to examine how the probability distributions of ζ and β, called $f(\zeta, \Gamma)$ and $g(\beta, \Gamma)$ respectively, evolve as a function of Γ. This evolution is tracked by the flow equations:

$$\begin{aligned} \frac{df}{d\Gamma} &= (1+\zeta)\frac{\partial f}{\partial \zeta} + g(0,\Gamma)\int d\zeta_1 d\zeta_2 f(\zeta_1,\Gamma) f(\zeta_2,\Gamma) \\ &\quad \times \delta(\zeta - \zeta_1 - \zeta_2 - 1) + f(\zeta,\Gamma)(f(0,\Gamma) + 1 - g(0,\Gamma)) \end{aligned} \tag{2.27}$$

$$\begin{aligned} \frac{dg}{d\Gamma} &= \frac{\partial g}{\partial \beta} + f(0,\Gamma)\int d\beta_1 d\beta_2 g(\beta_1,\Gamma) f(\beta_2,\Gamma)\delta(\beta - \beta_1 - \beta_2) \\ &\quad + f(\zeta,\Gamma)(g(0,\Gamma) - f(0,\Gamma)) \end{aligned} \tag{2.28}$$

These integrodifferential equations look imposing, but in fact, all the terms have straightforward interpretations. In each case, the first term comes from the redefinition of ζ or β as Γ grows. Meanwhile, the second term encodes the creation of new links or clusters via the decimation processes (2.25) and (2.26). Finally, the last term

accounts for the removal of sites and/or bonds and keeps the distributions normalized.

Altman et al. found that there exists an attractor set of solutions to the equations (2.27) and (2.28)[4]:

$$f(\zeta, \Gamma) = f_0(\Gamma) e^{-\zeta f_0(\Gamma)} \tag{2.29}$$

$$g(\beta, \Gamma) = g_0(\Gamma) e^{-\beta g_0(\Gamma)} \tag{2.30}$$

It is worthwhile to keep in mind what this means in terms of the unscaled couplings J and U: the distribution of Josephson couplings flows to a power law with exponent $g_0 - 1$, and the distribution of charging energies flows to $\frac{1}{U^2}$ with an exponential cutoff at small U. The latter property indicates that clusters become exponentially rare at small energies. The parameters $f_0(\Gamma)$ and $g_0(\Gamma)$ satisfy the differential equations:

$$\frac{df_0}{d\Gamma} = f_0(1 - g_0) \tag{2.31}$$

$$\frac{dg_0}{d\Gamma} = -f_0 g_0 \tag{2.32}$$

The quantity:

$$\epsilon = f_0 - g_0 + \ln g_0 + 1 \tag{2.33}$$

is an invariant of the flows (2.31) and (2.32). It thus serves as a natural tuning parameter for the problem. Altman et al. worked out the flows for various values of ϵ, and the resulting flow diagram is shown in Figure 2.4.

[4]For these to be solutions to flow equations, it is necessary to drop the 1 inside the delta function in equation (2.27). This is a reasonable approximation if most cluster formation occurs far away from the RG scale (i.e., involving sites with high ζ or, equivalently, low charging energies).

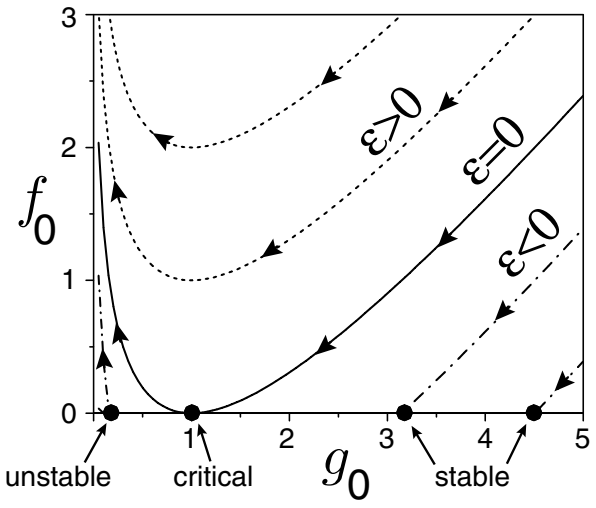

Figure 2.4: Schematic flow diagram for the one-dimensional disordered rotor model, as drawn in a paper by Altman, Kafri, Polkovnikov, and Refael [7]. See the text for details.

When $\epsilon = 0$, the flow terminates at an unstable fixed point at $(f_0, g_0) = (0, 1)$. To the left of the flow diagram, for $\epsilon > 0$, flows propagate towards an stable fixed point at $(f_0, g_0) = (\infty, 0)$. The system partitions into many small and disconnected superfluid islands, with no phase coherence between them. This is an insulating regime. Meanwhile, to the right of the fixed point, flows terminate at a line of fixed points with $f_0 = 0$. The RG builds ever larger and larger clusters, eventually absorbing most bare sites into a spanning cluster[5]. These fixed points correspond to superfluids with vanishing charging energies and various power law distributions of Josephson couplings. Interestingly, this point *also* corresponds to a model with vanishing charging energies (i.e., a *classical* model). This implies that the transition can be tuned at arbitrarily weak interaction strength by varying the distribution of Josephson couplings. It is also important to note that, at the unstable fixed

[5]In 1D, this cluster will not have true long-range order but rather power-law correlations [83].

point, the Josephson coupling distribution is *flat*. This is *not* an infinite-randomness distribution, and the SDRG is *not* asymptotically exact near criticality. Instead, it is an approximation whose predictions ought to be checked by other methods.

The most important predictions of Altman et al. concern the universality of the strong-disorder transition. In the vicinity of the fixed point, we can define $v(\Gamma) = \sqrt{f_0(\Gamma)}$ and expand $g_0(\Gamma)$ around one to find the flow equations:

$$\frac{dv}{d\Gamma} = \frac{v}{2}(1 - g_0) \tag{2.34}$$

$$\frac{dg_0^{-1}}{d\Gamma} = v^2 \tag{2.35}$$

This rewriting reveals that the flow equations take on the Kosterlitz-Thouless form near criticality [83]. We have already noted that the weak-disorder transition of Giamarchi and Schulz is also of Kosterlitz-Thouless type, and the two transitions share this property [58]. However, elsewhere there are important technical and conceptual differences. In particular, Altman et al. used the SDRG to calculate the effective Luttinger parameter at criticality and found:

$$K = \pi \sqrt{\frac{2}{\exp \frac{4}{\Gamma_0} - 1}} \tag{2.36}$$

Here, Γ_0 is a measure of the disorder strength in the bare system (i.e., before application of the SDRG). The Luttinger parameter (2.36) is therefore *nonuniversal*. This conflicts with the picture of Giamarchi and Schulz, where the critical Luttinger parameter in the particle-hole symmetric model is always 2. Altman et al. argued that

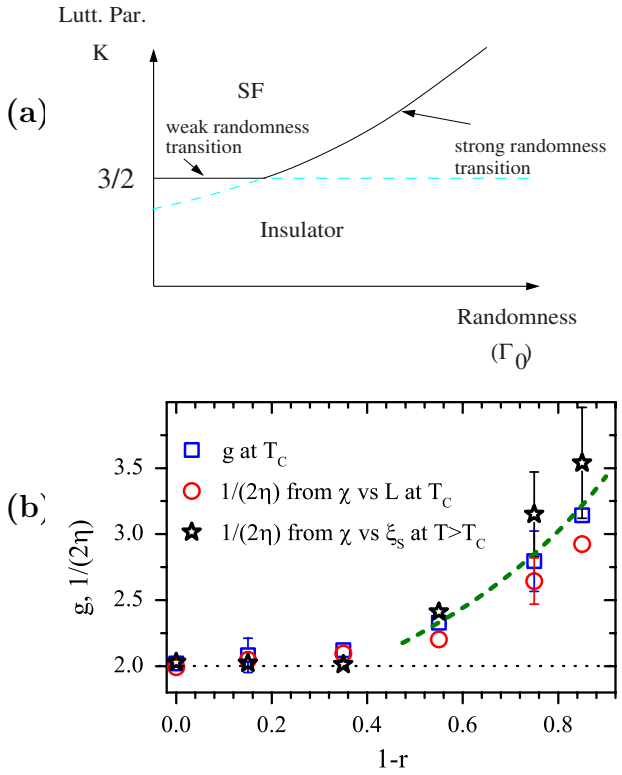

Figure 2.5: In panel (a), a schematic depiction of the crossover between weak and strong-disorder transitions, as depicted in a paper by Altman et al. [9]. Note that this figure shows the generic scenario and that, in the particle-hole symmetric case, the weak-disorder value of K is 2, not $\frac{3}{2}$. In panel (b), Monte Carlo results from Hrahsheh and Vojta that are suggestive of the crossover [70]. The quantity $1-r$ is a measure of the initial disorder in Hrahsheh and Vojta's simulations, and and $\frac{1}{2\eta}$ are proxies for the Luttinger parameter.

the strong-disorder transition may take over when the Luttinger parameter (2.36) exceeds the Giamarchi-Schulz value. Panel (a) of Figure 2.5 shows this scenario schematically, while panel (b) presents the numerical results of Hrahsheh and Vojta that seem to support the existence of the two transitions [9, 70].

What would this imply for the physics driving the transition? The transition of Giamarchi and Schulz is driven by phase slips in the superfluid order parameter. The strong-disorder transition therefore occurs in the regime in which phase slips are irrelevant, and another process must be driving the transition. Altman et al. interpreted

the site decimation step as the formation of a "phase-slip dipole," consisting of a phase slip and anti-phase slip pair, on adjacent bonds. They then reasoned that the strong-disorder transition to the insulating state is driven by the proliferation of these dipoles, as opposed to the proliferation of individual phase slips. As such, the question of whether the strong-disorder transition exists is a very important, fundamental question about how superfluidity can be destroyed in a disordered system [7, 8, 9].

2.3.3 Insulating Phases

Let us momentarily consider the impact of introducing random offsets to the filling of the disordered rotor model:

$$\hat{H}_{\rm rot} = -\sum_{\langle jk \rangle} J_{jk} \cos(\hat{\phi}_j - \hat{\phi}_k) + \sum_j U_j (\hat{n}_j - \bar{n}_j)^2 \tag{2.37}$$

Altman et al. found that the character of the superfluid-insulator transition is *independent* of the properties of the random offsets $\bar{n}_j$. It is always of Kosterlitz-Thouless type, with a Luttinger parameter that depends on the bare disorder strength. However, they found that the symmetry properties of the $\bar{n}_j$ are crucial for identifying the glassy Griffiths phase that intervenes between the superfluid and Mott insulator. Before proceeding to calculate the susceptibility at criticality, we will discuss the various glassy phases that can be realized.

Consider the structure of the on-site charging spectrum:

$$E_j(n_j) = U_j (n_j - \bar{n}_j)^2 \tag{2.38}$$

The integer value of n_j that minimizes this energy changes at half-integer values of $\bar{n}_j$. Suppose $\bar{n}_j$ are randomly distributed in the interval $[-\frac{1}{2}, \frac{1}{2})$[6]. Then, the introduction of an arbitrarily small global chemical potential shift will bring a finite fraction of sites j arbitrarily close to these density changing points. Thus, a finite density of particles will be added to the system, making it compressible. Nevertheless, these particles will be localized by the disordered environment, leaving the system globally insulating. This is the mechanism behind the formation of the Bose glass [55].

The situation changes at two special particle-hole symmetric points. For example, consider the case where $\bar{n}_j$ is restricted to be integer ($\bar{n}_j = 0$) or half-integer ($\bar{n}_j = \frac{1}{2}$). Then, at the half-integer sites, there is a degeneracy: $E_j(0) = E_j(1)$. Decimation of a half-integer site thus leaves two possibilities for the local particle number, and this can be treated as an effective spin-$\frac{1}{2}$ degree of freedom. These "spins" can form singlets over arbitrarily large distances; physically, this corresponds to delocalization of a boson between two distant sites. The resulting glassy phase was termed a random-singlet glass by Altman et al., in analogy to the phase found in spin systems by Fisher [52, 8]. Like the Bose glass, the random-singlet glass has a diverging superfluid susceptibility (1.21). It differs from the Bose glass in that the compressibility (1.20) also diverges with system size.

By setting all $\bar{n}_j = 0$, we exclude the possibility of these effective spins and return to our original particle-hole symmetric model (2.9). Now, if the U_j are distributed

[6]There are two points to note here. First, it is not necessary to consider offsets outside of these bounds, because offsets of greater magnitude can be removed by a redefinition of the local filling. Thus, considering offsets within the specified interval is generic. We also should consider the case where the offsets are randomly distributed in a subset of the interval $[-\frac{1}{2}, \frac{1}{2})$. Then, cluster formation will induce effective offsets that span the whole interval, again leading to a Bose glass via the mechanism described in the text.

such that $U_{\min} > 0$, the on-site charging spectrum always has a unique minimum that is protected by a gap. A cluster needs to become infinitely large to have a vanishing local gap. The resulting insulating phase has only a finite susceptibility and a *vanishing* compressibility in the thermodynamic limit. Nevertheless, it is gapless because arbitrarily large clusters can appear in the thermodynamic limit. This exotic gapless, incompressible phase has been termed a *Mott glass* [7]. Such a phase was first found in disordered fermionic systems by Giamarchi, Le Doussal, and Orignac [57], but the work of Altman et al. indicated that it can be found in bosonic systems as well.

Thus, there are three possible glassy phases in the one-dimensional dirty boson problem: arranged in order of increasing "proximity" to superfluidity, they are the Mott glass, the Bose glass, and the random-singlet glass. The Mott glass is most directly relevant to the work that we will report in Chapter 3. We direct the reader to the work of Altman et al. for a more complete description of scaling behavior in all three glassy phases [8].

2.4 Calculation

We now proceed to describe our calculation of the superfluid susceptibility (1.21) at the critical point of the disordered, particle-hole symmetric rotor model (2.9).

2.4.1 General Structure of the Calculation

Our general strategy for calculating the susceptibility is based on the following intuition: the SDRG builds clusters of various sizes (through link decimation steps) and then removes these clusters from the chain (through site decimation steps). Each cluster that is removed from the chain behaves as an independent superfluid island that contributes to the susceptibility (1.21). Our goal is to keep track of all the contributions to χ as the RG proceeds.

This approach motivates writing the formula:

$$\chi = \frac{1}{L}\left[X_f + \int_0^{\Gamma_f} d\Gamma \rho(\Gamma) X_{\text{clust}}(\Gamma)\right] \tag{2.39}$$

Here, $\rho(\Gamma)d\Gamma$ is the density of clusters that are decimated at a scale Γ and Γ_f is the RG time necessary to renormalize an L-site chain to a single site. The quantity $X_{\text{clust}}(\Gamma)$ is the *extensive* susceptibility[7] of these decimated clusters. Thus, the integral accounts for the contribution of all but one superfluid cluster to χ. We must take that final cluster into account separately, and this contribution is represented as X_f[8].

To use equation (2.39), we need a way of calculating $X_{\text{clust}}(\Gamma)$. One way to do this is to track the evolution of the perturbation (1.22) under the RG, and then evaluate the susceptibility of the isolated cluster at the moment of site decimation[9]. This

[7] The extensive susceptibility is the susceptibility that we introduced in equation (1.21) but without normalization by the system size.

[8] We will see below that the final cluster needs to be treated rather differently from the others. This ultimately stems from the fact that the analytical RG does not know anything about the finite size of the system.

[9] We will use a variant of this approach to calculate the susceptibility numerically in Chapter 3. See Appendix 3.B for details.

implies $X_{\rm clust}(\Gamma) \sim \frac{\bar{s}(\Gamma)^2}{\Omega}$. Such a calculation neglects quantum fluctuations *within* the cluster, which are expected to be very important in 1D. It is, therefore, preferable to work with a Hamiltonian that describes the internal degrees of freedom and directly calculate the susceptibility. Since this Hamiltonian will itself be disordered, this could be as challenging as calculating the susceptibility of the full chain. There is, however, one possible simplification: the absence of phase slips within the cluster implies that the effective cluster Hamiltonian is quadratic. A *uniform* quadratic Hamiltonian would be easy to treat analytically, and if we could devise an appropriate uniformization procedure, we could analytically approximate $X_{\rm clust}(\Gamma)$. We will follow such an approach below. At first glance, this method may seem fundamentally contradictory to the assumption of strong disorder. In fact, it is only necessary to assume that X is a self-averaging quantity, not that the disorder is weak. Numerical checks[10] on systems of size $L = 100$ suggest that this assumption is valid, but for very strong disorder, the self-averaging may require a prohibitively large amount of statistical sampling. We will elaborate on this point later. It is important to keep in mind that the uniformization procedure is an uncontrolled approximation whose ultimate justification is that it produces a result that is consistent with the numerical observations of Hrahsheh and Vojta.

At various stages of this calculation, it will be useful to check our progress by comparison to a numerical implementation of the strong-disorder renormalization group[11]. In this chapter, we always start our numerical RG with distributions that

[10] In these checks, we compare numerical predictions of correlation functions in the disordered, quadratic Hamiltonian to the analogous predictions in the uniformized model.

[11] This is the same numerical implementation that we use to study the 2D problem in Chapter 3.

have the universal forms (2.29) and (2.30). Furthermore, we focus exclusively on critical flows with $\epsilon = 0$ and tune the initial value of g_0 to explore how quantities depend nonuniversally on the bare disorder strength. We study lattices of sizes $L = 1000$ to $L = 32000$ and always pool 100 realizations of the randomness[12].

2.4.2 Ingredients

Here, we describe the preparatory work that is needed to piece together the calculation (2.39).

2.4.2.1 Critical Flows, the Cluster Density $\rho(\Gamma)$, and the Renormalization Time Γ_f

We begin by calculating the critical flows $f_0(\Gamma)$ and $g_0(\Gamma)$. We imagine beginning at $\Gamma = 0$ somewhere on the critical manifold in Figure 2.4 and near to the unstable fixed point. Using $\epsilon = 0$ in equation (2.33) and expanding $g_0(\Gamma) = 1 + y(\Gamma)$, we can find the flow equation:

$$\frac{dy}{d\Gamma} \approx -\frac{1}{2}y^2 \tag{2.40}$$

This differential equation is easily solved:

$$y(\Gamma) \approx \frac{2}{\Gamma + \frac{2}{y_i}} \equiv \frac{2}{\gamma(\Gamma)} \tag{2.41}$$

[12]In Chapter 3, we will do far more sampling to get better statistics, but here we only do numerics to verify our theoretical predictions.

where $y(0) = y_i$ is the initial deviation of g_0 from one, and we have implicitly defined the "shifted" RG time:

$$\gamma(\Gamma) \equiv \Gamma + \frac{2}{y_i} \tag{2.42}$$

Translating back to the flows in terms of $f_0(\Gamma)$ and $g_0(\Gamma)$:

$$f_0(\Gamma) \approx \frac{2}{\gamma(\Gamma)^2} \tag{2.43}$$

$$g_0(\Gamma) \approx 1 + \frac{2}{\gamma(\Gamma)} \tag{2.44}$$

Using the solutions (2.43) and (2.44), we can compute the number of effective sites in the chain at a time Γ:

$$\begin{aligned} N(\Gamma) &= L \exp\left[-\int_0^\Gamma d\Gamma'(f_0(\Gamma) + g_0(\Gamma))\right] \\ &\approx L \exp\left[-\Gamma - 2\ln\left(1 + \frac{y_i\Gamma}{2}\right)\right] \\ &= L\frac{4e^{\frac{2}{y_i}}}{y_i^2}\frac{e^{-\gamma(\Gamma)}}{\gamma(\Gamma)^2} \end{aligned} \tag{2.45}$$

We can then reason that, in the time interval $(\Gamma, \Gamma + d\Gamma)$, a proportion $f_0(\Gamma)$ of the effective sites get decimated. Thus, the *density* of clusters decimated in this time interval is:

$$\begin{aligned} \rho(\Gamma)d\Gamma &= \frac{f_0(\Gamma)N(\Gamma)}{L}d\Gamma \\ &\approx \frac{8e^{\frac{2}{y_i}}}{y_i^2}\frac{e^{-\gamma(\Gamma)}}{\gamma(\Gamma)^4}d\Gamma \end{aligned} \tag{2.46}$$

Finally, using equation (2.45), we can deduce how long the RG needs to be run to renormalize the entire chain to a single site by setting $N(\Gamma_f) = 1$. This yields an expression that can be iteratively inverted to yield:

$$\gamma(\Gamma_f) \approx \ln\left(\frac{4e^{\frac{2}{y_i}}}{y_i^2}L\right) - 2\ln\ln\left(\frac{4e^{\frac{2}{y_i}}}{y_i^2}L\right) \tag{2.47}$$

2.4.2.2 Internal Variables: Cluster Sizes and Internal Coupling Sums

In order to calculate the susceptibility, we will also need to compute some "internal" properties of the superfluid clusters formed by the RG. For instance, we will need to know the typical size of clusters that are decimated at RG time Γ. Moreover, to implement the uniformization procedure for the cluster susceptibility, we will also need to assign a uniformized charging energy and Josephson coupling to each cluster; this requires keeping track of internal couplings in the cluster. We will now show how these quantities can be calculated analytically in the SDRG.

To assign a uniformized charging energy to a cluster of size s, we simply ask which uniform value of U satisfies:

$$\frac{1}{U_{\text{unif}}} \equiv \frac{1}{s}\sum_{j\in c}\frac{1}{U_j} \tag{2.48}$$

where the sum runs over all bare sites that belong to the cluster. For clusters at the RG scale, the sum on the right-hand side of equation (2.48) is simply $\frac{1}{\Omega}$. There is no variance in this sum at all, because having effective charging energy Ω is precisely

the condition for being decimated. Hence:

$$U_{\text{unif}} = s\Omega \tag{2.49}$$

for clusters at the RG scale. It should be noted that this is essentially the inverse compressibility of one of these clusters:

$$\kappa = \frac{1}{2U_{\text{unif}}} \tag{2.50}$$

The uniformized Josephson coupling for a cluster is defined as the average:

$$\frac{1}{J_{\text{unif}}} \equiv \frac{1}{s}\sum_{j\in c}\frac{1}{J_j} \equiv \frac{1}{s}R \tag{2.51}$$

where the sum runs over all Josephson couplings that are "internal" to the cluster. In other words, these couplings were those involved in the link decimations that built the cluster[13]. This definition makes J_{unif} equivalent to the superfluid stiffness of the cluster:

$$\rho_s = J_{\text{unif}}. \tag{2.52}$$

We would like to keep track of the distribution of the sum R in equation (2.51) as the RG proceeds. This requires extending the RG to follow the distribution $\tilde{f}(\zeta, R, \Gamma)$. Altman et al. wrote an integrodifferential equation that describes the evolution of

[13]These Josephson couplings may themselves have been effective couplings that were generated by one or more previous site decimations.

this distribution:

$$\begin{aligned}\frac{d\tilde{f}}{d\Gamma} &= (1+\zeta)\frac{\partial \tilde{f}}{\partial \zeta} + \tilde{f}(\zeta, R, \Gamma)(f_0(\Gamma) + 1 - g_0(\Gamma)) \\ &\quad + g_0(\Gamma)\int d\zeta_1 d\zeta_2 dR_1 dR_2 \tilde{f}(\zeta_1, R_1, \Gamma)\tilde{f}(\zeta_2, R_2, \Gamma) \\ &\quad \times \delta(\zeta - \zeta_1 - \zeta_2 - 1)\delta\left(R - R_1 - R_2 - \Omega^{-1}\right) \end{aligned} \tag{2.53}$$

They then solved this equation for the quantity:

$$\bar{R}(\Gamma) = \int_0^\infty d\zeta \int_0^\infty dR R\tilde{f}(\zeta, R, \Gamma) \tag{2.54}$$

which is an average over clusters of all charging energies. Since our calculation of the susceptibility focuses on processes occurring at the RG scale, we are instead interested in the average:

$$\bar{R}(\zeta, \Gamma) = \frac{\int dR R\tilde{f}(\zeta, R, \Gamma)}{\int dR \tilde{f}(\zeta, R, \Gamma)} \equiv \frac{Q_R(\zeta, \Gamma)}{f(\zeta, \Gamma)} \tag{2.55}$$

evaluated at $\zeta = 0$.

We begin by focusing on the quantity $Q_R(\zeta, \Gamma)$, which was implicitly defined on the right-hand side of equation (2.55). We can obtain an equation governing the evolution of $Q_R(\zeta, \Gamma)$ by differentiating the numerator of equation (2.55) and inserting equation

(2.53):

$$\begin{aligned}\frac{dQ_R}{d\Gamma} &= (1+\zeta)\frac{\partial Q_R}{\partial \zeta} + Q_R(\zeta,\Gamma)(f_0(\Gamma)+1-g_0(\Gamma)) \\ &\quad +2g_0(\Gamma)\int_0^\zeta d\zeta' Q_R(\zeta',\Gamma)f(\zeta-\zeta',\Gamma) \\ &\quad +\frac{g_0(\Gamma)}{\Omega}\int_0^\zeta d\zeta' f(\zeta',\Gamma)f(\zeta-\zeta',\Gamma) \end{aligned} \tag{2.56}$$

Now, using the definition in (2.55) to write $Q_R(\zeta,\Gamma) = f(\zeta,\Gamma)\bar{R}(\zeta,\Gamma)$, we can write:

$$\begin{aligned}\frac{d\bar{R}}{d\Gamma} &= (1+\zeta)\frac{\partial \bar{R}}{\partial \zeta} + \bar{R}(\zeta,\Gamma)\left[f_0(\Gamma)+1-g_0(\Gamma)+(1+\zeta)\frac{\partial \ln f}{\partial \zeta} - \frac{\partial \ln f}{\partial \Gamma}\right] \\ &\quad +2f_0(\Gamma)g_0(\Gamma)\int_0^\zeta d\zeta' \bar{R}(\zeta',\Gamma) + \zeta\frac{f_0(\Gamma)g_0(\Gamma)}{\Omega} \end{aligned} \tag{2.57}$$

We now perform a power series expansion of $\bar{R}(\zeta,\Gamma)$ in ζ:

$$\bar{R}(\zeta,\Gamma) = R_0(\Gamma) + \zeta R_1(\Gamma) + \frac{\zeta^2}{2}R_2(\Gamma) + \cdots \tag{2.58}$$

Plugging the expansion (2.58) into (2.57), we can write the differential equations:

$$\frac{dR_0}{d\Gamma} = R_1 \tag{2.59}$$

$$\frac{dR_1}{d\Gamma} = R_1 + R_2 + f_0(\Gamma)g_0(\Gamma)R_0 + \frac{f_0(\Gamma)g_0(\Gamma)}{\Omega_i}e^\Gamma \tag{2.60}$$

$$\frac{dR_2}{d\Gamma} = 2R_2 + R_3 \tag{2.61}$$

and so on. Note that the initial condition is $R_j(0) = 0$ for all j, because at time $\Gamma = 0$, all sites are bare. Therefore, $R_j(\Gamma) = 0$ for *all* Γ for all $j \geq 2$. Remarkably, the power series *terminates*. We can therefore solve equations (2.59) and (2.60) to find:

$$\bar{R}(\zeta, \Gamma) \approx \frac{e^{\Gamma}}{\Omega_i}(1+\zeta)\left[\exp\left(y_i - \frac{2}{\gamma(\Gamma)}\right) - 1\right] \tag{2.62}$$

The expression (2.62) implies a relationship between $\bar{R}(0, \Gamma)$ and $\Omega = \Omega_i e^{-\Gamma}$: we use our numerical implementation of the RG to check this relationship in panel (a) of Figure 2.6. The plot shows qualitative agreement. However, the analytical prediction clearly overestimates the numerical result.

The reason for this discrepancy reveals an important aspect of the SDRG. The clusters that are represented in panel (a) of Figure 2.6 are the largest clusters in their respective lattices, because for small values of y_i and on finite lattices, smaller clusters rarely get decimated. They do form at intermediate stages of the RG, but they almost always get absorbed into the largest cluster. This largest cluster only gets decimated at the very end of the numerical procedure, when it is the only cluster remaining. Nevertheless, from the perspective of the analytical RG this cluster is in fact embedded in an infinite lattice in which *still larger* clusters exist. This assumption can distort statistical predictions from their numerical reality on finite-size systems. Meanwhile, this reasoning implies that the relationship (2.62) should hold quantitatively for the smaller clusters that form in intermediate stages of the RG. We check this in panel (b) of Figure 2.6 and find that the analytical relationship holds well. We will return to the question of how to estimate $\bar{R}$ for the final cluster in

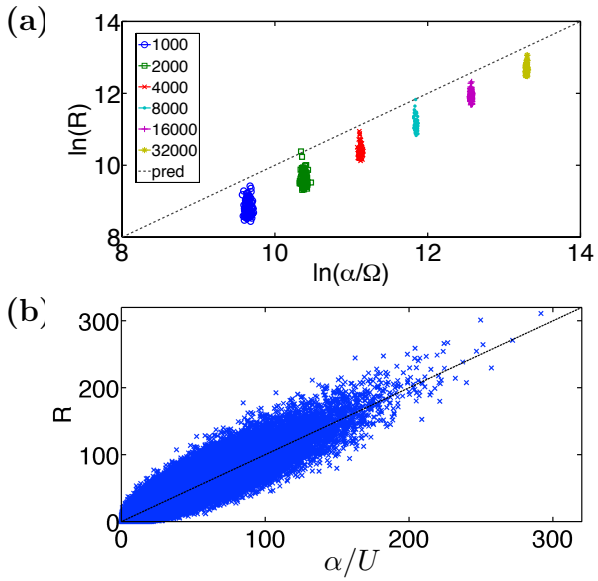

Figure 2.6: Check of the result (2.62) for $\bar{R}(\zeta,\Gamma)$. In panel (a), we monitor the local gap (which is equivalent to the RG scale Ω) and value of R for all clusters decimated in the numerical RG for $y_i = 0.01$. Note that $\alpha = \left[\exp\left(y_i - \frac{2}{\gamma(\Gamma)}\right) - 1\right]$, and the legend in panel (a) refers to different values of the lattice size L. There is a factor of 2 discrepancy with the theoretical prediction, which is explained in the text. In panel (b), we stop the numerical RG on an $L = 32000$ lattice after 28800 RG steps. Then, $\bar{R}$ closely follows the theoretical prediction as a function of ζ, or equivalently the charging energy U.

Section 2.4.2.3 below.

It is also important to consider the variance of the sum R for clusters at the RG scale, because this will, in turn, tell us whether it makes sense to assign all of these clusters the same uniformized Josephson coupling (2.51). To evaluate this variance, we need to compute the second moment of the distribution:

$$u_R(\zeta,\Gamma) \equiv \bar{R^2}(\zeta,\Gamma) = \frac{\int dR R^2 f(\zeta,R,\Gamma)}{\int dR \tilde{f}(\zeta,R,\Gamma)} \tag{2.63}$$

Proceeding along the same lines as above, we can construct an equation that governs

the evolution of u_R:

$$\begin{aligned}\frac{du_R}{d\Gamma} &= (1+\zeta)\frac{\partial u_R}{\partial \zeta} + u_R(\zeta,\Gamma)\left[f_0(\Gamma)+1-g_0(\Gamma)+(1+\zeta)\frac{\partial \ln f}{\partial \zeta}-\frac{\partial \ln f}{\partial \Gamma}\right] \\ &+f_0(\Gamma)g_0(\Gamma)\left(2\int_0^\zeta d\zeta' u_R(\zeta',\Gamma)+2\int_0^\zeta d\zeta' \bar{R}(\zeta',\Gamma)\bar{R}(\zeta-\zeta',\Gamma)\right) \\ &+\frac{4}{\Omega}f_0(\Gamma)g_0(\Gamma)\int_0^\zeta d\zeta' \bar{R}(\zeta',\Gamma)+\zeta\frac{1}{\Omega^2} \end{aligned} \tag{2.64}$$

The power series expansion for $u_R(\zeta,\Gamma)$ terminates at second order in ζ. Numerical integration of the differential equations obtained from the power series expansion[14] suggests that the standard deviation of $R(0,\Gamma)$ settles to a small fraction of the mean. Hence, the $\ln \bar{R}(0,\Gamma)$ is very well defined.

A simpler calculation of the type outlined above yields the behavior of the mean cluster size as a function of ζ and Γ:

$$\bar{s}(\zeta,\Gamma) \approx \frac{y_i^2}{2}(1+\zeta)e^\Gamma \tag{2.65}$$

In Figure 2.7, we use the numerical RG to check this result and find excellent agreement. Constructing the analog of equation (2.64), we can furthermore show that the standard deviation of $s(0,\Gamma)$ settles to a small fraction of the mean at late Γ. As such, there is also a very well-defined relationship between RG time Ω and $\ln \bar{s}(0,\Gamma)$. Indeed, this is consistent with the numerical observations in Figure 2.7 as the numerical results more tightly cluster around the theoretical prediction for larger lattice sizes. Furthermore, for $\bar{s}$, the analytical prediction holds even for all clusters, subleading

[14]These are the analogs of equations (2.59)-(2.60).

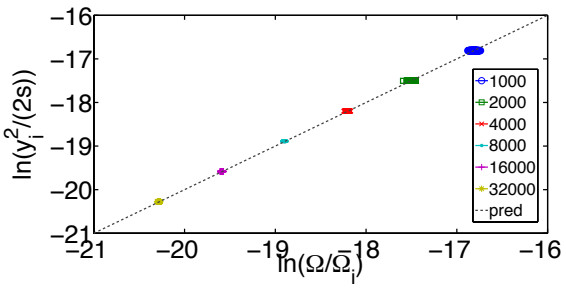

Figure 2.7: A check of the relationship (2.65) between $\bar{s}(0,\Gamma)$ and Ω. We monitor the local gap (which is equivalent to the RG scale Ω) and the size all clusters decimated in the numerical RG for $y_i = 0.01$. The legend refers to different values of the lattice size L, and the dashed line follows the theoretical prediction.

and final. At least for the subleading clusters, the relationships (2.62) and (2.65) can be used to find:

$$J_{\text{unif}}(\Gamma) = \frac{\bar{s}(0,\Gamma)}{R(0,\Gamma)} \tag{2.66}$$

$$U_{\text{unif}}(\Gamma) = \bar{s}(0,\Gamma)\Omega \tag{2.67}$$

2.4.2.3 Adjustments for the Final Cluster

We saw above that our prediction for $\bar{R}(\zeta,\Gamma)$ qualitatively matches the behavior of all decimated clusters. However, it *quantitatively* matches only the subleading clusters, while there is a discrepancy for the final cluster (see Figure 2.6). Here, we show how to fix our prediction $\bar{R}$ for this final cluster.

As a preliminary step, we ask how many sites of the bare chain get incorporated into the final cluster. The fraction of bare sites that are decimated as parts of subleading clusters can be easily calculated via the integral:

$$f = \int_0^{\Gamma_f} d\Gamma \rho(\Gamma)\bar{s}(0,\Gamma) = \int_{\frac{2}{y_i}}^{\gamma(\Gamma_f)} d\gamma \rho(\Gamma)\bar{s}(0,\Gamma) \tag{2.68}$$

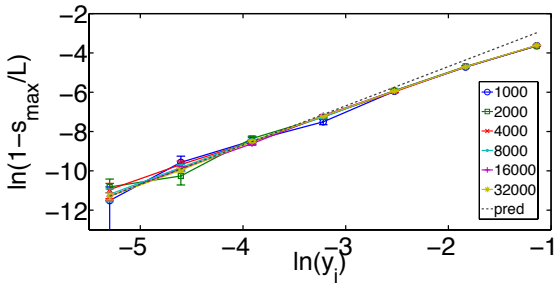

Figure 2.8: A check of the theoretical prediction (2.69) of the fraction of the lattice included in the largest cluster. We plot the fraction *not* covered by this cluster versus y_i for various system sizes. The legend refers to different values of the lattice size L, and the dashed line follows the theoretical prediction.

Plugging in equation (2.65)[15], we find that:

$$f \approx \frac{1}{2} y_i^2 \tag{2.69}$$

Therefore, the final cluster must include:

$$s_{\max} = (1-f)L \approx \left(1 - \frac{1}{2} y_i^2\right) L \tag{2.70}$$

sites. We use the numerical RG to check that this approximates the maximal cluster size well in Figure 2.8 and find good agreement.

The key point here is that, in the small y_i limit, the largest cluster includes nearly all bare sites. Hence, the average of $\frac{1}{J}$ for bonds within this cluster can be approximated by the mean in the bare distribution, and then this can be inverted to

[15]In fact, we also use the fact that $\bar{s}(0,\Gamma)$ stays pinned at 1 until a time of order $\ln\left(\frac{2}{y_i^2}\right)$ until it begins growing according to equation (2.65). This turns out to be important for the numerical prefactor multiplying y_i^2 in equation (2.69).

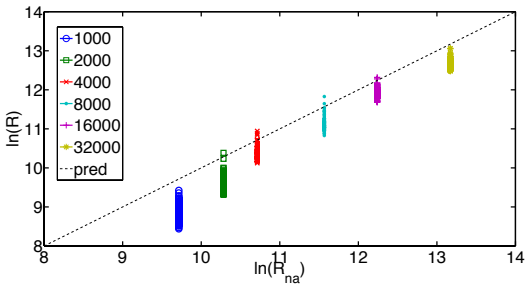

Figure 2.9: A comparison of R for clusters decimated in the numerical RG vs. $R_{\rm na}$ estimated by naive averaging over the bare distribution of $\frac{1}{J}$. Here, $y_i = 0.01$, and in both the numerical RG and the naive averaging, we use 100 samples. The legend refers to different values of the lattice size L, and the dashed line follows the theoretical prediction.

yield $J_{\rm unif}$ for the cluster:

$$J_{\rm unif} \approx \frac{\Omega_i y_i}{1+y_i} \tag{2.71}$$

If we wish, we can also extract the corresponding value of $\bar{R}$ using equations (2.71) and (2.70):

$$\bar{R} \approx \frac{1+y_i}{\Omega_i y_i}\left(1-\frac{1}{2}y_i^2\right)L \tag{2.72}$$

In interpreting both (2.71) and (2.72), it is important to note that the bare distribution of $\frac{1}{J}$ is a power law with a very long tail at high values of $\frac{1}{J}$: $P(\frac{1}{J}) \propto \left(\frac{1}{J}\right)^{-2-y_i}$. Consequently, the average of this quantity will converge very slowly to its mean. Therefore, to check the approximation (2.72), we directly sample from the bare distribution, naively average, and compare to the same data from numerical RG that we used in Figure 2.6. We show this in Figure 2.9, where we can see that the quantitative agreement has improved.

We can also naively average the bare charging energy distribution to find $U_{\rm unif}$:

$$U_{\rm unif} \approx \frac{\Omega_i}{1+\frac{2}{y_i^2}} \tag{2.73}$$

This quantity, however, agrees quantitatively with the predictions of Section 2.4.2.2.

2.4.2.4 The Cluster Susceptibility $X_{\rm clust}(\Gamma)$

We now adopt the uniformization perspective to calculate the susceptibility of a cluster decimated at some stage of the RG. We can write a Hamiltonian for each of these clusters:

$$H_{\rm clust} = \frac{J_{\rm unif}}{2}\sum_{j=1}^{\bar{s}}(\hat{\phi}_j - \hat{\phi}_{j+1})^2 + U_{\rm unif}\sum_{j=1}^{\bar{s}}\hat{n}_j^2 \tag{2.74}$$

Note that we have expanded the cosine in the Josephson coupling term because the link decimation procedure implies the absence of phase slips within the cluster. The Hamiltonian (2.74) is intended to represent some set of bare sites which are effectively isolated from the rest of the lattice, and therefore, we ought to impose open boundary conditions on this subsystem. However, the advantage of the uniformization procedure is that it restores translational invariance; for analytical convenience, we will use periodic boundary conditions under the assumption that the difference does not matter for sufficiently large clusters.

Our goal is now to calculate the ground-state susceptibility (1.21) for the uniformized model (2.74). This can be achieved by adding the perturbation (1.22) and

then constructing the path integral for the zero-temperature partition function:

$$Z_{\text{clust}} = \int D\phi e^{-S_{\text{clust}}} \tag{2.75}$$

where the action S_{clust} is:

$$S_{\text{clust}} = \sum_{j=1}^{\bar{s}} \int_0^\infty d\tau \left[\frac{J_{\text{unif}}}{2}(\phi_{j+1}(\tau) - \phi_j(\tau))^2 + \frac{1}{4U_{\text{unif}}}(\partial_\tau \phi_j)^2 - h\cos\phi_j(\tau) \right] \tag{2.76}$$

We can then apply the definition of the superfluid susceptibility (1.21) to find:

$$X_{\text{clust}} = \frac{1}{2}\sum_{j=1}^{\bar{s}}\sum_{m=1}^{\bar{s}} \int_{-\infty}^{\infty} d\tau \langle e^{i(\phi_j(0)-\phi_m(\tau)} \rangle_{h=0} \tag{2.77}$$

where the expectation value is taken in the unperturbed Hamiltonian. Standard techniques lead to:

$$\begin{aligned} -\ln\langle e^{i(\phi_j(0)-\phi_m(\tau)}\rangle &= \frac{2U_{\text{unif}}}{\pi\bar{s}} \int_0^\infty d\omega \frac{1-\cos(\omega\tau)}{\omega^2} \\ &\quad + \frac{1}{\pi\bar{s}} \sum_{q\neq 0} \int_0^\infty d\omega \frac{1-\cos\left(\frac{2\pi q}{\bar{s}}(j-m)+\omega\tau\right)}{2J_{\text{unif}}(1-\cos\left(\frac{2\pi q}{\bar{s}}\right)) + \frac{1}{2U_{\text{unif}}}\omega^2} \end{aligned} \tag{2.78}$$

The second sum runs over values of q from $-Q$ to Q where $Q \equiv \frac{\bar{s}-1}{2}$, excluding $q = 0$ (which is taken care of in the first sum). The discrete sum over momenta q is due to the finite size of the cluster. The $q = 0$ term in this sum must be treated separately, because it is responsible for the exponential decay of correlations in imaginary time. For the remaining momenta, the discreteness is not so important (at least for large

clusters) and the sum may be replaced by an integral. This yields:

$$-\ln\left\langle e^{i(\phi_j(0)-\phi_m(\tau)}\right\rangle \approx \frac{U\tau}{\bar{s}} + \eta \ln\left(\pi e^{\gamma_E}\sqrt{(j-m)^2 + 2U_{\text{unif}}\tau^2}\right) \tag{2.79}$$

where $\gamma_E \approx 0.5772$ is the Euler gamma and :

$$\eta \equiv \frac{1}{2\pi}\sqrt{\frac{2U_{\text{unif}}}{J_{\text{unif}}}} = \frac{1}{2\pi}\sqrt{\frac{1}{\rho_s \kappa}} \tag{2.80}$$

Finally, we can obtain the cluster susceptibility when η is small:

$$X_{\text{clust}} \approx \frac{\left(\pi e^{\gamma_E}\right)^{-\eta}\bar{s}^{3-\eta}}{U_{\text{unif}}}\left[1 + \eta\ln\eta + \eta\ln\left(\pi e^{\gamma_E}\right)\right] \tag{2.81}$$

Equation (2.81) is the first hint of the role of η (2.80). Roughly, when we apply the formula (2.39) to calculate the susceptibility of the full chain, we will find that the exponent of L in the result will inherit the exponent of $\bar{s}$ in equation (2.81).

2.4.3 Piecing Together the Susceptibility

Using equations (2.71), (2.73), and (2.80, we can identify η for the final cluster. This coincides with the value of η_{sd} that we reported at the beginning of the chapter in equation (2.2). Meanwhile, using results from Section 2.4.2.2, we can derive the value of η for subleading clusters as a function of Γ:

$$\eta(\Gamma) \approx \frac{1}{2\pi}\sqrt{\left[2\left(\exp\left(y_i - \frac{2}{\gamma(\Gamma)}\right) - 1\right)\right]} \tag{2.82}$$

Interestingly, for small y_i, $\lim_{\Gamma\to\infty}\eta(\Gamma)=\sqrt{2}\eta_{\rm sd}$.

First, we compute the final cluster's contribution to the susceptibility. We can use equation (2.77) to find:

$$\chi_f=\frac{X_f}{L}\approx\frac{2}{y_i^2\Omega_i}(\pi e^{\gamma_E})^{-\eta_{\rm sd}}\left(1-\frac{y_i^2}{2}\right)^{3-\eta_{\rm sd}}L^{2-\eta_{\rm sd}}[1+\eta_{\rm sd}\ln(\pi e^{\gamma_E})+\eta_{\rm sd}\ln\eta_{\rm sd}] \quad (2.83)$$

On the other hand, evaluating the contribution of the subleading clusters requires evaluating the integral:

$$\begin{aligned}\chi_{\rm sl} &= \frac{X_{\rm sl}}{L}\\ &= \int_0^{\Gamma_f}d\Gamma\rho(\Gamma)X_{\rm clust}(\Gamma)\\ &\approx \frac{2}{\Omega_i y_i^2}\int_{\frac{2}{y_i}}^{\gamma(\Gamma_f)}d\gamma(\pi e^{\gamma_E})^{-\eta(\gamma)}\left(\frac{y_i^2}{2}\right)^{3-\eta(\gamma)}\frac{e^{(2-\eta(\gamma))\gamma}}{\gamma^4}\\ &\qquad\times[1+\eta(\gamma)\ln(\pi e^{\gamma_E})+\eta(\gamma)\ln\eta(\gamma)] \end{aligned}\quad (2.84)$$

At small y_i, the density of subleading clusters is strongly suppressed. These clusters rarely form and, therefore, have little opportunity to contribute to the susceptibility. In any case, in the thermodynamic limit, the contribution of the final cluster diverges as $L^{2-\eta_{\rm sd}}$, whereas the contribution of subleading clusters diverges as $L^{2-\sqrt{2}\eta_{\rm sd}}$. Hence, the contribution of the final cluster dominates and the anomalous exponent $\eta_{\rm sd}$ is the principal result of our work.

Nevertheless, it may be difficult to cleanly observe the anomalous exponent $\eta_{\rm sd}$ in numerics. The reason was already mentioned in Section 2.4.2.3 above. The distribution of $\frac{1}{J}$ within the cluster is well approximated by the bare Josephson coupling

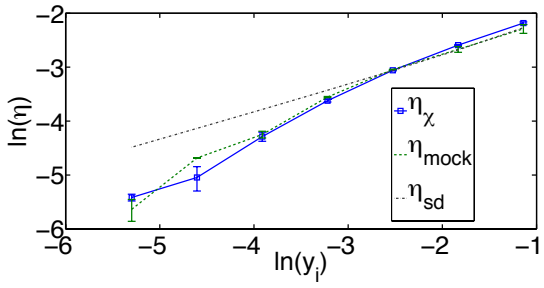

Figure 2.10: The anomalous exponent η_χ extracted from measurements of $\chi(L)$ in a numerical implementation of the SDRG. We explore lattices from $L = 1000$ to $L = 32000$ with various values of y_i at criticality ($\epsilon = 0$). We always average over 100 realizations of the randomness. Two reference lines are included in the plot. The dashed-dotted line follows the theoretical prediction (2.2). The dashed reference line follows the prediction of η from a mock susceptibility calculated by naively averaging the bare coupling distributions and applying equation (2.83). This reference line follows the line from numerical RG closely, suggesting that the discrepancy with the theoretical prediction is due to slow convergence under sampling of the random distributions. See the text for details.

distribution, which is very wide for small y_i. A mean of $\frac{1}{J}$ exists for $y_i > 0$, but estimates of this mean converge very slowly as a function of the number of samples. When averaging over a moderate number of samples, we will typically obtain an underestimate of the mean, leading to an effectively lower value of η. We illustrate this in Figure 2.10. Note that a lower value of the anomalous exponent is actually further from the Giamarchi-Schulz value of $\eta = \frac{1}{4}$ and closer to the scaling result of $\eta = 0$.

2.5 Discussion

It is instructive to compare our result (2.2) to the Luttinger parameter (2.36) found by Altman, Kafri, Polkovnikov, and Refael. To do so, we need to note the correspondence between our measure of initial disorder, y_i, and theirs, Γ_0. In the papers of Altman et al., Γ_0 is essentially the RG time that it takes for the initial

nonuniversal distributions to evolve into their universal forms. Our critical flows (2.43) and (2.44) suggest that the appropriate correspondence is:

$$y_i = \frac{2}{\Gamma_0} \tag{2.85}$$

In fact, if we insert this relationship into the expression for the Luttinger parameter (2.36), we find:

$$\begin{aligned} \frac{1}{2K} &\approx \frac{1}{2\pi}\sqrt{\frac{e^{2y_i}-1}{2}} \\ &\approx \frac{1}{2\pi}\sqrt{y_i} \\ &\approx \eta_{\mathrm{sd}} \end{aligned} \tag{2.86}$$

To leading order, our result thus matches the Luttinger parameter of Altman et al. through $\eta = \frac{1}{2K}$. This relationship was also reflected in the Monte Carlo results of Hrahsheh and Vojta [70].

The agreement of our result with that of Altman et al. is not accidental or even surprising. In calculating the susceptibility, we have essentially used the same machinery that Altman et al. used in their calculation of K. The quantities $J_{\mathrm{unit}}(\Gamma)$ and $U_{\mathrm{unit}}(\Gamma)$ are essentially the superfluid stiffness and inverse compressibility of the clusters being decimated at a scale Γ. The stiffness and compressibility determine the effective Luttinger parameter K, and in a clean system, $\eta = \frac{1}{2K}$ follows from using linear response theory for the susceptibility. Our uniformization procedure naturally guarantees that this relationship is maintained in our result. What is more surprising

is that the relationship between η and K is reflected in the Monte Carlo data of Hrahsheh and Vojta [70]. This implies that the one uncontrolled approximation in our calculation, the uniformization procedure, may be reasonable.

In turn, our work has potentially important implications for interpreting the numerics of Hrahsheh and Vojta. Specifically, as we have noted above, our results suggest that Monte Carlo estimates of η cannot necessarily be taken to reflect the "true" thermodynamic value of the anomalous dimension η_{sd}. Numerical values of this exponent will inevitably be plagued by strong finite sampling corrections. It is worthwhile to get a sense of how large of a system or how much statistical sampling is necessary to obtain an accurate prediction of the anomalous dimension. We do this in Figure 2.11 by asking how many times we need to randomly draw $\frac{1}{J}$ from the bare Josephson coupling distribution for the mean to converge. For $y_i = 0.32$, the mean converges relatively quickly, and on the order of 10 samples of an $L = 3200$ lattice are probably sufficient to observe η_{sd}. For $y_i = 0.04$, the mean once reaches its analytic value, presumably due to the random sampling of an anomalously large value of $\frac{1}{J}$, but it thereafter drops below and does not recover. For $y_i = 0.005$, we would expect considerable underestimation of the mean, even if we were able to study 1000 samples of an $L = 3200$ system, which is on the order of what Hrahsheh and Vojta explore using Monte Carlo [70]. Therefore, we ought to expect underestimation of the critical exponent in the strong-disorder limit, as we saw in Figure 2.10. Naively, it may then seem strange then that the relationship $\eta = \frac{1}{2K}$ is captured reasonably well in the Monte Carlo results, even if we assume that it holds in disordered systems. However,

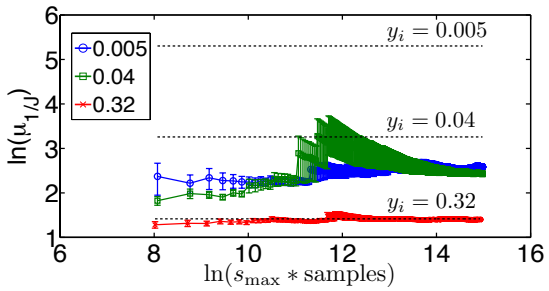

Figure 2.11: A test of the circumstances in which the exponent ηsd can be cleanly observed. We sample $\frac{1}{J}$ from the bare Josephson coupling distribution and follow the convergence of the mean to the analytical value, indicated by the dashed reference lines. We plot data for the three values of y_i given in the legend. See text for details.

the argument above suggests that η and K will be affected by finite sampling results in the same way, suggesting that the relationship between them will be maintained even as these quantities slowly drift to their true values.

2.6 Conclusion

In this chapter, we have presented a detailed calculation of the superfluid susceptibility at the critical point of the one-dimensional disordered rotor model, within the strong-disorder renormalization framework. Our calculation implies that the susceptibility diverges with system size as $\chi \sim L^{2-\eta_{\rm sd}(y_i)}$, where the strong-disorder anomalous dimension $\eta_{\rm sd}(y_i)$ depends upon the bare disorder strength. This prediction contrasts sharply with the situation at the weak-disorder transition of Giamarchi and Schulz, where the anomalous dimension is always the universal Kosterlitz-Thouless value of $\frac{1}{4}$.

In addition to our principal result, we found several other interesting properties of the strong-disorder critical point. For instance, we were able to find an expression,

$\bar{s}(\zeta, \Gamma)$, for the mean size of clusters with any charging energy. Furthermore, we found evidence that the variance of $\bar{s}(0, \Gamma)$ vanishes at late RG time, suggesting a well-defined relationship between the size of superfluid islands and their local charging gap. We were similarly able to find a well-defined relationship between the gap and the superfluid stiffness, at least for the subleading superfluid islands that "dress" the largest cluster.

On the other hand, we also found an important qualification of our principal result: the "true" thermodynamic value of η_{sd} may not be observed on systems that are accessible to numerics. The observed value of η is polluted by a finite sampling error which dies off very slowly due to the wide distribution of internal $\frac{1}{J}$. This distribution gets wider and wider as we tune the transition progressively closer to the unstable fixed point[16]. Hence, this effect must be kept in mind when interpreting Monte Carlo results. Since experiments on ultracold atoms, for instance, work with only moderately sized systems, this difficult could affect interpretation of these experiments as well.

[16]Indeed, as $y_i \to 0$, the mean of $\frac{1}{J}$ diverges.

Chapter 3

Mott Glass to Superfluid Transition for Dirty Bosons in Two Dimensions

3.1 Introduction

3.1.1 Motivation

As we noted in Chapter 2, the strong-disorder renormalization group (SDRG) of Altman, Kafri, Polkovnikov, and Refael has raised the exciting possibility of a strong-disorder transition in the one-dimensional dirty boson problem [7, 9], and recent numerical simulations may be capturing the emergence of this novel criticality [70]. This naturally motivates the question of what the SDRG can tell us about the superfluid-insulator in higher dimensions.

This problem is tantalizing for several reasons: first, experiments are obviously not restricted to one-dimensional systems, so it is imperative to refine our understanding of higher-dimensional bosonic systems. Second, when implemented numerically, SDRG can access larger system sizes than many other methods. Therefore, it can

be a valuable tool in extracting universal properties of the superfluid-insulator transition. Finally, suppose we consider the two-dimensional version of the rotor model that we studied in Chapter 2. The clean version of this model is characterized by an exponent $\nu \approx 0.663$ that *violates* the Harris criterion (1.18) [59]. Hence, in contrast to the one-dimensional case, we expect even weak disorder to be relevant, perhaps indicating the presence of strong-disorder physics.

In this chapter, we numerically extend the SDRG of Altman et al. to study the disordered rotor model in two dimensions. Our work is aimed at resolving two physical questions. Most importantly, we want to understand the character and universal properties of the superfluid-insulator transition. Here, "universal" refers to those properties that are independent of the structure of the microscopic distributions of couplings in the model. We also want to identify the glassy phase that intervenes between the superfluid and the Mott insulator. Intuition from the one-dimensional problem suggests that, in this particle-hole symmetric model, we ought to expect the incompressible Mott glass to appear in place of the Bose glass [7]. We would like to confirm this through the SDRG analysis[1].

Application of the SDRG in $d > 1$ has historically been rare[2]. In part, this is because analytical approaches are usually not tractable. More importantly, there are few known transitions that exhibit infinite randomness, the property that ensures that

[1]Note that the numerical SDRG is a very ideal method for examining the effects of rare regions of superfluid ordering, because it explicitly constructs these regions in real space. Furthermore, because the method can reach larger systems than other numerical techniques, Griffiths effects can be more easily observed.

[2]Examples which do apply the method in $d > 1$ include papers by Motrunich et al. and Kovács and Iglói, which treat the transverse field Ising model [104, 84].

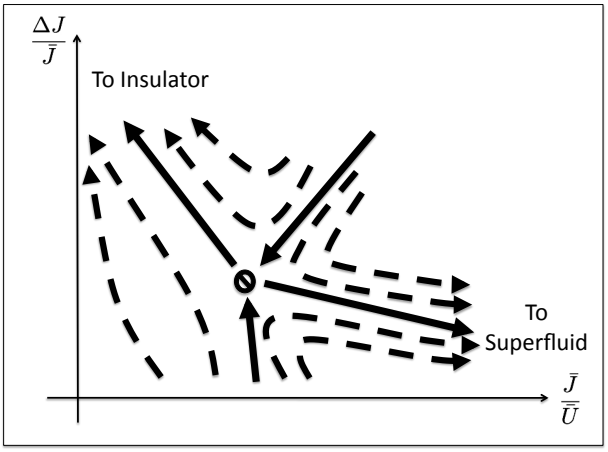

Figure 3.1: A schematization of the universal features of the proposed flow diagram. The x-axis gives the ratio of the mean of the renormalized Josephson coupling distribution to the mean of the renormalized charging energy distribution. The y-axis gives the ratio of the standard deviation of the Josephson coupling distribution to its mean. In this context, the Josephson coupling distribution only includes the dominant $2\tilde{N}$ couplings in the renormalized J distribution, where $\tilde{N}$ is the number of sites remaining in an effective, renormalized lattice. See the text of Section 3.2 for the reasoning behind the exclusion of weaker Josephson couplings from statistics.

the SDRG becomes asymptotically exact near criticality[3]. The superfluid-insulator transition for dirty bosons is not expected to be of the infinite randomness class: indeed, the one-dimensional transition of Altman et al. occurs at finite disorder [7], and as we will see below, the same is true of the transition that we find in two dimensions. Hence, in addition to the physical questions that we posed above, our work also aims to address a methodological question: might the SDRG give useful information about a model, even when confronted with the twin difficulties of higher dimensionality and the absence of infinite randomness?

3.1.2 Preview of the Results

Our main results are as follows: we present numerical evidence for the existence of an unstable finite-disorder fixed point of the RG flow, near which the distributions of Josephson couplings and charging energies in the rotor model flow to universal forms. A schematic picture of this unstable fixed point and the flows in its vicinity is given in Figure 3.1.

To the left of the diagram, flows propagate towards a regime in which the ratio of $\bar{J}$, the mean of the Josephson couplings, to $\bar{U}$, the mean of the charging energies, vanishes; meanwhile, the ratio of ΔJ, the width of the Josephson coupling distribution, to $\bar{J}$ grows very large. These flows terminate in one of two insulating phases. The first is a conventional Mott insulator, in which it is energetically unfavorable for the particle number to fluctuate from the large filling at any site. The other is a glassy phase, in which there exist rare Griffiths regions of superfluid ordering. As the thermodynamic limit is approached, arbitrarily large rare regions appear, driving the gap for charging the system to zero. However, the density of the largest clusters decays exponentially in their size, and the size of the largest cluster in a typical sample does not scale extensively in the size of the system. Moreover, the largest clusters are so rare that they cannot generate a finite compressibility. Thus, the phase is a Mott glass.

This insulating phase gives way to global superfluidity when the rare regions of superfluid ordering percolate, producing a macroscopic cluster of superfluid ordering.

[3] See 1.4.1.3 for a discussion of infinite-disorder fixed points.

The appearance of the macroscopic cluster is associated with flows that propagate towards the lower right of Figure 3.1, indicating that the unstable fixed point governs the glass-superfluid transition. Our numerical implementation of the SDRG allows us to extract estimates for the critical exponents that characterize this transition. We are thus able to construct a compelling picture of the superfluid-insulator transition: a picture that must, however, be checked by other methods because of the perils of employing the SDRG method in the vicinity of a finite-disorder fixed point.

3.1.3 Organization of the Chapter

In Chapter 2, we have already introduced the disordered rotor model and the SDRG method. As such, we begin our two-dimensional study in Section 3.2 by describing the modifications that are necessary to apply the SDRG in 2D. Next, in Section 3.3, we present data collected from our numerical implementation of the SDRG. Our data is suggestive of striking universality in the superfluid-insulator transition, and in Section 3.4, we proceed to use the data to characterize the phase transition and the phases it separates. We then conclude in Section 3.5 by summarizing the results, making connections to experiments, and giving an outlook.

Appendices 3.A-3.D report supplementary material. In Appendix 3.A, we elaborate on an important difference between our numerical implementation of the SDRG and those used in previous studies [4]. Then, in Appendix 3.B, we provide technical details regarding our calculation of physical properties in the SDRG. Appendix 3.C is devoted to details regarding a choice of microscopic distributions that we use exten-

[4]The significance of this difference will become more apparent in Section 3.2

sively in the body of the paper. Finally, in Appendix 3.D, we take up the important tasks of identifying and testing potential weaknesses of the SDRG and comparing to exact diagonalization of small systems.

3.2 Methodology: Numerical Application of the SDRG

In this Section, we describe our numerical application of the SDRG to the two-dimensional disordered rotor model. The fundamentals of the SDRG approach have already been introduced in Chapter 2. Here, we begin by describing the modifications to this procedure that are necessary to move to higher dimensions. Then, we describe an approach for including the effects of quantum fluctuations within clusters. Finally, we collect some details of the initial distributions that we use in our numerical simulations[5].

3.2.1 Adaptations for 2D

Many of the challenges in applying the SDRG in dimensions greater than one originate in the RG's modification of the lattice structure. Suppose we begin with a two-dimensional model on a square lattice. Site and link decimations will immediately destroy this structure locally, as shown in Figures 3.2 and 3.3. Over time, the global square lattice structure will be lost, and the RG will instead yield a more

[5]Throughout this section, we use notation that was introduced in Sections 2.2 and 2.3. It may be helpful to refer back to those sections to recall the conventions.

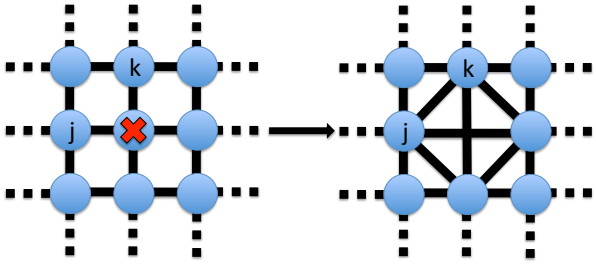

Figure 3.2: The site decimation RG step in 2D: The site marked with the X has the dominant charging energy and is decimated away, generating bonds between neighboring sites j and k. The new local structure of the lattice is shown to the right.

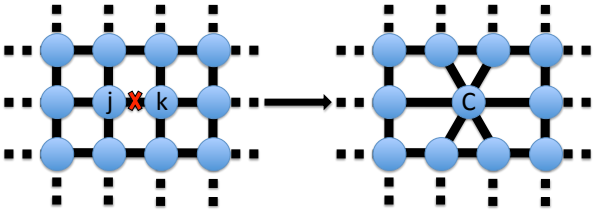

Figure 3.3: The link decimation RG step in 2D: The crossed link has the dominant Josephson coupling. The two sites it joins are merged into one cluster, resulting in the effective lattice structure shown to the right.

general, highly connected network structure. This forces adaptations in the numerical implementation of the SDRG and in the analysis of the data.

3.2.1.1 Sum Rule

As the first example, consider the situation depicted in Figure 3.4: decimation of a site X produces an effective Josephson coupling between two sites j and k that are

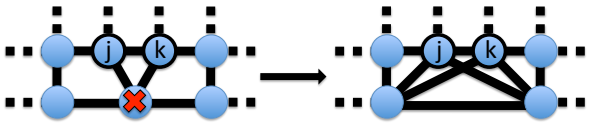

Figure 3.4: Site decimation with the sum rule: The site with the dominant charging energy (marked with an X) is coupled to two sites (j and k) that are already coupled to one another. After site decimation, the effective Josephson coupling between sites j and k is the sum of the old coupling and the effective coupling generated through decimation of site X (see equation (3.1)).

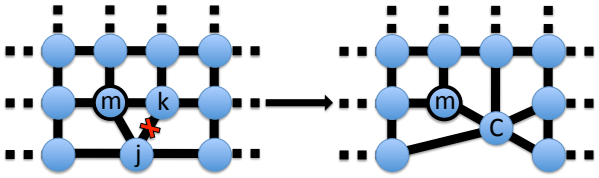

Figure 3.5: Link decimation with the sum rule: The two sites connected by the dominant Josephson coupling (sites j and k) are both coupled to a third site m. Following link decimation, the effective Josephson coupling between site m and the cluster C is the sum of the two preexisting couplings between site m and sites j and k (see equation (3.2)), up to corrections coming rom Debye-Waller factors.

already linked to one another by a preexisting Josephson coupling. How should we treat this situation? In our numerics, we sum the preexisting and new coupling to form the effective coupling between the remaining sites:

$$\tilde{J}_{jk} = J_{jk} + \frac{J_{jX} J_{Xk}}{U_X} \tag{3.1}$$

This choice to sum the preexisting and effective couplings is known as the *sum rule*. The sum rule also may also need to be invoked during link decimation. In Figure 3.5, a cluster is formed by two sites j and k, each of which is connected to a third site m. We choose the effective coupling between the cluster and site m to be:

$$\tilde{J}_{mC} = J_{mj} + J_{mk} \tag{3.2}$$

The sum rule is not always used in implementations of the SDRG. Some authors replace it with a *maximum rule* [104, 84, 53]. The motivation behind the maximum

rule is that, in the strong-disorder limit:

$$\max\left[J_{jk}, \frac{J_{jX}J_{Xk}}{U_X}\right] \approx J_{jk} + \frac{J_{jX}J_{Xk}}{U_X} \tag{3.3}$$

This should be a good approximation in an infinite-disorder context. For our model however, we find that the sum rule increases the class of distributions that find the unstable fixed point depicted schematically in Figure 3.1. For further discussion of the difference between the sum and maximum rules, please consult Appendix 3.A.

3.2.1.2 Thresholding

In dimensions greater than one, there is a tendency for the numerics to slow down considerably if the renormalization procedure involves a lot of site decimation. If a site X is decimated, then effective links are generated between each pair of sites that were previously coupled to X. Thus, site decimation generates many new couplings, increasing the coordination number of the effective lattice. At the same time, the site decimation step takes quadratic time in the coordination number of the site being decimated. To apply the procedure to large lattices, it is necessary to find a way to circumvent this difficulty.

At the beginning of the RG, we specify a thresholding parameter, which we call α. During a site decimation, if a new Josephson coupling is created between sites j and k such that:

$$\tilde{J}_{jk} = \frac{J_{jX}J_{Xk}}{U_X} < \alpha U_X = \alpha\Omega \tag{3.4}$$

then the coupling is thrown away. For convenience in implementation, the new bond is ignored only if it does not sum with a preexisting Josephson coupling. If α is chosen to be very small, then ignoring the coupling will hopefully not affect the future course of RG. However, to be more careful, it is better to perform an extrapolation in the threshold α to see if the numerics converge. Using this thresholding procedure, we are able to reach lattices up to size 300×300 if we additionally require averaging over a reasonably large number of disorder samples. In this chapter, unless otherwise stated, we always use 10^3 samples for any given choice of distributions.

3.2.1.3 Distribution Flows

Typically, in an application of the SDRG, it is interesting to monitor the flow of the distributions of the various couplings as the RG proceeds. As we saw in Chapter 2, this is straightforward for a 1D chain. In higher dimensions, there is yet again a complication from the evolving lattice structure. As the renormalization proceeds, it is possible to generate very highly connected lattices. Many of the effective Josephson couplings will, however, be exceedingly small. Incorporating these anomalously small couplings into the statistics can be misleading. Despite the large number of weak bonds, there may exist a number of strong bonds sufficient to produce superfluid clusters. In fact, including the weak bonds in statistics is analogous to polluting the statistics with the inactive next-nearest neighbor Josephson couplings of the original lattice. It is more appropriate to follow Motrunich et al. and focus on the largest $O(\tilde{N})$ Josephson couplings, where $\tilde{N}$ is the number of sites remaining in the effective lattice [104]. In the remainder of the chapter, the "Josephson coupling distribution"

will therefore refer solely to the dominant $2\tilde{N}$ effective Josephson couplings at any stage in the SDRG, and all statistics will be done only on these $2\tilde{N}$ couplings.

3.2.2 Fluctuations within Clusters: Debye-Waller Factors

Consider the link decimation procedure, as discussed in Section 2.3 of Chapter 2. There we approximated the relative phase between the two clustered sites as zero: $\phi_R \equiv \phi_j - \phi_k \approx 0$ in Figure 3.3. We later noted that this is not a good approximation in low dimensions: true long range order is absent due to the fluctuations within clusters. We incorporated these fluctuations into our calculation by evaluating the susceptibility for a cluster directly from a uniformized Hamiltonian (2.74). However, in a numerical application of the SDRG, these fluctuations can be incorporated directly into the RG rules. We will describe this approach below.

Higher order corrections to the link decimation procedure, arising from harmonic vibrations of the phases that make up the cluster, can be obtained by considering the part of the local Hamiltonian (2.20) involving the relative coordinate:

$$\hat{H}_R = (U_j + U_k)\hat{n}_R^2 - J_{jk}\cos(\hat{\phi}_R) \tag{3.5}$$

These terms act approximately like a simple harmonic-oscillator Hamiltonian on the basis of $\hat{n}_R$ eigenstates[6]. Thus, the ground state for the relative coordinate can be

[6]Note that clustering implies that there will not be phase slips between site j and k. Hence, ϕ_R may be treated as noncompact.

approximated by a simple harmonic-oscillator ground state:

$$|\psi_R\rangle \approx \frac{\gamma^{\frac{1}{8}}}{\pi^{\frac{1}{4}}} \int_{-\infty}^{\infty} dn_R \exp\left[-\frac{\gamma^{\frac{1}{2}}}{2} n_R^2\right] |n_R\rangle \tag{3.6}$$

with:

$$\gamma = \frac{2(U_j + U_k)}{J_{jk}} \tag{3.7}$$

We use this approximation in our numerics to compute *Debye-Waller factors* that modify Josephson couplings entering the newly formed cluster. Quantum fluctuations of ϕ_R weaken the phase coherence of the cluster, and consequently, suppress these Josephson couplings. Mathematically, the Debye-Waller factors arise because, in writing down the local, two site Hamiltonian (2.17), we have neglected that $\hat{\phi}_R$ also appears in the other links penetrating the two sites j and k. Consider a Josephson coupling from a third site m to the site j. This corresponds to a term in the full Hamiltonian (2.9):

$$\begin{aligned}
\cos\left(\hat{\phi}_m - \hat{\phi}_j\right) &= \cos\left(\hat{\phi}_m - \hat{\phi}_C - \mu_1 \hat{\phi}_R\right) \\
&= \cos\left(\hat{\phi}_m - \hat{\phi}_C\right) \cos\left(\mu_1 \hat{\phi}_R\right) \\
&\quad + \sin\left(\hat{\phi}_m - \hat{\phi}_C\right) \sin\left(\mu_1 \hat{\phi}_R\right) \\
&\approx \cos\left(\hat{\phi}_m - \hat{\phi}_C\right) \langle \cos\left(\mu_1 \hat{\phi}_R\right) \rangle + \\
&\quad \sin\left(\hat{\phi}_m - \hat{\phi}_C\right) \langle \sin\left(\mu_1 \hat{\phi}_R\right) \rangle
\end{aligned} \tag{3.8}$$

with:

$$\mu_j = \frac{U_j}{U_j + U_k} \tag{3.9}$$

The angle brackets in the final line of equation (3.8) refer to averages taken in the relative coordinate ground state (3.6). The expectation value of the sine vanishes, and the expectation value of the cosine yields the Debye-Waller factor:

$$c_{\mathrm{DW},j} \approx \frac{\sin^2(\pi\mu_j)}{\pi^2} \sum_{q=-\infty}^{\infty} \frac{(q^2+\mu_j^2)}{(q^2-\mu_j^2)^2} \exp\left(-\frac{\gamma^{\frac{1}{2}}}{4} q^2\right) \tag{3.10}$$

In the numerics, we truncate the calculation of this sum at a specified order, $|q_{\max}| = 20$, and multiply the Josephson coupling J_{mj} by the result to find the new Josephson coupling $\tilde{J}_{mC}$ penetrating the cluster. Note that the Debye-Waller factor for links penetrating the site k is, in general, not equal to $c_{\mathrm{DW},j}$, but its calculation is completely analogous.

3.2.3 Initial Disorder Distributions

At the beginning of the numerical procedure, we specify the bare "microscopic" couplings of the rotor model by sampling them from various distributions. We specifically choose the initial distributions $P_i(U)$ and $P_i(J)$ to have one of the following forms:

1. Gaussian distributions truncated at three standard deviations:

$$P_i(x) \propto \exp\left[-\frac{(x-x_0)^2}{2\sigma_x^2}\right] \tag{3.11}$$

for $x \in (x_0 - 3\sigma_x, x_0 + 3\sigma_x)$.

2. Power law distributions with upper and lower cutoffs:

$$P_i(x) = \frac{\eta + 1}{x_{\max}^{\eta+1} - x_{\min}^{\eta+1}} x^{\eta} \tag{3.12}$$

for $x \in (x_{\min}, x_{\max})$.

3. Flat distributions with upper and lower cutoffs:

$$P_i(x) = \frac{1}{x_{\max} - x_{\min}} \tag{3.13}$$

for $x \in (x_{\min}, x_{\max})$. Of course, this is just a power law with exponent $\eta = 0$.

4. "Bimodal" distributions consisting of two flat peaks centered at x_ℓ and x_h:

$$P_i(x) = \frac{1}{2\delta x} \tag{3.14}$$

for $x \in (x_\ell - \frac{\delta x}{2}, x_\ell + \frac{\delta x}{2})$ and $x \in (x_h - \frac{\delta x}{2}, x_h + \frac{\delta x}{2})$.

All of these distributions have positive lower and upper cutoffs ($x_{\min}$ and $x_{\max}$ respectively) and have zero weight for x outside of these bounds. This restriction avoids the complications of frustration in the phase degrees of freedom and the pathologies of the particle sinks that result from on-site charging spectrums that are unbounded from below.

3.3 Numerical Results

We now proceed to present numerical data collected from our implementation of the SDRG. First, we explore the SDRG flows of the distributions of charging energies and Josephson couplings. This investigation points to the existence of an unstable fixed point of the RG flow. We find that the presence of this fixed point is robust to many different changes in the choices of the initial distributions. Next, we examine the distributions generated by the RG near this fixed point and find that universal physics arises in its vicinity. Subsequently, we proceed away from the fixed point to study properties of the phases of the disordered rotor model. We find phases that we tentatively identify as Mott insulating, glassy, and superfluid, and furthermore, we find that the unstable fixed point governs the putative glass-superfluid transition. We defer detailed interpretation of the data and analysis of the transition to Section 3.4.

3.3.1 Flow Diagrams and the Finite Disorder Fixed Point

In Figures 3.6-3.8, we plot flows of quantities that characterize the Josephson coupling and charging energy distributions. We emphasize again that, in the context of our study of distribution flows, the "Josephson coupling distribution" actually only includes the greatest $2\tilde{N}$ Josephson couplings, where $\tilde{N}$ is the number of sites remaining in the effective lattice. After M decimation steps of the RG, we stop the procedure and look at the remaining charging energies and these dominant Josephson couplings. We then use these values to update estimates for the mean and standard deviation

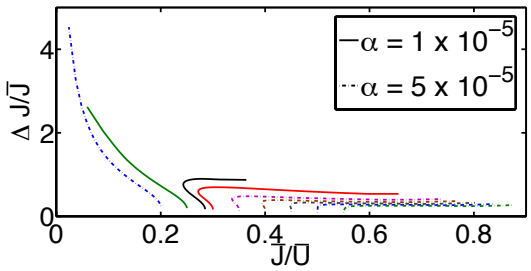

Figure 3.6: The projection, in the $\Delta J/\bar{J}$ vs. $\bar{J}/\bar{U}$ plane, of the flows of the coupling distributions at different stages of the RG. The initial distributions $P_i(U)$ and $P_i(J)$ are both truncated Gaussians, and J_0 (the center of the initial J distribution) is used as the tuning parameter. Each flow corresponds to a different choice of the tuning parameter. The flows start at the bottom of the figure and go up and to the left or up and to the right. A smaller value of the thresholding parameter is used near criticality as indicated by the legend. All runs were done on $L = 100$ lattices.

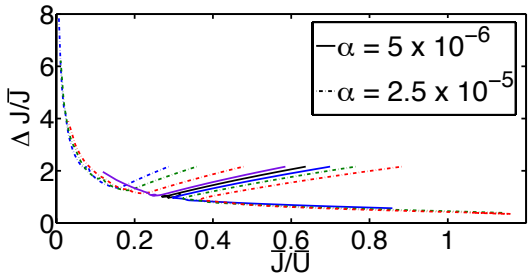

Figure 3.7: Same as Figure 3.6, except $P_i(U)$ is Gaussian and $P_i(J) \propto J^{-1.6}$, with cutoffs chosen to make the latter distribution very broad. The parameter U_0 is used to tune through the transition. The flows begin near the center of the figure. To the left of the figure, flows initially propagate towards the bottom left but eventually turn around and propagate towards the top left. To the right of the figure, flows initially propagate towards the bottom left but eventually turn around and propagate towards the bottom right. All runs were done on $L = 300$ lattices.

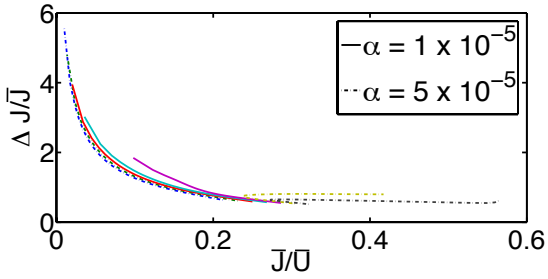

Figure 3.8: Same as Figure 3.6, except $P_i(U) \propto U^5$ and $P_i(J) \propto J^{-3}$. The tuning parameter is $J_{\min}$, the lower cutoff of $P_i(J)$. All runs were done on $L = 150$ lattices.

corresponding to that step in the RG. For each realization of the randomness (i.e., each sample), we do this for many different choices of M, and we repeat the process for 10^3 realizations of the randomness. Ultimately, this procedure gives a "flow" that characterizes the disorder-averaged evolution of the distributions at different stages of the renormalization.

The x-axes of Figures 3.6-3.8 give the ratio of the means of the two distributions. Meanwhile, the y-axes give the ratio of the standard deviation of the Josephson coupling distribution to the mean of the distribution. The plots actually show 2D projections of flows that occur in the space of all possible distributions. At the very least, these plots imply the existence of a third axis, namely $\frac{\Delta U}{\bar{U}}$, which may carry important information. Nevertheless, these highly simplified 2D pictures are surprisingly effective in describing the fate of the model with different choices of parameters. In interpreting Figures 3.6-3.8, the reader will likely find it helpful to glance back to Figure 3.1, which shows a schematization of the flows.

Figure 3.6 specifically corresponds to flows for initial distributions $P_i(U)$ and $P_i(J)$ that are Gaussian. The center of the Josephson coupling distribution is used as the tuning parameter. To the left of the plot are two flows that propagate to the top left of the diagram, towards small $\frac{\bar{J}}{\bar{U}}$ and large $\frac{\Delta J}{\bar{J}}$. Because these flows propagate towards high $\bar{U}$, it is tempting to identify them as flowing towards an insulating regime. Meanwhile, to the right of the plot, there are seven flows that propagate towards high $\bar{J}$, and it is tempting to identify these as propagating towards a superfluid regime. At the interface between these two behaviors, the flows "slow down" and travel a shorter

distance in the plane. This behavior is suggestive of a separatrix flow that terminates at an unstable fixed point, as shown in Figure 3.1.

Our next goal will be to show that the behavior indicating the presence of this unstable fixed point is robust to changes in the choice of the initial distributions. In Figure 3.7, $P_i(U)$ is a Gaussian, and $P_i(J) \propto J^{-1.6}$. The center of the charging energy distribution, U_0, is used as the tuning parameter. The numerical choices place the flows initially above and to the right of the location of the unstable fixed point in the previous figure. From the point of view of the SDRG procedure, this choice of initial distributions is advantageous: the flows begin in a regime of high disorder in J, where the procedure is more accurate. Later in the chapter, after presenting evidence of the universal physics that emerges in the disordered rotor model, we will focus on this choice of distributions exclusively. Therefore, we have collected additional details about these distributions in Appendix 3.C. Note that the leftmost flows in Figure 3.7 initially propagate towards the lower left hand corner of the figure; then, they turn upward, continuing onward to lower $\frac{\bar{J}}{U}$ but now also towards high $\frac{\Delta J}{\bar{J}}$. Hence, they share the same qualitative fate as the leftmost flows in Figure 3.6. To the right of Figure 3.7, the flows initially also propagate towards the lower left; however, these flows ultimately turn around and propagate towards high $\frac{\bar{J}}{U}$. The separatrix that divides these two classes of flows appears to terminate in the same critical region that was seen in Figure 3.6.

In Figure 3.8, we make yet another choice of initial distributions. Now, $P_i(J) \propto J^{-3}$ and $P_i(U) \propto U^5$. The resulting flow diagram again suggests the presence of an

unstable fixed point in the same critical region. It would be misleading, however, to suggest that every flow diagram generated by the RG will have the nice properties of Figures 3.6-3.8. We provide a counterexample in Figure 3.9, in which $P_i(U)$ is bimodal and $P_i(J) \propto J^{-1.6}$. Panel (a) shows the extremely complicated behavior of some of the flows. These features are reflections of the structural details of the bimodal distribution. We will see shortly that, at least in the vicinity of criticality, the RG works to wash away these details and construct universal distributions. After these universal distributions are somewhat well approximated, the flows should be more well behaved, but in Figure 3.9, we see a nonuniversal era of the flows, where the complexities of the initial distributions can manifest in complicated flows. To bring out this point more clearly, we have removed data for the early stages of the RG in panel (b). Now, the behavior at late RG time falls more nicely in line with what is seen in Figure 3.7.

3.3.2 Universal Distributions

Near the unstable finite-disorder fixed point of the RG flow, we expect universal physics to emerge. Certain aspects of the critical behavior should be independent of microscopic details, including the structure of the initial distributions. The universality of the fixed point should become evident in the forms of the *renormalized* distributions generated through the RG: whatever the initial distributions may be, they should evolve towards universal forms, provided that they put the system near criticality.

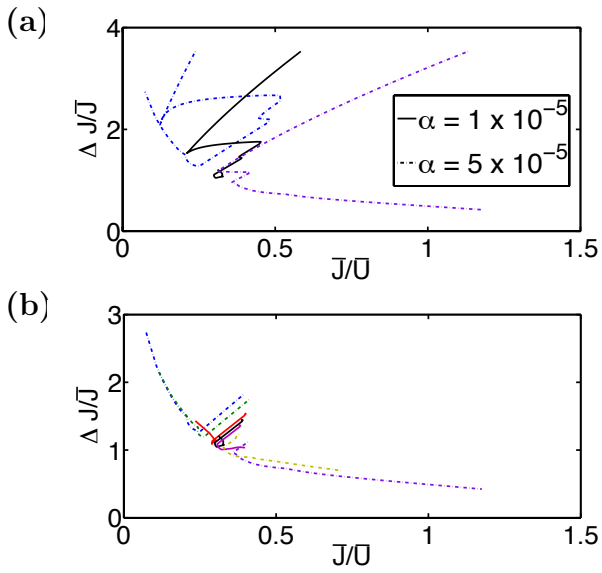

Figure 3.9: In these numerical flow diagrams, $P_i(U)$ is bimodal and $P_i(J) \propto J^{-1.6}$, with cutoffs chosen to make the latter distribution very broad. In panel (a), we show a few sample flows that start near the top of the figure and initially propagate towards the lower left hand corner. The complex features of these flows reflect the structural details of the bimodal U distribution. In panel (b), we exhibit the flows at late RG times, when the procedure has had an opportunity to renormalize away the microscopic details of the initial distributions. Then, the flows are, at least qualitatively, more similar to those seen in Figure 3.7. All runs were done on $L = 200$ lattices.

We first focus on determining the universal form of the fixed point Josephson coupling distribution. Figure 3.10 shows data for the four different choices of the initial distributions that we explore in this chapter. In panels (a), (c), and (d), $P_i(U)$ and $P_i(J)$ have the same qualitative form as one another, and in panel (b), $P_i(J) \propto J^{-1.6}$ and $P_i(U) \propto U^{1.6}$. We tune the parameters characterizing the distributions such that the flows propagate near the unstable fixed point, run the numerics on 100×100 lattices, and plot the initial distributions alongside the renormalized distributions when 100 sites remain in the effective lattice. For the renormalized distributions, we again only include the dominant $2\tilde{N} = 200$ Josephson couplings for each sample. The renormalized distributions suggest that the RG indeed washes away the details of the initial choices, leaving a power law in each case. The universality of this power law is more striking in Figure 3.11, where we plot the renormalized distributions for the four cases together. In this plot, we scale J for each of the four cases by the mean RG scale Ω when only 100 sites remain. This scaling causes the distributions to nearly collapse onto one another, revealing the universal form:

$$P_{\text{univ}}\left(\frac{J}{\Omega}\right) \propto \left(\frac{J}{\Omega}\right)^{-\varphi} \tag{3.15}$$

We will momentarily defer providing a numerical estimate of φ, in anticipation of presenting higher quality data, taken from runs on larger lattices, below.

We have not plotted the renormalized charging energy distributions for the four cases shown in Figure 3.10. Were we to do so, we would see that these distributions, while showing hints of universality, are not as strikingly universal as the corresponding

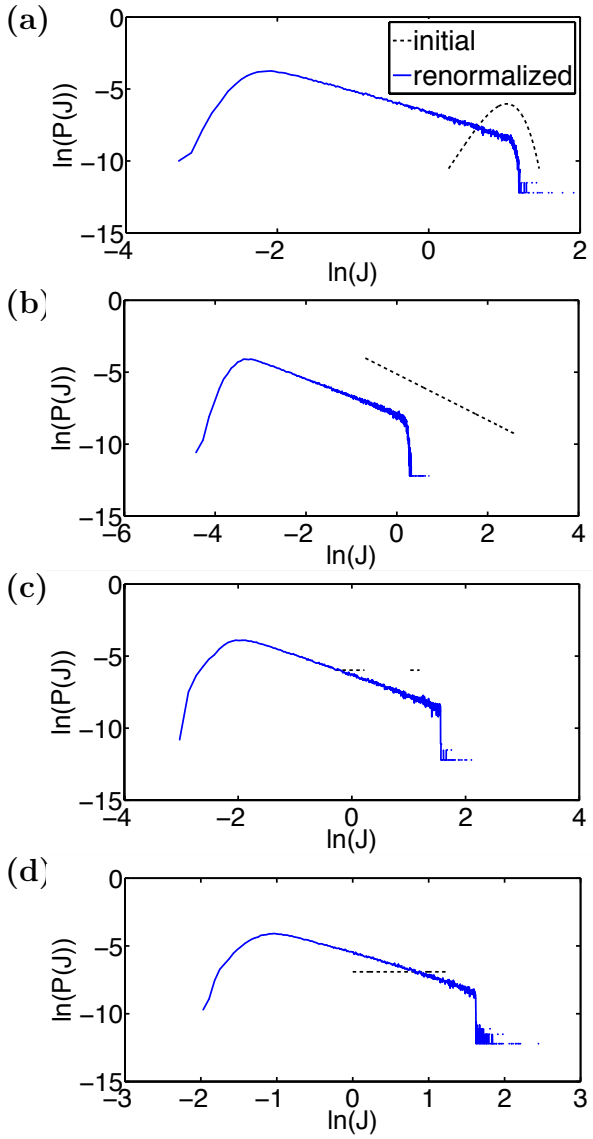

Figure 3.10: Log-log plots of initial and renormalized Josephson coupling distributions for near-critical flows. All runs were done on $L = 100$ lattices with $\alpha = 5 \times 10^{-6}$. Each plot shows the initial distribution and the distribution when the effective lattice has 1% of the original number of sites. The initial distributions have four different forms, but the distributions after renormalization show a universal power law. Note that the plots of initial distributions in these plots were not constructed from actual data (i.e., actual numerical sampling of the distributions), but were instead constructed by hand.

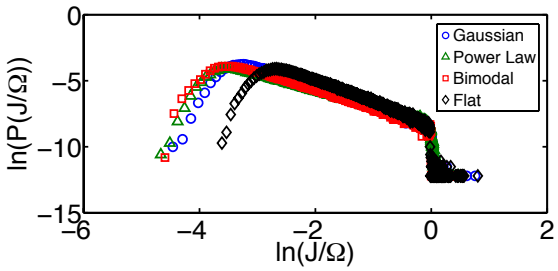

Figure 3.11: The distributions from Figure 3.10, with the Josephson coupling strength scaled by the mean RG scale Ω at the corresponding stage of the RG. The near collapse of the distributions reveals the universal power law form of the Josephson coupling distribution near the unstable fixed point. A refined version of this plot, showing data for larger lattices (but also different choices of initial distributions), can be found in panel (b) of Figure 3.13.

renormalized Josephson coupling distributions. The reason for this is the following: in three out of the four cases, the initial distributions have $J_{\max} < U_{\max}$. Several of the initial distributions we study in this chapter satisfy this property, because in dimensions greater than one, interesting choices of distributions typically have most bare charging energies greater than most bare Josephson couplings. Otherwise, they almost certainly yield superfluid behavior. Consequently, for three out of the four cases in Figure 3.10, the RG begins with only site decimations. These site decimations dramatically modify the Josephson coupling distribution, but the charging energy distribution is, to a large extent, only truncated from above by the renormalization scale. Later on in the procedure, after many sites have been decimated away, the RG enters a regime where the charging energy and Josephson coupling scales compete. Only then do link decimations begin to occur, and only then can the charging energy distribution begin to evolve in a nontrivial way. However, by this point, there is far less RG time remaining for the fixed point distribution to emerge.

There are two ways to circumvent this difficulty. One strategy is to note that this problem of insufficient RG time would not arise if we had access to arbitrarily large lattices. We could follow the renormalization as long as necessary to construct the universal distributions. Thus, we can try to explore larger lattices up to the limits set by our computational capabilities. On the other hand, another solution is to work with very wide initial distributions of Josephson couplings. These are distributions which have large $\frac{\Delta_J}{\bar{J}}$. As such, they correspond to flows that begin above the unstable fixed point in our $\frac{\Delta_J}{\bar{J}}$ vs. $\frac{\bar{J}}{\bar{U}}$ flow diagrams. Using such distributions, it is possible to engineer situations where most bare charging energies exceed most bare Josephson couplings, but where, due to the presence of a small fraction of anomalously large Josephson couplings, $J_{\max} > U_{\max}$ at the beginning of the RG. If the parameters are chosen appropriately, the renormalization procedure will begin with a few link decimations until the charging energy and Josephson coupling scales meet. After this point, the RG will feature an interplay of site and link decimations. Thus, both the Josephson coupling and charging energy distributions will evolve nontrivially.

To target the fixed point charging energy distribution, we apply both of the strategies. We proceed to 200×200 lattices and compute renormalized distributions when the effective lattice has 200 sites remaining. Additionally, we work with very wide initial Josephson coupling distributions. In order to achieve large initial $\frac{\Delta_J}{\bar{J}}$, we restrict our attention to power law distributions of Josephson couplings of the form $P_i(J) \propto J^{-1.6}$. We vary the choice of the initial charging energy distribution and tune the parameters near criticality. The results are shown in Figure 3.12. Now,

the RG does generate a universal form for the charging energy distribution, and in Figure 3.13, we scale the renormalized distributions by the corresponding RG scales to expose the universality more clearly. Figure 3.13 suggests that the functional form of the fixed point charging energy distribution may be:

$$P_{\text{univ}}\left(\frac{U}{\Omega}\right) \propto \left(\frac{\Omega}{U}\right)^{\beta} \exp\left[-\left(\frac{f_U}{\Omega}\right) \times \left(\frac{\Omega}{U}\right)\right] \tag{3.16}$$

where f_U is an energy scale below which the charging energies are exponentially rare. We have been unable to extract a good estimate of β. Panel (a) of Figure 3.13 presupposes $\beta \approx 1$, and the linearity of the plots suggests that this may be close to the actual value. Taking $\beta = 1$ and focusing on the case where $P_i(U)$ is Gaussian (because that is the choice of initial distributions for which we have most accurately targeted criticality), we fit:

$$\frac{f_U}{\Omega} \approx 0.66 \pm 0.02 \tag{3.17}$$

We should note that qualitatively similar charging energy distributions to those seen in Figure 3.12 still emerge near criticality if we relax the restrictions of initially power law J distributions and initially high $\frac{\Delta J}{\overline{J}}$. This is true of the distributions studied in Figure 3.10, even in the flat and bimodal cases where $J_{\max} < U_{\min}$ initially and clusters can only form due to the use of the sum rule. As argued above, the additional restrictions we impose on $P_i(J)$ in Figure 3.12 simply allow the RG to construct the fixed point distributions more cleanly.

In the lower panel of Figure 3.13, we additionally present data for the Josephson

coupling distributions at the same stage of the RG. Again focusing on the data for the case where $P_i(U)$ is Gaussian, we can estimate:

$$\varphi \approx 1.15 \pm 0.01 \tag{3.18}$$

Before proceeding, we should note that the form of the fixed point J distribution allows us to construct an argument for the validity of the RG near criticality. We expand upon this argument greatly in Appendix 3.D, but we will sketch the basic premise here. Essentially, we should consider the implications of the fixed point distributions for the reliability of each of the RG steps. The validity of site decimation rests on the reliability of the perturbative treatment of the Josephson couplings entering the site with the dominant charging energy. The form of the critical Josephson coupling distribution immediately guarantees that most Josephson couplings are much weaker than the RG scale. For the Gaussian case in Figure 3.13, the ratio of the median J to the RG scale is $\frac{J_{\rm typ}}{\Omega} \approx 0.11 \pm 0.01$. Hence, the site decimation is usually very safe. In the case of link decimation, a similar argument allows us to ignore, to leading order, other Josephson couplings penetrating the sites joined by the dominant coupling. However, the structure of the critical *charging energy* distribution, shown in Figure 3.13, actually suggests that there can be a large number of charging energies of the same order as the RG scale; in particular, the ratio of the median U to the RG scale is $\frac{U_{\rm typ}}{\Omega} \approx 0.67 \pm 0.01$. Consequently, the question of the reliability of link decimation reduces to the following: in a single junction (or two-site) problem, how reliable is cluster formation when both charging energies are weaker than the

Josephson coupling but potentially of the same order-of-magnitude? We address this question in Appendix 3.D and find that the link decimation step also seems to be reasonably safe.

3.3.3 Physical Properties and Finite Size Scaling

To determine a preliminary classification of the phases of the model, we now measure four physical properties. First, we measure $s_{\max}$, the size of the largest cluster formed by link decimations during the renormalization procedure. This corresponds to the largest domain of superfluid ordering in the system. We also measure s_2, the size of the second largest cluster. We will see that the behavior of this quantity differs dramatically from that of $s_{\max}$ in the superfluid phase, and therefore, both are interesting quantities to measure.

The third quantity that we measure is the charging gap for the system, $\Delta_{\min}$. We can estimate this quantity by simply measuring the charging energy of the final site remaining in the RG.

Finally, we measure the superfluid susceptibility χ (1.21). Contributions to χ are calculated during site decimation, in a process that closely parallels that one that we used analytically in Chapter 2. Perturbation theory gives a single-site susceptibility of $\frac{1}{U_X}$, where X is the site being decimated. Neglecting corrections from harmonic fluctuations, we can find the contribution of a cluster to the susceptibility by multiplying the perturbative result by s^2, where s is the size of the cluster. One factor of s arises from the fact that the cluster represents s terms in the original sum (1.21), and

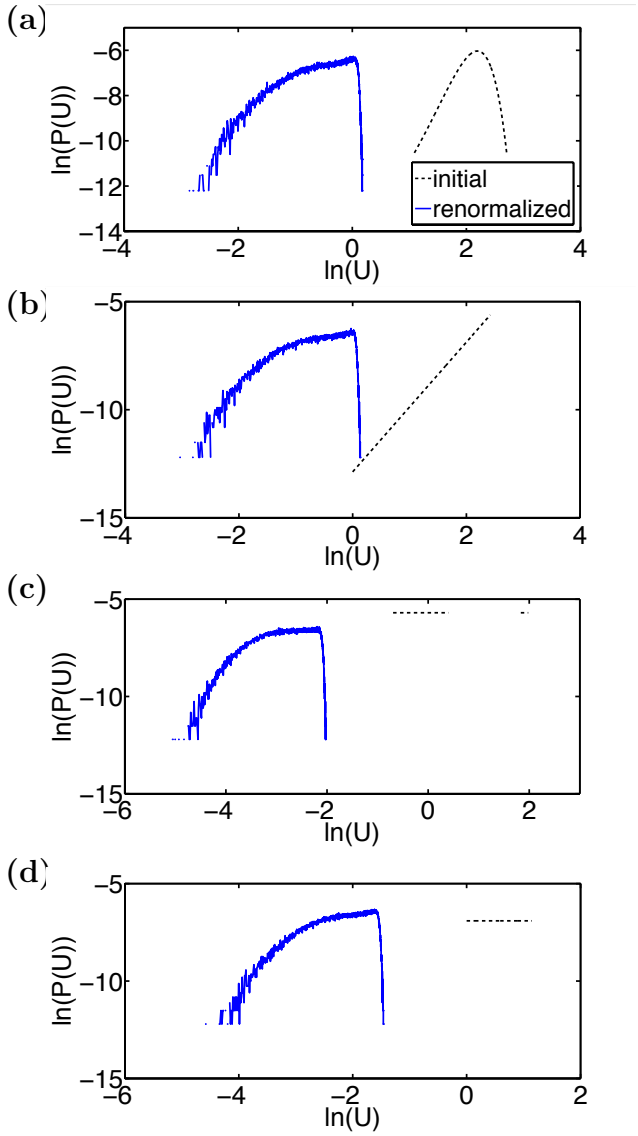

Figure 3.12: Log-log plots of initial and renormalized charging energy distributions for near-critical flows. All runs were done on $L = 200$ lattices with $\alpha = 5 \times 10^{-6}$. Each plot shows the initial distribution and the distribution when the effective lattice has 0.5% of the original number of sites. The initial charging energy distributions have four different forms, but the distributions after renormalization show a universal form. Note that the plots of initial distributions in these plots were not constructed from actual data (i.e., actual numerical sampling of the distributions), but were instead constructed by hand.

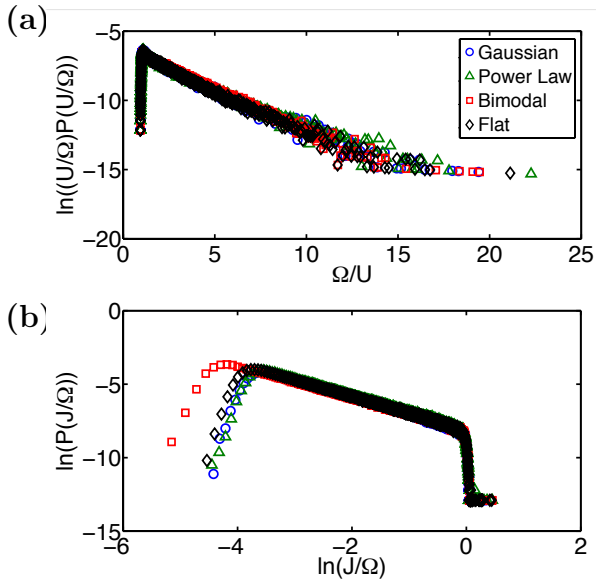

Figure 3.13: In panel (a), the renormalized distributions from Figure 3.12 are plotted together, with the charging energies U scaled by the mean RG scale Ω. In panel (b), we provide a similar plot for the renormalized Josephson coupling distributions produced by the runs in Figure 3.12.

the other follows from the fact that the effective field coupling to the cluster phase is enhanced s times when s phases rotate together. When harmonic fluctuations are taken into account, both of these factors of s should be replaced by a renormalized factor which we denote as b. This b-factor accounts for the fact that quantum fluctuations weaken the phase coherence of the cluster. For a bare site, $b = 1$, and when two sites j and k merge, the renormalized b-factor for the cluster C is:

$$b_C = b_j c_{DW,j} + b_k c_{DW,k} \tag{3.19}$$

where $c_{DW,j}$ and $c_{DW,k}$ are the Debye-Waller factors (3.10). Hence, the total contri-

bution of the cluster to the susceptibility, before the normalization by $\frac{1}{L^2}$, is:

$$\chi_C = \frac{b_C^2}{U_C} \tag{3.20}$$

where b_C and U_C are the b-factor and charging energy of the cluster respectively. Further details of this calculation can be found in Appendix 3.B.

For each lattice size, we obtain an estimate for the four quantities $s_{\max}$, s_2, $\Delta_{\min}$, and χ. Then, we examine how these estimates vary as we raise L. For certain types of distributions, computational limitations force us to work on relatively small lattices. This occurs, for example, when both $P_i(U)$ and $P_i(J)$ are bimodal, and we study this case in Figure 3.14.

In panel (a) of Figure 3.14, there is no cluster formation whatsoever. Hence, $s_{\max} = s_2 = 1$. This results in a gap $\Delta_{\min}$ that remains asymptotically constant as it cannot be lower than the lower bound of the initial charging energy distribution. The susceptibility χ also remains asymptotically constant.

Next, in panel (b), we find a regime in which link decimations do occur and clusters do form. Moreover, $s_{\max}$ and s_2 grow with system size, with what appears to be a power law for the largest lattice sizes that we explore. This growth is, however, subquadratic in L, meaning that $s_{\max}$ grows subextensively with lattice size. Meanwhile, $\Delta_{\min}$ decays with a power slower than L^{-2}, and the susceptibility χ remains constant with L.

In panel (c), all quantities show power law behavior out to $L = 100$, including the susceptibility which appears to grow with a very slow power. The growth of $s_{\max}$ is

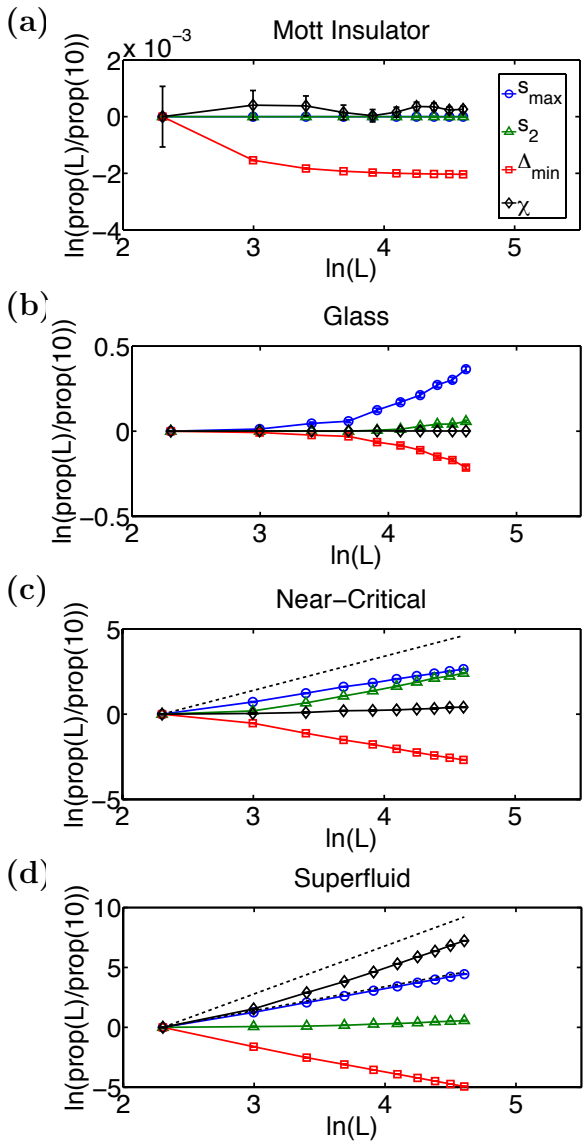

Figure 3.14: Four characteristic behaviors of $s_{\max}$, s_2, $\Delta_{\min}$, and χ as a function of system size. Here, $P_i(U)$ and $P_i(J)$ are both bimodal, and J_h, the center of the higher peak of the Josephson coupling distribution, is used as the tuning parameter. All quantities have been normalized by their value for $L = 10$, and data is shown for $L = 10$ to $L = 100$. In panel (a), the quantities reflect the purely local physics of the Mott insulator. In panel (b), $s_{\max}$ and s_2 grow subextensively with system size, with what appears to be a power law. The gap also decays with a slow power, and the susceptibility remains constant. In panel (c), all quantities show power law behavior out to $L = 100$. The reference line shows L^2 growth. Panel (d) reflects the macroscopic clusters of the superfluid phase. The cluster size $s_{\max}$ is parallel to the L^2 reference line, and the susceptibility χ is parallel to the L^4 reference line for large L.

still slower than L^2, so the largest cluster is still subextensive.

Finally, in panel (d), we find a regime in which $s_{\max}$ grows as L^2, reflecting the formation of macroscopic clusters that scale extensively with the size of the lattice. The gap $\Delta_{\min}$ decays as L^{-2}, and the susceptibility shows an approximately L^4 growth for large L. Perhaps surprisingly, although s_2 grows with system size, it does so more slowly than it does in panel (c).

We now turn to a class of distributions for which we can reach larger lattice sizes. In particular, we return to the data set described in Appendix 3.C, in which $P_i(U)$ is Gaussian and $P_i(J) \propto J^{-1.6}$. Data for this choice of distributions is shown in Figure 3.15.

Panel (a) of this figure qualitatively reproduces the features of panel (a) of Figure 3.14. Panel (b) of the new figure, on the other hand, differs from panel (b) of Figure 3.14 in an important way. For large L, the power law behaviors of $s_{\max}$, s_2 and $\Delta_{\min}$ are lost, and all three quantities vary more slowly. In Figure 3.14, this effect may have been hidden by the use of smaller system sizes.

If the parameters are tuned such that the corresponding flow propagates very close to the unstable fixed point, then we *do* find a regime in which all quantities show nearly power law behavior out to $L = 300$. This regime is depicted in panel (c) of Figure 3.15.

Tuning past this point, we enter a regime in which macroscopic clusters form. Panel (d) of Figure 3.15 shows the behavior of the four quantities in this regime, and we see that most of the essential features of the corresponding panel of Figure 3.14

are reproduced. An important feature of the plot in panel (d) is that we can clearly see that the behavior of s_2 in this regime is closer to that observed in panel (b) than in the intervening panel (c).

The plots in Figures 3.14 and 3.15 are suggestively labeled with their corresponding phase identifications. We will provisionally use these labels for convenience in referring to these regimes, in advance of presenting arguments for these classifications in Section 3.4. In the flow diagrams that we presented earlier, choices of initial distributions that correspond to the Mott insulating and glassy behaviors shown in panels (a) and (b) of Figures 3.14 and 3.15 flow to the stable insulating region where $\frac{\bar{J}}{U} \to 0$. The superfluid behavior in panel (d) of the figures corresponds to flows towards the high $\frac{\bar{J}}{U}$ regime of the flow diagrams. The pure power law behavior of panel (c) emerges for flows that propagate very close to the proposed unstable fixed point. This suggests that this fixed point may control the glass to superfluid transition of the disordered rotor model.

For now, we assume that this is the case and investigate more closely the behavior of physical quantities in the vicinity of this proposed transition. Having provided evidence of the universality that emerges near the critical point, here and in the remainder of this chapter, we will focus exclusively on the choice of distributions detailed in Appendix 3.C. In Figure 3.16, we show that plots of physical quantities vs. tuning parameter, taken for different L, can be collapsed onto universal curves. We will use this scaling collapse in Section 3.4 to determine the critical exponents governing the putative transition between glassy and superfluid phases.

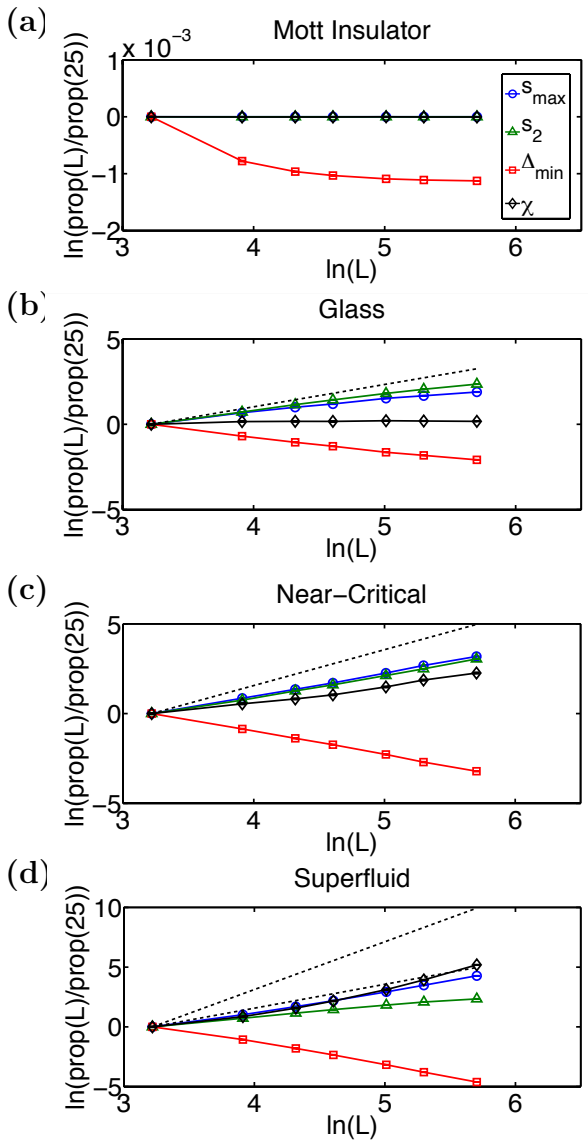

Figure 3.15: Four characteristic behaviors of $s_{\max}$, s_2, $\Delta_{\min}$, and χ as a function of system size. The initial distributions are those described in Appendix 3.C. All quantities have been normalized by their value for $L = 25$, and data is shown for $L = 25$ to $L = 300$. In the four panels, $U_0 = 400$, 9.2, 8.97, and 8.8 respectively. Panel (a) reflects the purely local physics of the Mott insulator. Panel (b) shows a glassy regime characterized by rare-regions clusters that grow subextensively in system size. The reference line shows the power law that $s_{\max}$ obeys near criticality. This nearly critical regime is shown in in panel (c). The reference line here shows L^2 growth. Finally, panel (d) shows the superfluid regime, in which $s_{\max}$ is asymptotically parallel to the L^2 reference line. The susceptibility χ is expected to grow as L^4 for very large L, but it does not quite reach this behavior (indicated with a reference line) for $L \leq 300$.

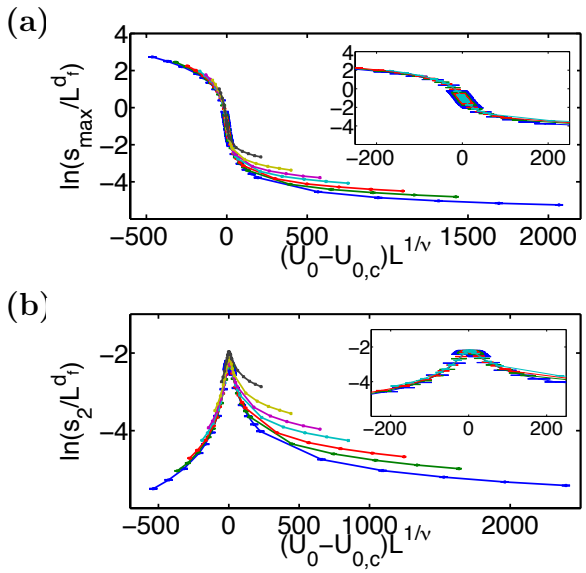

Figure 3.16: Panel (a) shows the scaling collapse of $s_{\max}$ as a function of tuning parameter, and panel (b) shows a similar collapse of s_2. Each line corresponds to a different value of the lattice size. We show data for $L = 25$, 50, 75, 100, 150, 200, and 300. The insets show magnified views of the vicinity of the critical point for the four largest lattice sizes. The error bars, which are difficult to see in the larger plots, are clearly visible in the insets. To see the values of the exponents ν and d_f used in each panel, consult equations (3.48), (3.49), (3.53), and (3.54).

3.3.4 Cluster Densities and b-factors

The physical property that distinguishes the Mott glass from the Bose glass is the compressibility. Later, we will show that, in order to calculate the compressibility, it is insufficient to consider just the minimum charging gap. Within each sample, the RG may form several clusters, each of which implies a local gap for adding and removing bosons. We will need to monitor all of these local gaps to find the compressibility.

More specifically, in this section, we will calculate the density (number per unit area) of clusters of a given size and of clusters with gaps in a given range of energy. We call these densities $\rho(s)$ and $\rho(\Delta)$ respectively. The latter quantity gives a "density of states" for addition of single bosons or holes to the system. For a single sample corresponding to a specific choice of initial distributions, we monitor the size and local charging gap for all the clusters formed during the renormalization, excluding bare sites that are not clustered by the RG. In the case of the local charging gap, we again estimate this quantity as the charging energy of a cluster at the time of decimation. In principle, we could also include perturbative corrections to this local gap, but we omit these and do not anticipate that they would affect the behavior of the density at low Δ. We pool data for 10^3 samples, compute histograms of gaps and cluster sizes, and normalize these histograms by the total simulated surface area: L^2 times the number of samples.

Our study of these densities will bring into focus the crucial difference between two types of clusters formed by the RG: rare-regions clusters and the macroscopic clusters that characterize the superfluid phase. We will, therefore, also take the opportunity

to examine how the b-factors, which quantify the effect of harmonic fluctuations on the susceptibility, vary as a function of s for the two types of clusters.

3.3.4.1 The Charging Gap Density $\rho(\Delta)$

Note that $\rho(s)$ and $\rho(\Delta)$ are not particularly interesting for choices of distributions and parameters that yield the Mott insulating behavior from Figures 3.14 and 3.15. The profile of $\rho(\Delta)$ will be identical to the profile of the initial charging energy distribution, and because we choose this distribution to be bounded from below by some positive $U_{\min}$, it can be shown that this always corresponds to an incompressible phase. Hence, we begin by focusing on the glassy regime.

Figure 3.17 shows the gap density for a choice of parameters in the glassy phase. As the size of the lattice is raised, the density profile remains essentially invariant at large Δ, but smaller gaps, corresponding to larger clusters, begin to appear. However, these smaller gaps simply fill out a decay to 0 as $\Delta \to 0$.

Now, we turn our attention to the putative superfluid phase. The gap density in this phase is shown in Figure 3.18. In panel (a), there is an invariant piece to the distribution, but at very low Δ, a second peak appears as well. Panel (b) of Figure 3.18 shows a magnified view of this low Δ peak. This peak propagates towards $\Delta = 0$ as the system size is raised. Accompanying the propagation is a shrinking: the integrated weight of the low Δ peak shrinks as L^{-2}. The consequences of these two effects need to be taken into account carefully to calculate the compressibility.

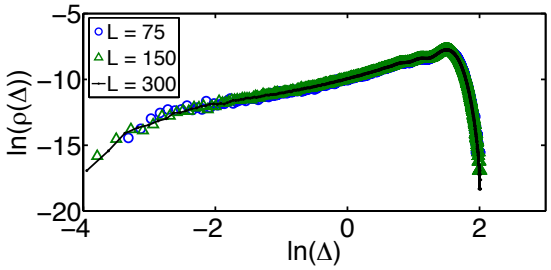

Figure 3.17: The number per unit area of clusters with local gap near Δ for a choice of parameters in the glassy phase. The initial distributions are those described in Appendix 3.C, with $P_i(U)$ Gaussian and $P_i(J) \propto J^{-1.6}$. Data is shown for $L = 75$, 150, and 300 lattices. The density decays faster than a power law at small Δ.

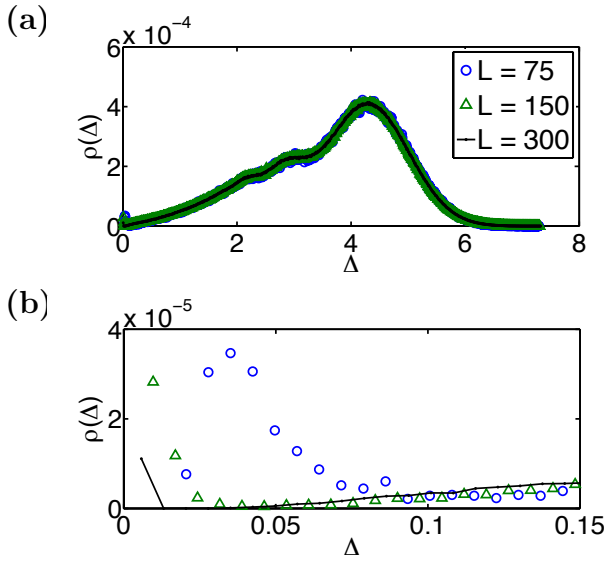

Figure 3.18: In panel (a), the density (number per unit area) of clusters with a given gap Δ for a choice of parameters in the superfluid phase. The initial distributions are those described in Appendix 3.C, with $P_i(U)$ Gaussian and $P_i(J) \propto J^{-1.6}$. Data is shown for $L = 75$, 150, and 300 lattices. The density profiles exhibit two peaks. The broad peak that is visible in panel (a) remains invariant with L. To expose the second peak, we provide a magnified view of the low Δ part of the density in panel (b). This peak simultaneously shrinks and propagates towards $\Delta = 0$ as the system size is raised.

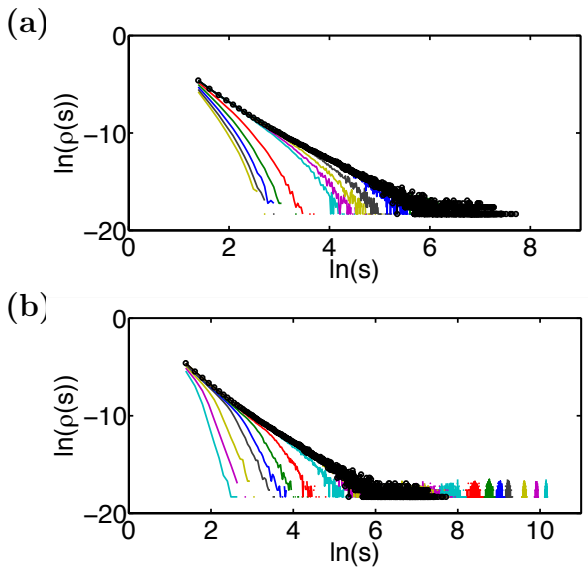

Figure 3.19: Sweeps of $\rho(s)$ as the system is tuned through the superfluid-insulator transition on $L = 300$ lattices. The initial distributions are those described in Appendix 3.C, with $P_i(U)$ Gaussian and $P_i(J) \propto J^{-1.6}$. The tuning parameter is U_0, the center of $P_i(U)$. Panel (a) shows the sweep from deep in the glassy phase ($U_0 = 20$) to the transition ($U_0 = 8.97$). The closest data set to the transition is indicated by the black line with large data point markers. This line is reproduced in panel (b), which shows the sweep from the transition into the superfluid phase (up to $U_0 = 6.5$). In the superfluid phase, the density plot is characterized by a peak at large s and a remnant decay at low s.

3.3.4.2 The Cluster Size Density $\rho(s)$

We now consider how $\rho(s)$, the density of clusters of size s, varies as we sweep through the glassy regime and into the superfluid. Panel (a) of Figure 3.19 shows the approach to the transition from the glassy side. Very close to the transition at $U_0 \approx 8.97$, $\rho(s)$ exhibits a striking power law decay. Proceeding into the proposed glassy phase, the power law decay persists at small s. However, this behavior is cut off by some scale $\tilde{s}$, beyond which $\rho(s)$ decays very rapidly, essentially exponentially.

In panel (b), we proceed in the other direction from the putative transition, into the superfluid regime. A peak, corresponding to the macroscopic clusters, appears at

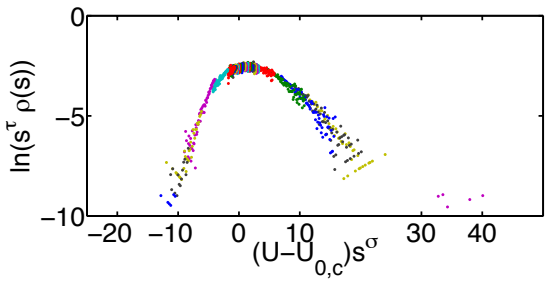

Figure 3.20: Scaling collapse of the $\rho(s)$ curves from Figure 3.19. Small cluster sizes ($s < 30$) need to be discarded, because they are nonuniversal. Large cluster sizes ($s > 100$) need to be discarded, because they are noisy. Then, the remaining $\rho(s)$ curves, taken for different values of the tuning parameter, collapse onto a universal curve.

large s. The macroscopic cluster in each sample is dressed by rare-regions clusters, and these clusters are represented by the remnant decay at small s. While the large s peak is related to the low Δ peak in Figure 3.18, the remnant decay at small s is the analogue of the high Δ feature that stays invariant with system size. The low s decay in panel (b) of Figure 3.19 qualitatively resembles the decay well inside the glassy regime. In summary, $\rho(s)$ exhibits a pure power law decay in the vicinity of the proposed glass-superfluid transition; tuning away from criticality in either direction, and excluding the macroscopic clusters of the superfluid phase, the power law form of $\rho(s)$ only survives up to a scale $\tilde{s}$. For $s > \tilde{s}$, clusters become exponentially rare.

A type of scaling collapse can be performed for the $\rho(s)$ curves from Figure 3.19, and this collapse is shown in Figure 3.20. We will see below that this collapse gives a complementary set of critical exponents. These exponents are connected by scaling relations to those that can be extracted from Figure 3.16.

3.3.4.3 Susceptibility b-factors

The data presented above for the cluster densities $\rho(s)$ and $\rho(\Delta)$ highlights the difference between the rare regions and macroscopic clusters generated by the RG. A study of how the b-factors for the clusters vary as a function of s can bring out another difference between the two types of clusters. Recall that the b-factor was introduced in equation (3.20) to quantity the effect of harmonic fluctuations on the susceptibility of a superfluid cluster. As such, understanding the behavior of the b-factors will also be essential in explaining the behavior of χ in the various phases of our model.

In Figure 3.21, we plot the average value of b for a cluster of size s and plot it against s. Again, we work with $L = 300$ lattices and the choice of distributions described in Appendix 3.C. Panel (a) shows data for the glassy phase and for the nonmacroscopic clusters of the superfluid phase. We see that $\bar{b}$ varies with s as a power law:

$$\bar{b}(s) \sim s^{\zeta} \tag{3.21}$$

and that the power is consistent for seven different choices of the tuning parameter U_0. We will provide an estimate of ζ in Section 3.4. Panel (b) of Figure 3.21 shows data for the macroscopic clusters when $U_0 = 8.8$. Now, $\bar{b}(s) \propto s$. This behavior can be anticipated from a simple picture of macroscopic cluster growth, which we will discuss in Section 3.4.

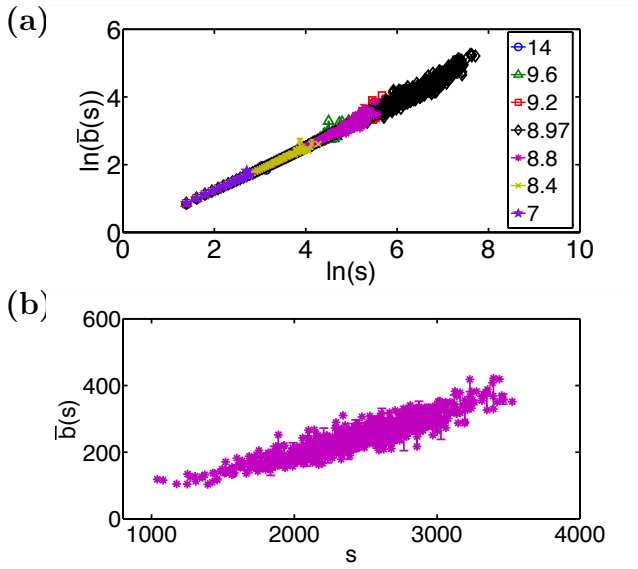

Figure 3.21: The behavior of the mean b-factor for clusters of size s as a function of s. In panel (a), we show data for the glassy regime and for the nonmacroscopic clusters of the superfluid regime. The initial distributions are those described in Appendix 3.C, and data is shown for seven values of the tuning parameter U_0 on $L = 300$ lattices. The log-log plot shows power law behavior of $\bar{b}(s)$ vs. s. In panel (b), we show data for the macroscopic clusters when $U_0 = 8.8$. The plot indicates that $\bar{b}(s) \sim s$ for macroscopic clusters. In both plots, the error bars indicate the error on the mean $\bar{b}(s)$ over all clusters of size s. For some of the largest and smallest values of s in each plot, the absence of an error bar indicates that only one cluster of that size was generated in all of the samples.

3.4 Phases and Quantum Phase Transitions

Having collected representative numerical data in the previous section, we now assess what the data tells us about the phases and phase transitions of our model. Our first task will be to confirm the association of phases with the behaviors of physical properties that we observed in Figures 3.14 and 3.15. To this end, we will have to preemptively introduce one of the main conclusions of the discussion below: that our data points to a superfluid-insulator transition driven by a percolation-type process. The rare-regions clusters of the glassy phase find one another, and their phases cohere, producing a macroscopic cluster of superfluid ordering and driving the transition to long range order and global superfluidity.

Motivated by this intuitive picture of the transition, the logic of the discussion below will be the following: in the proposed glassy and superfluid regimes of Figure 3.15, the cluster size density $\rho(s)$ exhibits the universal features that we would expect from nonpercolating and percolating phases of standard models of percolation. We can use these correspondences to extrapolate the behaviors seen in Figure 3.15 to the thermodynamic limit, in the process showing that these phases will indeed have the properties expected of glassy and superfluid phases. Furthermore, we can analyze the critical point and extract the critical exponents that govern the superfluid-insulator transition. After characterizing this transition, we will finally return to the question of the identity of the glassy phase and determine if a Mott glass is present in two dimensions.

Before proceeding, we should clarify that, although our RG produces clusters with

size distributions similar to models of percolation, our transition is not the standard percolation transition. Indeed, the exponents that we recover are significantly different from the percolation exponents on a 2D square lattice. However, the analogy to percolation allows us to easily identify the relations between the various exponents and the scaling functions for the observables.

3.4.1 Phases of the Model

3.4.1.1 Mott Insulator

We briefly digress to describe the phase of our model in which the percolation picture is irrelevant, simply because there are no regions of superfluid ordering. In a Mott insulating phase, the system wants to pin the number fluctuation to zero on each site to avoid the energetic costs of charging. Hence, $s_{\max}$ and s_2 simply stay pinned at one as L is raised. Meanwhile, $\Delta_{\min}$ should asymptote to a constant, reflecting the fact that the Mott insulator is gapped. A phase without any cluster formation can be described by completely local physics. This means that the susceptibility can be approximated by averaging over disorder in a single site problem. In other words, χ should also stay constant as the system size is increased. Thus, in a Mott insulating phase, we expect exactly the behavior seen in panel (a) of Figures 3.14 and 3.15.

3.4.1.2 Glass

In nonpercolating phases of standard models of percolation, the cluster size density is expected to go as:

$$\rho(s) = cs^{-\tau} f\left(\frac{s}{\tilde{s}}\right) \tag{3.22}$$

where c is a constant. The function $f(x)$ is expected to be approximately constant for $x < 1$ and to rapidly decay for $x > 1$ [130]. Equation (3.22) is consistent with what we have observed in our proposed glassy phase in panel (a) of Figure 3.19 and is also consistent with the expectation that, in a Griffiths phase, the frequency of occurrence of large rare regions decays exponentially in their size [140]. In our numerics, $f(x)$ seems to follow a pure exponential decay $f(x) \sim e^{-x}$, but the conclusions we present below would be qualitatively unchanged if, for example, $f(x) \sim e^{-x^{\lambda}}$.

With the form of equation (3.22) in hand, we can now formulate a simple argument for the asymptotic behavior of $s_{\max}$. In particular, the order-of-magnitude of $s_{\max}$ is set by the condition [130]:

$$L^2 \sum_{s=s_{\max}}^{L^2} \rho_L(s) \approx 1 \tag{3.23}$$

If the left hand side of equation (3.23) is less than 1, then it is unlikely that even one cluster of size greater than or equal to $s_{\max}$ will be present in a sample of size L^2. In equation (3.23), $\rho_L(s)$ is the finite size approximant to the function $\rho(s)$ that appears in equation (3.22). The upper limit of the sum in equation (3.23) is taken as L^2 because larger clusters simply cannot appear in the finite size sample. With enough sampling of the random distributions, it should, in principle, be able to accurately

represent the approximant $\rho_L(s)$ out to $s = L^2$. The data indicates that $\rho_L(s)$ will simply reproduce $\rho(s)$ out to this value of s, so in the calculations below, we can replace the approximant $\rho_L(s)$ by $\rho(s)$. This will *not* be possible in the superfluid phase.

For large L, where we also expect $s_{\max} > \tilde{s}$, we can use equation (3.23) to find $s_{\max}$ by computing

$$\begin{aligned} L^2 \sum_{s=s_{\max}}^{L^2} \rho(s) &\approx L^2 \int_{s_{\max}}^{L^2} ds \rho(s) \\ &= cL^2 \int_{s_{\max}}^{L^2} ds \exp\left[-\frac{s}{\tilde{s}} - \tau \ln(s)\right] \\ &\approx cL^2 \int_{s_{\max}}^{L^2} ds \exp\left[-\frac{s}{\tilde{s}}\right] \\ &= c\tilde{s}L^2 \left(e^{-\frac{s_{\max}}{\tilde{s}}} - e^{-\frac{L^2}{\tilde{s}}}\right) \end{aligned} \tag{3.24}$$

Setting this expression equal to 1 and inverting for $s_{\max}$, we find that, asymptotically in L:

$$s_{\max} \sim \ln L \tag{3.25}$$

For large clusters, the link decimation rule for addition of charging energies (2.21) implies[7] that $U \sim s^{-1}$, and therefore:

$$\Delta_{\min} \sim \frac{1}{\ln L} \tag{3.26}$$

An entirely analogous condition to equation (3.23) can be formulated for s_2. We

[7]Note that we justified this result analytically in 1D in equation (2.65).

simply replace the right hand side of equation (3.23) with two, indicating that we want to find the value of s such that there are likely to be two clusters of that size or greater in a typical sample. However, the remainder of the calculation is qualitatively unaffected by this change. Therefore, s_2 should also grow logarithmically in this regime:

$$s_2 \sim \ln L \tag{3.27}$$

Finally, we turn to the susceptibility. This can be computed as follows:

$$\begin{aligned} \chi &= \frac{1}{L^2}\sum_C \frac{b_C^2}{U_C} \\ &\sim \frac{1}{L^2}\sum_C s_C b_C^2 \\ &\sim \frac{1}{L^2}\sum_{s=1}^{s_{\max}} \rho(s) L^2 (\bar{b}(s))^2 s \\ &= \sum_{s=1}^{s_{\max}} \rho(s)(\bar{b}(s))^2 s \end{aligned} \tag{3.28}$$

In this calculation, the sum over C is a sum over clusters, with s_C being the size of the cluster. Then, the sum over s is, as before, a sum over cluster sizes, and $\bar{b}(s)$ is the average value of the b-factor for a cluster of size s. Figure 3.21 shows that, in the glassy regime, $\bar{b}$ varies as a power of s and that this power remains the same all the way up to criticality. While we do not have a complete understanding of this behavior, we can still understand the asymptotic behavior of χ by reasoning that $\bar{b}$ can, at most, grow linearly in s. This follows from an interpretation of the b-factor as the effective number of rotors that order with the field in a cluster of size s. Because

$\rho(s)$ decays exponentially for large s while $\bar{b}(s)$ grows at most as a power, the sum (3.28) converges, and χ should be asymptotically constant:

$$\lim_{L\to\infty} \chi = \chi_\infty \tag{3.29}$$

All of these behaviors are consistent with what has been observed numerically in panel (b) of Figure 3.15. Moreover, because logarithmic behavior can be difficult to discriminate from a slow power law at low L, they are also consistent with panel (b) of Figure 3.14. Thus, under the numerically justified assumption that this regime corresponds to a nonpercolating phase, we can deduce that, as L increases, arbitrarily large rare regions of superfluid ordering will appear, driving the gap to zero. However, the typical size of these regions grows extremely slowly (i.e., logarithmically) with system size. The fraction of sites occupied by the largest cluster in a typical sample vanishes as $L \to \infty$, so there is no long range order. The behavior of this phase for large L corresponds to what we should expect for a glassy phase.

3.4.1.3 Critical Region

At the critical point of a percolation transition, the cluster size scale in equation (3.22) is expected to diverge as:

$$\tilde{s} \sim |g - g_c|^{-\frac{1}{\sigma}} \tag{3.30}$$

where g is the tuning parameter for the transition. This divergence is related to the divergence of a correlation length that indicates the typical linear extent of the largest clusters:

$$\xi \sim |g - g_c|^{-\nu} \tag{3.31}$$

By equation (3.30), $\rho(s)$ is a power law at criticality:

$$\rho(s) = cs^{-\tau} \tag{3.32}$$

This means that the critical point is characterized by a scale invariant, fractal structure of clusters [130]. In turn, this implies that ξ and $\tilde{s}$ are related by a fractal dimension:

$$\tilde{s} = \xi^{d_f} \tag{3.33}$$

Equations (3.30), (3.31), and (3.33) together imply:

$$\sigma = \frac{1}{\nu d_f} \tag{3.34}$$

We will use this scaling law in our analysis of the transition below [130].

For the present purposes, note that equation (3.32) is once again consistent with what we have observed numerically in Figure 3.19. Now, the calculation for $s_{\max}$ becomes:

$$L^2 \sum_{s=s_{\max}}^{L^2} \rho(s) \approx cL^2 \int_{s_{\max}}^{L^2} ds s^{-\tau} = 1 \tag{3.35}$$

which, when inverted, yields:

$$\begin{aligned} \ln s_{\max} &= \frac{2}{\tau-1}\ln L - \frac{1}{\tau-1}\ln\left(\frac{\tau-1}{c}\right) \\ &\quad + \ln\left(1+\frac{c}{(\tau-1)L^{2(\tau-2)}}\right) \end{aligned} \tag{3.36}$$

Asymptotically, as long as $\tau > 2$, this simply corresponds to a power law growth of $s_{\max}$:

$$s_{\max} \sim L^{\frac{2}{\tau-1}} \tag{3.37}$$

On the other hand, because d_f is the exponent that connects length and cluster size scales at the transition, equation (3.37) yields another scaling relation:

$$d_f = \frac{2}{\tau-1} \tag{3.38}$$

Equations (3.34) and (3.38) are the usual scaling laws connecting exponents at a percolation transiiton [130].

In accordance with equation 3.37, the gap should close as:

$$\Delta_{\min} \sim L^{-\frac{2}{\tau-1}} = L^{-d_f} \tag{3.39}$$

Furthermore, as in the glassy regime, the qualitative behavior of the second largest cluster size s_2 should be identical to that of $s_{\max}$:

$$s_2 \sim L^{\frac{2}{\tau-1}} \tag{3.40}$$

Turning to the susceptibility, we find that it no longer converges to a finite value. At criticality, power law behavior of $\bar{b}(s)$ follows naturally from scale invariance. When $\rho(s) \sim s^{-\tau}$ and $\bar{b}(s) \sim s^{\zeta}$:

$$\begin{aligned} \chi &= \sum_{s=1}^{s_{\max}} \rho(s)(\bar{b}(s))^2 s \\ &\sim \int_1^{s_{\max}} ds s^{1+2\zeta-\tau} \\ &\sim s_{\max}^{2+2\zeta-\tau} \\ &\sim L^{\frac{4+4\zeta-2\tau}{\tau-1}} \\ &= L^{d_f(1+2\zeta)-2} \end{aligned} \tag{3.41}$$

From Figure 3.21, we can estimate the exponent ζ:

$$\zeta \approx 0.68 \pm 0.01 \tag{3.42}$$

With another choice of initial distributions ($P_i(U)$ bimodal and $P_i(J) \propto J^{-1.6}$), we have found a similar value for ζ. If $d_f(1+2\zeta) > 2$, then χ asymptotically diverges as a power law, as seen in Figure 3.15. We will provide an estimate of d_f shortly in equation (3.49). For now, we note that the power observed for χ vs. L in Figure 3.15 is consistent with this estimate of d_f and the estimate for ζ that is given above. In Figure 3.14, the relatively small system sizes likely put us out of the scaling regime for χ, and this is probably responsible for the extremely slow growth of χ with L.

3.4.1.4 Superfluid

The percolating phase is characterized by the presence of a macroscopic cluster that grows with the size of the lattice, so trivially:

$$s_{\max} \propto L^2 \tag{3.43}$$

and therefore:

$$\Delta_{\min} \propto L^{-2} \tag{3.44}$$

It is possible to construct an argument along the lines of condition (3.23) for the behavior of $s_{\max}$ in equation (3.43), but in this case, it is important *not* to substitute the infinite lattice density $\rho(s)$ for the finite size approximant $\rho_L(s)$. The subtlety lies in the high s peak observed in the density plots in panel (b) of Figure 3.19. Consistent with their low Δ counterparts in the plots of the gap density, the location of these peaks propagates as L^2 towards high s as L is raised. Simultaneously, the integrated weight of the peaks shrinks as L^{-2}, reflecting the fact that there is only one macroscopic cluster in each sample. Balancing the shrinking and propagation, it is possible to see that, in order to achieve the condition (3.23), $s_{\max}$ must scale as shown above in equation (3.43).

The reasoning above has important implications for the behavior of s_2. Because the weight of the high s peak of $\rho_L(s)$ shrinks as L^{-2}, the second largest cluster must be drawn from the remnant low s decay. The physical picture behind this low s decay is the following: suppose we remove the sites belonging to the macroscopic cluster

from the original lattice. In doing so, we remove some of the lowest charging energies and highest Josephson couplings from the bare lattice. Consequently, the remnant lattice is globally insulating. Nevertheless, rare regions of superfluid ordering can form exactly as in the glassy phase. It follows that s_2 will grow with L as in the glassy phase:

$$s_2 \sim \ln L \tag{3.45}$$

Note that s_2 actually grows fastest at criticality (i.e., as a power law) and slower on either side of the transition (i.e., logarithmically). This is responsible for the peak in s_2 at criticality.

To calculate the susceptibility, we first take into account the contribution of the macroscopic cluster. The behavior of $\bar{b}(s)$ for a macroscopic cluster can be inferred from a simplified picture of a large cluster merging with single neighbors. As the macroscopic cluster grows, its charging energy becomes smaller, driving the Debye-Waller factor for the cluster to one. At the same time, the cluster's Josephson couplings to other sites grow through link addition processes. This means that the Debye-Waller factor for another site that is merging with the macroscopic cluster also approaches unity. Therefore, the b-factor addition rule (3.19) approximately becomes $\tilde{b}_C = b_C + 1$, and $b_C \sim s_C$ follows. Then:

$$\begin{aligned} \chi_{mac} &= \frac{1}{L^2} \times \frac{(\bar{b}(s_{\max}))^2}{\Delta_{\min}} \\ &\sim \frac{1}{L^2} \times \frac{s_{\max}^2}{\Delta_{\min}} \\ &\sim L^4 \end{aligned} \tag{3.46}$$

	Mott Insulator	Glass	Critical	Superfluid
$s_{\max}$	const.	$\ln(L)$	L^{d_f}	L^2
s_2	const.	$\ln(L)$	L^{d_f}	$\ln(L)$
$\Delta_{\min}$	const.	$\frac{1}{\ln(L)}$	L^{-d_f}	L^{-2}
χ	const.	const.	$L^{d_f(1+2\zeta)-2}$	L^4

Table 3.1: Behavior of physical properties in the three phases and at the critical point for large lattice sizes L.

The rare-regions clusters dressing the macroscopic cluster add a subleading contribution to the susceptibility, which we call χ_{rr}. The reasoning of equation (3.29) indicates that this contribution should be asymptotically constant. Thus, the quartic behavior of equation (3.46) is the correct asymptotic behavior. Finite size corrections from χ_{rr} will modify this behavior, however, and this is probably why we do not quite see χ reach the L^4 behavior in panel (d) of Figure 3.15.

We have now provided arguments for the behaviors observed in each panel of that figure, and we summarize this information in Table 3.1.

3.4.2 Quantum Phase Transitons

3.4.2.1 Glass-Superfluid Transition

Earlier, we remarked in passing that systems exhibiting the behaviors that we have now identified as Mott insulating and glassy eventually propagate towards the putative insulating region to the top left of the numerical flow diagrams. Correspondingly, the systems exhibiting superfluid behavior propagate towards the putative superfluid region to the bottom right. We can now verify our tentative identifications of these stable regions of the diagram. We can also determine that the unstable fixed point

controls the transition between the glass and the superfluid, the superfluid-insulator transition of our disordered rotor model. This allows us to draw the schematic picture of the flow diagram that we presented in Figure 3.1.

We will now focus on the critical region and extract critical exponents governing the glass-superfluid transition. Estimating these exponents requires formulating scaling hypotheses for the behavior of physical quantities in the critical region. In the case of $s_{\max}$ [130]:

$$s_{\max} = L^{d_f} h_1\left(\frac{L}{\xi}\right) = L^{d_f}\tilde{h}_1(L^{\frac{1}{\nu}}(g - g_c)) \tag{3.47}$$

Exactly at criticality, $s_{\max} \sim L^{d_f}$ asymptotically, so plotting $\frac{s_{\max}}{L^{d_f}}$ vs. $(g-g_c)$ generates a crossing of the curves for different lattice sizes. Slightly away from criticality, the power law behavior should persist if $L < \xi$. For $L > \xi$, the system recognizes that it is not critical and we should see a crossover to logarithmic growth on the glassy side and L^2 growth on the superfluid side. Hence, $\frac{L}{\xi}$ is the important ratio near criticality, and this motivates scaling form (3.47). The scaling form, in turn, implies that we can produce scaling collapse by plotting $\frac{s_{\max}}{L^{d_f}}$ vs. $(g - g_c)L^{\frac{1}{\nu}}$. This is what we have done in panel (a) of Figure 3.16. This scaling collapse leads to the estimates:

$$\nu \approx 1.09 \pm 0.04 \tag{3.48}$$

$$d_f \approx 1.3 \pm 0.2 \tag{3.49}$$

These estimates and error bars are obtained through the following procedure:

first, to find an estimate of the critical value of the tuning parameter g, we examine the behavior of the sample average of s_2 vs. g. Since s_2 varies as a power law in L exactly at criticality and grows more slowly within the phases, s_2 should exhibit a maximum at g_c. We can find the error Δg_c on our estimate of g_c by taking the window of values of g around g_c for which the sample average of s_2 is within one error bar of the maximum. To proceed to obtain estimates for ν and d_f, we now partition our data into bins of size N_b samples. For example, the first bin may consist of the first $N_b = 50$ samples for each value of the tuning parameter and each lattice size L. Our immediate goal is to find the best values of ν and d_f for this subset of our data. We first focus on the data for $g = g_c$ and do linear regression to find the best exponent that describes the power law growth of the average value of $s_{\max}$ with L. This gives an estimate of d_f for the bin. Next, for two lattice sizes (typically, $L = 150$ and $L = 300$), we compute an average of $s_{\max}$ over the samples in the bin for several values of the tuning parameter around g_c. Then, using the previous estimate of d_f for the bin and the scaling hypothesis (3.47), we test various values of ν until we achieve the best collapse. Thus, we also obtain an estimate of ν for the bin. From the distribution of estimates for the different bins, we can find mean values and error bars for d_f and ν. However, these error bars do not take into account the error on our estimate of the critical point. To propagate this error, we need to repeat the critical point estimation procedure using $g_c + \Delta g_c$ and $g_c - \Delta g_c$ as our estimates of the critical point.

We have repeated the scaling collapse of $s_{\max}$ for a different choice of initial distri-

butions: bimodal $P_i(U)$ and $P_i(J) \propto J^{-1.6}$. Ultimately, we have been able to recover estimates for ν and d_f which are consistent with (3.48) and (3.49):

$$\nu \approx 1.1 \pm 0.1 \tag{3.50}$$

$$d_f \approx 1.2 \pm 0.2 \tag{3.51}$$

A completely analogous scaling hypothesis can be made for s_2:

$$s_2 = L^{d_f} h_2\left(\frac{L}{\xi}\right) = L^{d_f} \tilde{h}_2(L^{\frac{1}{\nu}}(g - g_c)) \tag{3.52}$$

Then, the exponents needed to produce the collapse in panel (b) of Figure 3.16 are:

$$\nu \approx 1.06 \pm 0.09 \tag{3.53}$$

$$d_f \approx 1.31 \pm 0.07 \tag{3.54}$$

Since all the estimates (3.48)-(3.54) are consistent, we will proceed using our tightest estimates on these exponents: (3.48) and (3.54).

We now note that a scaling hypothesis can also be formulated for the density $\rho(s)$. From equation (3.22), we see that, for fixed lattice size L, $s^\tau \rho(s)$ should depend only on the combination $\frac{s}{\tilde{s}} \sim s(g - g_c)^{\frac{1}{\sigma}}$. Hence, by plotting $s^\tau \rho(s)$ vs. $s(g - g_c)^{\frac{1}{\sigma}}$ and tuning until the curves for different choices of g collapse, we ought to be able to extract estimates for τ and σ. On the other hand, τ and σ are related to ν and d_f through the scaling relations (3.38) and (3.34). Therefore, from the estimates (3.48)

and (3.54), we can infer:

$$\tau \approx 2.53 \pm 0.08 \tag{3.55}$$

$$\sigma \approx 0.70 \pm 0.04 \tag{3.56}$$

In Figure 3.20, we attempt to produce collapse of $\rho(s)$ for $L = 300$ lattices using these estimates of τ and σ. To produce this plot, we need to discard data points for small cluster sizes ($s < 30$), because they are nonuniversal, and for large cluster sizes ($s > 100$), because they are noisy. Once we do this, the collapse works fairly well, indicating that we have found a consistent set of critical exponents obeying the necessary scaling relations.

3.4.2.2 Comments on the Glass-Insulator Transition

Before proceeding further, we should note that our numerics do not accurately capture the boundary between the Mott insulator and the glass. The "theorem of inclusions" of Pollet et al. suggests that we should expect glassy behavior to occur whenever the ratio of the largest bare Josephson coupling to the lowest bare charging energy $\frac{J_{\max}}{U_{\min}}$ exceeds the value of the clean transition, because this condition allows for the presence of regions in which the system locally looks like it is in the superfluid phase [113]. However, in the SDRG treatment, some distributions that satisfy this criterion still produce Mott insulating behavior out to the largest lattice sizes that we investigate. Because the glassy phase occurs due to rare regions or Griffiths effects, a finite size system will only look glassy if it is large enough for the rare regions to appear. This suggests that some choices of parameters that yield Mott insulating be-

havior on finite size lattices may actually yield glassy behavior in the thermodynamic limit. Of course, this difficulty necessarily afflicts all numerical methods (except those that rely on mean-field type approximations [25]), since they are confined to operate on finite size systems.

We will not comment on this transition further, but we take this opportunity to refer the reader to chapters by Krüger et al. and Pollet et al., which present two viewpoints on the Mott insulator to Bose glass transition [86, 113].

3.4.3 Identifying the Glass

Finally, we return to the question of the identity of the glassy phase. Is the phase a Bose glass or a Mott glass? A definitive diagnosis requires a measurement of the compressibility (1.20) The compressibility is more subtle to measure than the quantities that we have already discussed. Strictly speaking, any finite size system is gapped and therefore incompressible. On the other hand, in the thermodynamic limit, the gap can vanish and the compressibility need not be zero.

How can we measure the compressibility of the glassy phase in the RG? In Figures 3.17 and 3.18, we presented data for the density (number per unit area) of clusters with a given gap Δ. With this density profile in hand, we can calculate the density

of particles introduced to the system by a small chemical potential shift μ:

$$\begin{aligned} \rho_{\text{ex}} &= \int_0^\mu d\Delta \rho(\Delta) n(\Delta) \\ &= \int_0^\mu d\Delta \rho(\Delta) \lfloor \frac{\mu}{2\Delta} - \frac{1}{2} \rfloor \\ &\approx \int_0^\mu d\Delta \rho(\Delta) \left(\frac{\mu}{2\Delta} - \frac{1}{2} \right) \end{aligned} \tag{3.57}$$

Here, $n(\Delta)$ is the number of particles added to a cluster with gap Δ if the chemical potential is μ. If $\rho(\Delta)$ stays finite as $\Delta \to 0$, the integral is divergent, and the system is infinitely compressible. Suppose alternatively that $\rho(\Delta)$ vanishes as Δ^β for small Δ. Then:

$$\int_0^\mu d\Delta \Delta^\beta \left(\frac{\mu}{2\Delta} - \frac{1}{2} \right) = \frac{\mu^{\beta+1}}{2\beta(\beta+1)} \tag{3.58}$$

The derivative of the integral vanishes at $\mu = 0$ for $\beta > 0$. Thus, the system is incompressible. Comparing to the data shown in Figures 3.17, we see that there is no numerical evidence for a finite compressibility in the glassy phase; the gap density appears to vanish even faster than a power law as $\Delta \to 0$. This is consistent with the behavior of $\rho(s)$ in equation (3.22), because Δ is expected to scale as s^{-1}. Hence, the numerics imply that the Mott glass intervenes between the Mott insulator and the superfluid in this model.

At first glance, the preceding argument may be troubling. Due to the shrinking of the low Δ peak in panel (b) of Figure 3.18, the gap density $\rho(\Delta)$ also appears to vanish as $\Delta \to 0$ in the superfluid phase. The caveat is that it is necessary to more carefully evaluate the competing effects of the shrinking and the propagation. The

low Δ peak in Figure 3.18 represents the macroscopic superfluid clusters that form in the superfluid phase. These clusters do not appear in proportion to the surface area of the sample, as is the case for rare-regions clusters; instead, one such macroscopic cluster appears in each sample. Therefore, the density of macroscopic clusters will go as $\frac{1}{L^2}$, and this is responsible for the shrinking of the low Δ peak. The propagation of the peak, meanwhile, reflects the fact that the gap closes as L^{-2}. For a fixed choice of μ, the number of bosons that will be added to these macroscopic clusters scales as:

$$\frac{\mu}{2\Delta} - \frac{1}{2} \propto \mu L^2 \tag{3.59}$$

for large L. Then, the *density* of particles introduced to the system is:

$$\rho_{\text{ex}} \propto \frac{1}{L^2} \times \mu L^2 = \mu \tag{3.60}$$

This directly implies that the compressibility (1.20) is a constant in the thermodynamic limit, so we recover the expected result that the superfluid phase is compressible.

3.5 Conclusion

While the interplay of disorder and interactions in bosonic systems has attracted considerable interest for nearly three decades, the dirty boson problem remains a fertile source of intriguing physics. In this chapter, we have investigated a particular model of disordered bosons, the two-dimensional rotor (or Josephson junction)

model. Our SDRG analysis suggests the presence of an unstable finite-disorder fixed point of the RG flow, near which the coupling distributions flow to universal forms. Furthermore, the SDRG indicates the presence of three phases of the model: the Mott insulating and superfluid phases of the clean model are separated in the phase diagram by an intervening glassy phase[8]. The unstable fixed point governs the transition between the superfluid and this glassy phase, and the transition is driven by a kind of percolation. The RG procedure also provides evidence that this glassy phase is, in fact, the incompressible Mott glass.

Our work is a numerical extension into two dimensions of the one-dimensional study by Altman, Kafri, Polkovnikov, and Refael [7]. The 2D fixed point, however, differs from the 1D fixed point in an important way. The 1D fixed point occurs at *vanishing* interaction strength (all charging energies $U_j = 0$). Thus, it corresponds to a completely classical model and reveals that the superfluid-insulator transition can be tuned by varying the width of the Josephson coupling distribution at arbitrarily small interaction strength. The 2D fixed point is, in contrast, fully quantum. Indeed, in the critical distributions generated by the SDRG, the charging energy distribution is peaked near the RG scale while the Josephson coupling distribution is peaked well below.

On the other hand, the fixed point that we have identified in this chapter is similar to its 1D counterpart in that it does not exhibit infinite randomness. Finite disorder

[8]As we noted in Section 1.4.1.2 of Chapter 1, there is, in principle, also a Griffiths phase with large rare regions of insulator embedded in a superfluid bulk. Indeed, we have seen numerical evidence for such a phase. However, since this phase is characterized by macroscopic superfluidity, it is conventionally just referred to as a superfluid.

fixed points are not optimal settings for strong-disorder renormalization analyses, because the procedure does not become asymptotically exact near criticality and is, in this sense, an uncontrolled approximation. We have proceeded with such an analysis anyway. In doing so, we have found a robust fixed point controlling the superfluid-insulator transition and phases exhibiting reasonable physical properties. While this may be surprising given the perils of the method, we have attempted to argue for the appropriateness of the method, as an *approximation*, through an analysis of the RG steps in light of the forms of the fixed point distributions $P_{\rm univ}(U)$ and $P_{\rm univ}(J)$. We certainly acknowledge that there are other subtleties due to the lack of infinite randomness; for example, the notion of a superfluid cluster is not completely sharp, and consequently, percolation of superfluid clusters can only be an approximate picture of the transition [71]. Nevertheless, the structure of the fixed point Josephson distribution (3.15) suggests that the picture may be a good approximation, and we take this opportunity to remind the reader that we extensively discuss the reliability of the RG, in light of the properties of the fixed point, in Appendix 3.D. Moreover, the self-consistency of our numerical results, especially the striking universality and robustness of the unstable fixed point, gives us confidence that our SDRG analysis provides useful information about the system. With the potential caveats in mind, we therefore turn to exploring connections with other theoretical, numerical, and experimental work.

The Mott glass phase of the two-dimensional model is the straightforward analogue of the phase found in one dimension by Altman et al. The charging gap, the energy

needed to add or remove a boson from the system, vanishes due to the presence of arbitrarily large rare regions of superfluid ordering. However, there is no true long range order because these rare regions grow subextensively with system size. If a small chemical potential shift is turned on, it becomes energetically preferable to add bosons somewhere in the system, specifically in the largest of the rare-region superfluid clusters. Nevertheless, these clusters do not occur with sufficient number to produce a finite density of bosons, and the glass remains incompressible. In a Monte Carlo study of a model similar to ours, Prokofiev and Svistunov previously found evidence for a glassy phase in which the compressibility vanishes for this reason [114]. Moreover, the Mott glass that was identified by Roscilde and Haas in a related spin-one model also relies on the same mechanism [119]. The original proposal of Giamarchi, Le Doussal, and Orignac is, however, fundamentally distinct [57]. In their Mott glass, the charging gap remains finite, guaranteeing a vanishing of the compressibility; however, gaplessness is achieved through the closing of a mobility gap for transport of particles between nearby insulating and superfluid patches. Sengupta and Haas have argued that particle-hole symmetry, a crucial ingredient in the formation of our Mott glass, is not necessary for the realization of the phase through this alternative mechanism [124].

In the superfluid phase, true long range order emerges because the largest cluster scales with the size of the lattice. In this sense, this cluster is macroscopic. Despite this, near the transition, the macroscopic cluster may only occupy a small fraction of the total number of lattice sites. Because the clustering procedure can merge sites

that are not nearest neighbors in the bare lattice, the fraction of insulating sites may actually exceed the standard 2D square lattice percolation threshold. Even with such a large fraction of insulating sites, a superfluid phase can still exist because virtual tunneling processes can carry supercurrent through the insulating sites, allowing for macroscopic superfluidity on the "depleted" lattice that forms when the insulating sites are removed from the lattice by site decimation.

Nevertheless, the Mott glass to superfluid transition of our model should be contrasted with transitions that arise when disorder is introduced to a 2D square lattice by bond or site depletion [119, 51, 27, 143, 144, 121]. Models of the latter type only have the opportunity to form long range ordered phases when the underlying lattice is percolating. This percolation is purely classical and exhibits all the critical properties expected of standard site or bond percolation on a square lattice [130]. The superfluid-insulator transition is, in general, distinct from this transition; once the underlying lattice percolates, the bosonic model defined on that lattice may still exhibit Mott insulating, glassy, and superfluid phases [119]. In contrast, the only percolation-type process in our model is the one that actually drives the superfluid-insulator transition. The critical properties of this transition differ dramatically from those of classical 2D square lattice percolation, because the transition is not a purely geometric phenomenon. Instead, there are quantum tunneling processes overlaid on top of a geometric structure [102].

In Section 1.5, we identified several contexts in which experimental groups are currently probing disordered bosonic systems. In light of the work that we have

discussed in this chapter, the experiment on tin-decorated graphene by Allain, Han, and Bouchiat is particularly intriguing. These authors find a value of the exponent ν for the superconductor-insulator transition that is consistent with our estimate (3.48)[9]. Furthermore, they tie the emergence of superconductivity in their samples to a percolation process [5]. Hence, the transition described in this chapter may be describing this system.

We should also highlight again the work on nickel-based spin systems by Yu et al. These authors are among the first to report experimental evidence of the exotic Mott glass, and their realization of this phase is, in fact, the particle-hole symmetric variant. It is exciting that this phase can be observed, given that it requires a special symmetry. Qualitative aspects of our transition may be relevant to the spin system studied by Yu et al., but their spin systems are three dimensional [148]. Therefore, an extension of our work is necessary to characterize their system more thoroughly.

One immediate extension of our study would be to consider adding random filling offsets to the disordered rotor model, as Altman et al. did in a follow-up to their work on the 1D model [9]. The intuition from 1D suggests that such a model would contain a Bose glass phase. On the other hand, recent Monte Carlo work by Wang et al. indicates that the Mott glass might survive the substitution of exact particle-hole symmetry with *statistical* particle-hole symmetry [145]. In one dimension, Altman et al. found that the universality class of the transition (but not the identity of the glassy phase) is independent of the symmetry properties of the random filling offsets, but the situation may differ in $d > 1$; this remains to be understood.

[9] Allain et al. actually report the exponent $z\nu \approx 1.18$. Their analysis seems to assume $z = 1$.

Another interesting extension may be to study the rotor model defined on random networks. Suppose we do not begin with a square lattice but rather with a generalized network of mean connectivity $z = 4$. At its critical point, would such a model flow to the same fixed point as the model defined on a square lattice? The fact that the SDRG modifies the initial lattice structure into a more general network suggests that this may be the case for at least some types of random networks. Next, suppose we vary the mean connectivity from $z = 4$. Is there a range of connectivities for which random network models access our fixed point?

Perhaps it would be better to precede such investigations with a better characterization of the fixed point itself. In one dimension, Altman et al. were able to write down master equations for the RG flow, solve them to find fixed point distributions, and then verify numerically that these distributions are stable [7]. In two dimensions, a direct analytical approach is more difficult, and it remains to be seen whether such an approach is tractable. Our work provides suggestive numerical evidence regarding the forms of the universal distributions that characterize the critical point of the disordered rotor model (2.9). A Lyapunov analysis of these distributions, in which the RG is used as a tool to identify irrelevant directions in the space of possible distributions, could be a useful step in clarifying the critical forms still further. Then, analytically *verifying* these forms as attractor solutions of the RG flow may be an easier task than analytically identifying them would have been in the absence of any numerical guidance.

Recent work by Vosk and Altman suggests yet another direction for connecting

the results of the RG to experiment. These authors have derived the 1D version of the rotor model as an effective description of continuum bosons. In doing so, they have established a connection between the SDRG treatment of Altman, Kafri, Polkovnikov, and Refael and cold atom experiments on rubidium-87. Remarkably, the distributions that Vosk and Altman derive are of the same form as the fixed point distributions found by the SDRG [7, 141]. If such a treatment can be extended to the 2D case, that would be very valuable, both as a clarification of the critical behavior and as an indication of the relevance of the results to current experiments.

3.A Sum Rule vs. Maximum Rule

In this appendix, we present a short argument for why it may be preferable to use the *sum rule* (see equations (3.1) and (3.2)) instead of the *maximum rule* (see equation (3.3)). In dimensions greater than one, it should be easier to form ordered (e.g., superfluid) phases. Indeed, the transition for the clean rotor model occurs when $\frac{J}{U}$ is substantially smaller than one [131]. Suppose we begin with the clean model at its critical point and then disorder it by increasing some Josephson couplings and decreasing some charging energies. Suppose, further, that we do this such that that the greatest Josephson coupling $J_{\max}$ is still less than the weakest charging energy $U_{\min}$. Then, using the maximum rule, the strong-disorder renormalization procedure will predict no cluster formation at all. In other words, it will predict the ground state to be a Mott insulator, and this result seems inconsistent with the location of the clean

transition[10]. With the sum rule, the procedure has a mechanism to circumvent this inconsistency. The Josephson coupling scale can actually grow through decimation and compete with the charging energy scale. Thus, there can be cluster formation, and the procedure can find ground states that are glassy or superfluid, even when all Josephson couplings of the bare model are less than the minimum bare charging energy. This indeed occurs in the numerics, as we have noted while presenting the numerical data above.

The notion of the Josephson coupling scale increasing through the RG may be a source of concern to some readers. The increase actually corresponds to the generation of multiple effective couplings between two sites through different paths in the lattice. This still happens when the maximum rule is used, but it is hidden through the discarding of certain paths. If the coupling through each path is treated as an independent Josephson coupling, then the Josephson coupling scale does decrease as the renormalization proceeds. However, when it is time to determine the next decimation step, it is necessary to consider all of the couplings between any two sites. For this reason, it makes sense to sum all the couplings into one effective coupling between the sites.

3.B Measuring Physical Properties in the RG

Here, we work out two examples of how estimates of physical properties can be extracted from the RG procedure.

[10]There will, at least, be rare-regions of superfluid ordering in this situation [113].

3.B.1 Particle Number Variance

First, consider the quantity:

$$V_n \equiv \sum_j \langle \hat{n}_j^2 \rangle \tag{3.61}$$

This particle number variance gives the mean squared number fluctuation away from the large filling, summed over all sites in the lattice. When normalized by $\frac{1}{L^2}$, we find numerically that this quantity stays constant as the system size is increased for all choices of distributions and parameters. As such, this quantity is completely uninteresting for discriminating between phases of the model, but we do calculate it for comparison to exact diagonalization in Appendix 3.D.

The calculation of the particle number variance (3.61) is most straightforward when clusters do not form, so let us first consider the case where some site X is not clustered with any other site during the renormalization. At some stage in the procedure, the site is decimated away. The number fluctuation on site X is locked to zero at leading order with corrections incorporated in second order perturbation theory. An approximation to the ground state value of $\langle \hat{n}_X^2 \rangle$ can be obtained from the perturbative expansion of the state. In particular:

$$\langle \hat{n}_X^2 \rangle \approx \frac{1}{2} \sum_k \frac{J_{Xk}^2}{(U_X + U_k)^2} \tag{3.62}$$

When clusters do form, the calculation is trickier, but it can be performed by carefully keeping track of how the operator that we are targeting is written in terms

of the cluster and relative number operators (2.18) introduced in link decimation. To illustrate this, suppose we are trying to measure the operator:

$$\langle a_j \hat{n}_j^2 + a_k \hat{n}_k^2 \rangle = a_j \langle \hat{n}_j^2 \rangle + a_k \langle \hat{n}_k^2 \rangle \tag{3.63}$$

The factors a_j and a_k are just numbers. In (3.61), all these factors are one, but we will motivate the inclusion of more general a factors here shortly. If sites j and k are merged into a cluster, then we switch to the operators $\hat{n}_C$ and $\hat{n}_R$ to find:

$$\begin{aligned} a_j \hat{n}_j^2 + a_k \hat{n}_k^2 &= \frac{a_j U_k^2 + a_k U_j^2}{(U_j + U_k)^2} \hat{n}_C^2 + \frac{a_j U_k - a_k U_j}{U_j + U_k} \hat{n}_R \hat{n}_C \\ &\quad + (a_j + a_k) \hat{n}_R^2 \end{aligned} \tag{3.64}$$

During link decimation, the relative coordinate is specified, so the expectation value of the final term can be found immediately from the harmonic approximation:

$$(a_j + a_k) \langle \hat{n}_R^2 \rangle \approx (a_j + a_k) \frac{1}{2\gamma^{\frac{1}{2}}} \tag{3.65}$$

Furthermore, the harmonic theory also predicts that the expectation value of the term linear in $\hat{n}_R$ will vanish. The calculation of the term proportional to $\hat{n}_C^2$ must be deferred to later in the renormalization procedure. Thus, we keep the operator $\hat{n}_C^2$ in the portion of the sum (3.61) that remains to be evaluated, where it appears just like the $\hat{n}^2$ for a bare site, but multiplied by a renormalized a_C coefficient. This was the motivation for including the a factors; though the bare values of these factors

are all equal, different values can be generated through cluster formation. If the cluster is merged with another cluster in a future link decimation, we repeat the procedure above. When the cluster is finally decimated in a site decimation, the cluster's contributions to the sum (3.61) are calculated through equation (3.62) and then multiplied by the appropriate a factor.

The procedure that we have outlined above can run into a difficulty that we can anticipate by thinking about the two-site problem. Suppose two sites, labelled 1 and 2, are connected by a Josephson coupling J_{12}. Furthermore, suppose $U_1 > U_2$, but both charging energies are greater than J_{12}. Then, we would decimate site 1 first and obtain an estimate for the site's contribution to the particle number variance (3.61) through equation (3.62). Next, we would decimate site 2 and find that it does not contribute at all to (3.61) because there are no remaining links. However, the contribution of site 2 should, in fact, be equal to that of site 1, so our estimate is off by a factor of two. We can verify this by adopting cluster and relative coordinates (2.18) and then calculating (3.61) by doing perturbation theory for the relative coordinate. To partially resolve this difficulty, we can keep track of which sites are unclustered by the SDRG process. At the end, we can return to the original lattice and calculate the contributions of these unclustered sites to (3.61) using the bare couplings. For sites that *are* clustered by the RG, we reason that the main contribution to (3.61) comes from internal fluctuations of the cluster; thus, this correction may not be so important.

It is possible to raise another objection to our calculation of (3.61). The pertur-

bation theory that leads to the result in (3.62) incorporates the charging energies on sites neighboring site X. However, the perturbation theory leading to the RG rule (2.16) does not. How do we resolve this contradiction? When we calculate effective Josephson couplings in the RG, what we want to calculate is the effective rate of tunneling, once a boson has left one neighboring site and before it arrives at another, through the link-site-link system. For this purpose, it is appropriate to treat the sites neighboring site X as fictitious charging energy-free islands. On the other hand, when calculating observable quantities like (3.61), it is important to account for the fact that the ability to move a particle or hole from site X to a neighboring site also depends on how hard it is to charge the neighboring site.

3.B.2 Susceptibility

We now turn to our numerical calculation of the susceptibility (1.21). As we mentioned earlier, our numerical procedure for calculating the susceptibility closely mirrors the analytical approach of Chapter 2: we accumulate contributions to the susceptibility from site decimations as the RG proceeds. Where our numerical calculation differs is in the way that we incorporate internal fluctuations in the superfluid clusters. We describe this approach below[11].

To calculate χ, we consider how to calculate the expectation value:

$$\sum_j b_j \langle \cos(\phi_j) \rangle \tag{3.66}$$

[11]Note that this method is completely analogous to the one that we used for the particle-number variance above.

in the presence of an infinitesimal ordering field h. As with the a-factors in the calculation of the particle number variance, all the bare $b_j = 1$. When clusters form, the effective b-factor for the cluster can differ from unity.

If a bare site is decimated, then perturbation theory in h gives:

$$\langle \cos(\phi_X) \rangle = \frac{h}{U_X} \tag{3.67}$$

Since $b_X = 1$ for a bare site, (3.67) is the contribution of site X to the quantity (3.66).

When a link decimation joins two sites into a cluster, the corresponding terms in the sum (3.66) merge as:

$$\begin{aligned} b_j \cos(\hat{\phi}_j) + b_k \cos(\hat{\phi}_k) &= b_j \cos(\hat{\phi}_C + \mu_j \hat{\phi}_R) \\ &\quad + b_k \cos(\hat{\phi}_C - \mu_k \hat{\phi}_R) \\ &\approx b_j \cos(\hat{\phi}_C) \langle \cos(\mu_j \hat{\phi}_R) \rangle \\ &\quad - b_j \sin(\hat{\phi}_C) \langle \sin(\mu_j \hat{\phi}_R) \rangle \\ &\quad + b_k \cos(\hat{\phi}_C) \langle \cos(\mu_k \hat{\phi}_R) \rangle \\ &\quad + b_k \sin(\hat{\phi}_C) \langle \sin(\mu_k \hat{\phi}_R) \rangle \\ &= (b_j c_{DW,j} + b_k c_{DW,k}) \cos(\hat{\phi}_C) \end{aligned} \tag{3.68}$$

In this calculation, μ_j and μ_k are the ratios introduced in equation (3.9), and $c_{DW,j}$ and $c_{DW,k}$ are precisely the Debye-Waller factors given in equation (3.10). We can read off the renormalized b-factor for the cluster from the calculation above, and the

resulting expression is given in equation (3.19).

Next, it is important to note that the ordering field terms in the Hamiltonian (1.22) transform in the same way. In other words, the term coupling to $\cos(\hat{\phi}_C)$ in the Hamiltonian should appear multiplied by b_C after a merger. Physically, this corresponds to the fact that, when the cluster phase rotates, s_C bare phases rotate semi-coherently. Complete coherence is lost due to quantum fluctuations, which are accounted for by Debye-Waller factors. The factor b_C can be thought of as the effective number of rotors that are coherently ordering with the field. The energetic cost of ϕ_C straying from the direction of the ordering field is therefore amplified by this factor. When the cluster is finally decimated, the perturbation is amplified by this amount. Furthermore, $\cos(\hat{\phi}_C)$ appears in the sum (3.66) multiplied by b_C, and the total contribution of the cluster to the sum is:

$$b_C \langle \cos(\phi_C) \rangle = \frac{h b_C^2}{U_C} \tag{3.69}$$

From equations (3.67) and (3.69), the calculation of the linear response to an infinitesimal h (i.e., the susceptibility χ) follows immediately.

3.C The $P_i(U)$ Gaussian, $P_i(J)$ Power Law Data Set

In the work reported above, we have explored the superfluid-insulator transition of the disordered rotor model using several different species of disorder (see equations (3.11)-(3.14)). In doing so, we have exposed universal features of the critical behavior.

After providing numerical evidence of this universality however, we focus on data for one particular choice of initial distributions. In this appendix, we describe this choice of distributions in greater detail.

The choice of distributions in question first appears in Figure 3.7. The initial distribution of charging energies $P_i(U)$ is taken to be a Gaussian with center at U_0 and width $\sigma_U = 2$. Hence, the form of the distribution is:

$$P_i(U) \propto \exp\left[-\frac{(U-U_0)^2}{8}\right] \tag{3.70}$$

Recall that the Gaussian distribution is truncated at $3\sigma_U = 6$, so the distribution only has weight in the interval $U \in (U_0 - 6, U_0 + 6)$. We leave U_0 unspecified for the moment, because it is the parameter that we use to tune through the transition.

The initial Josephson coupling distribution $P_i(J)$ is a power law of the form $J^{-1.6}$. We choose the cutoffs so that the distribution has weight for $J \in (0.5, 100)$. This is a very wide power law distribution, and the corresponding flows begin well above the unstable finite disorder fixed point in the numerical flow diagrams. Explicitly:

$$P_i(J) \approx 0.413 J^{-1.6} \tag{3.71}$$

For this choice of distributions, we have acquired data for $U_0 = 400$, 20, 18, 16, 14, 12, 10, 9.8, 9.6, 9.4, 9.2, 8.8, 8.6, 8.4, 8.2, 8, 7.5, 7, and 6.5. To more finely target the transition, we have probed the interval $9.1 \geq U_0 \geq 8.9$ in increments of 0.01. We have always acquired data for $L = 25$, 50, 75, 100, 150, 200, and 300. In all cases, we

have pooled data for 10^3 disorder samples. The peak in the data for s_2 vs. U_0 gives the following estimate of the critical point:

$$U_{0,c} \approx 8.97 \pm 0.02 \tag{3.72}$$

Close to criticality, it is better to use a lower value of the thresholding parameter α. For several values of U_0 such that $10 \geq U_0 \geq 8$, we have run the RG with $\alpha_h = 10^{-5}$ and $\alpha_\ell = 5 \times 10^{-6}$ to test for convergence. Further away from criticality, we have instead used $\alpha_h = 5 \times 10^{-5}$ and $\alpha_\ell = 2.5 \times 10^{-5}$. Figure 3.22 shows a test of the convergence of the maximum cluster size $s_{\max}$ in the thresholding parameter α. We plot:

$$\upsilon = \frac{s_{\max}(\alpha_h)}{s_{\max}(\alpha_\ell)} \tag{3.73}$$

vs. the tuning parameter U_0. The ratio (3.73) is essentially always within two error bars of unity. We take this as an indication that physical properties have converged.

3.D Arguments for the Use of the SDRG in a Finite-Disorder Context

This appendix is devoted to exploring, in further detail, the validity of the RG procedure. We first expand upon the argument, introduced earlier in the chapter, for the reliability of the RG near criticality. Then, we move away from criticality and assess the performance of the RG in the various phases of the disordered rotor model

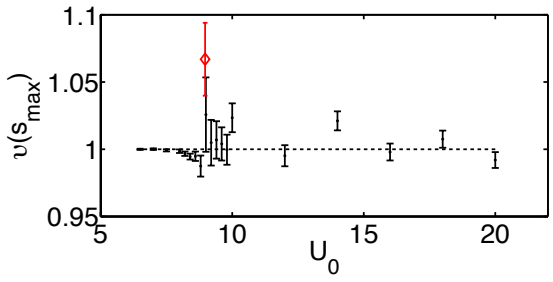

Figure 3.22: A test of the convergence of $s_{\max}$ in the thresholding parameter α. The variable υ is the ratio of the estimate of $s_{\max}$ for α_h, the less conservative value of the thresholding parameter, to α_ℓ, the more conservative value. We plot υ against the tuning parameter U_0. The closest data point to criticality ($U_0 = 8.97$) is indicated with a diamond. Note that $\upsilon > 1$ typically indicates convergence since less conservative thresholding (higher α) corresponds to throwing away more bonds and, therefore, biases the system away from cluster formation. Smaller values of α_h and α_ℓ are used in the vicinity of the transition. See the text of Appendix 3.C for details.

(2.9). Next, we focus on each of the RG steps, consider circumstances in which they may fail to capture important physics, and formulate tests to ensure that the RG is trustworthy in these situations. Finally, we present a comparison of the RG to exact diagonalization of small systems.

3.D.1 Review of the Argument for the RG at Criticality

Our confidence in the RG procedure near criticality rests on the form of the critical Josephson coupling distribution, reported in equations (3.15) and (3.18). Infinite randomness develops when $P(J) \propto J^{-1}$, and our numerical evidence suggests that the critical distribution of the disordered rotor model decays even more strongly [53]. Nevertheless, the critical distribution does not exhibit infinite randomness, because as seen in Figure 3.11, it is cut off from below. Recall that the lower cutoff of the "Josephson coupling distribution" is set by our choice to retain, *in statistics*, only the dominant $2\tilde{N}$ Josephson couplings, where $\tilde{N}$ is the number of sites remaining in the

effective lattice. Then, the appropriate way to interpret the distributions in Figure 3.10 is the following: penetrating any given site in the effective lattice, there are likely to be on the order of four Josephson couplings drawn from the depicted distribution. There may be other links penetrating the site, but these will be even weaker. We can estimate the typical strength of the four dominant links by comparing the median of the critical Josephson coupling distribution to the maximum. For the closest approach to criticality with the initial distributions described in Appendix 3.C, we find that the ratio $\frac{J_{\mathrm{typ}}}{\Omega}$ is approximately 0.11 ± 0.01 near the fixed point. Hence, the typical link is quantitatively weak compared to $\frac{J}{U} \approx 0.345$ at the clean transition [131].

The considerations above form a strong argument for the validity of the site decimation RG step. Here, we seek out the dominant effective charging energy in the effective lattice, and treat the links penetrating the site in perturbation theory. This perturbation theory is likely a very good approximation, because the Josephson couplings penetrating the site in question are usually extremely weak.

Now, we turn to the link decimation step, in which we seek out the largest Josephson coupling in the lattice and merge the corresponding sites into a cluster. Although, the Josephson coupling being decimated is the largest energy scale in the system, there is a high probability that all the other links penetrating the two sites being joined will be very small. However, the critical distribution of charging energies is *not* peaked at low U. The structure of the distributions plotted in Figure 3.12 suggests that it is quite likely that one or both sites being merged will have a charging energy of the same order as the RG scale, thus violating the strong-disorder hypothesis. Our treat-

ment of the quantum fluctuations of the relative phase of the cluster is based on the harmonic approximation (3.6). Is this approximation appropriate when the charging energies of the two-site problem are comparable in magnitude to the Josephson coupling? Alternatively, do the quantum fluctuations grow so large that the clustering becomes meaningless? We address this question as follows: using the fact that the remaining links are weak, we isolate the two site problem and solve it exactly, treating the remaining links via second order perturbation theory. Comparing the results of the RG with the exact solution, we find that, even in this worst case scenario, the RG produces reasonably accurate effective couplings. The evidence for this claim is given below in Section 3.D.3 of this appendix.

3.D.2 Reliability of the RG in the Phases

How reliable is the RG when we move away from criticality and into the phases of the disordered rotor model? To gain some insight into this question, we can consult Figure 3.23, which expands upon Figures 3.11 and 3.13 by plotting renormalized J and U distributions away from criticality.

Proceeding into the glassy regime, the arguments presented above for the validity of the RG near criticality generally become stronger. The primary reason for this is that the renormalized J distributions become progressively broader than at criticality. In the flow diagrams (Figures 3.6-3.9), this is reflected in the apparent divergence of $\frac{\Delta J}{\overline{J}}$. Consequently, the assumption of isolating local degrees of freedom becomes better as we get deeper into the glass. One complication is that the renormalized

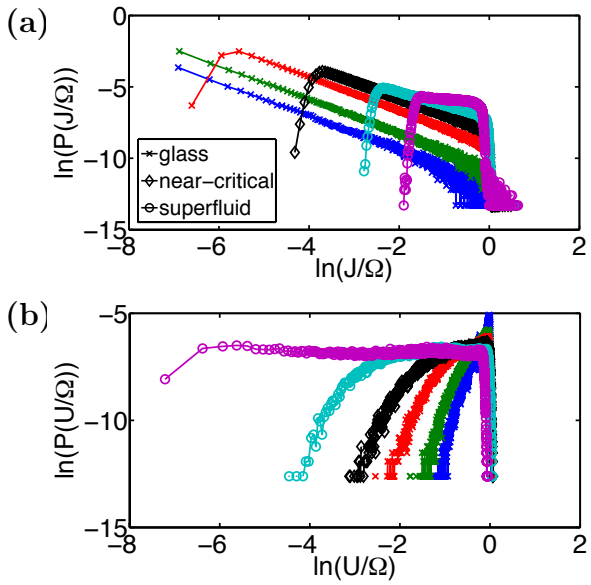

Figure 3.23: In panel (a), a sweep of renormalized J distribution through the glass and into the superfluid. The initial distributions are those described in Appendix 3.C: $P_i(U)$ is Gaussian and $P_i(J) \propto J^{-1.6}$. All data is taken for $L = 300$ lattices, and the renormalized distribution is computed when 300 sites remain in the effective lattice. The values of U_0 shown are 18, 12, 9.6, 8.97 (near-critical), 8.4, and 7.5. Panel (b) shows the corresponding sweep of the renormalized U distribution at the same stage of the RG.

U distribution becomes more strongly peaked near the RG scale. This may pose trouble for the link decimation step and makes it especially important to consider the reliability of this step when a strong J_{jk} couples two sites with charging energies $U_j \approx U_k \approx J_{jk}$. As mentioned previously, we will study this worst-case scenario in Section 3.D.3 of this appendix and find that the RG still works reasonably well.

Now, we turn to the superfluid phase. As we proceed away from criticality, the renormalized charging energy distribution becomes flatter and broader. The broadening of this distribution implies that the likelihood of encountering a strong Josephson coupling that connects sites with comparably strong charging energies decreases as we get deeper into the superfluid phase. However, the J distribution *also* seems to become flat deep in the superfluid, and this is problematic. For example, during link decimation, it may not be a reasonable approximation to isolate the two-site problem centered on the strongest Josephson coupling. As we proceed into the superfluid, it is necessary to be more dubious of the RG results; nearer to criticality however, the arguments used at the critical point are probably approximately valid.

The plots in Figure 3.23 attempt to elucidate systematics in the behavior of the renormalized distributions in the insulating and superfluid regimes, but this figure should be interpreted with some care. For the choice of distributions in Appendix 3.C, flows terminating in the insulating or superfluid regions of the flow diagram nevertheless propagate in the vicinity of the unstable fixed point along the way. This can be seen, for instance, in Figure 3.7. For certain choices of initial distributions, there can be flows towards the insulating or superfluid regimes which never propagate any-

where near the unstable fixed point. Consequently, the RG never has an opportunity to wash away the details of the initial distributions and allow the universal properties of the fixed point to emerge. Hence, the renormalized distributions generated along such flow trajectories are unlikely to exhibit the clean systematic properties seen in Figure 3.23.

3.D.3 Analysis of Potential Problems with the RG

Next, we address some potential difficulties with the arguments presented above for the reliability of the renormalization procedure in the critical region. These difficulties are rooted in the lack of strong randomness in the critical charging energy distribution.

Consider the lattice geometry shown in Figure 3.24. Suppose that all links except the Josephson coupling between sites 1 and 2 are perturbatively weak. Now, suppose further that J_{12}, U_1, and U_2 are comparable in magnitude, but J_{12} is the largest of the three. We would like to formulate a test of whether the RG appropriately handles this situation. Within the RG, a link decimation would merge sites 1 and 2 into a cluster. All links penetrating sites 1 and 2 would be modified by their corresponding Debye-Waller factors (3.10) $c_{\mathrm{DW},1}$ and $c_{\mathrm{DW},2}$. Then, because all the remaining links are assumed to be very weak, the cluster of sites 1 and 2 will be decimated, producing an effective coupling between sites 3 and 5:

$$\tilde{J}_{35,RG} = c_{\mathrm{DW},1} c_{\mathrm{DW},2} \frac{J_{13} J_{25} (U_1 + U_2)}{U_1 U_2} \tag{3.74}$$

Another approach to calculating this effective coupling would be the following: take the two-site problem of sites 1 and 2 and exactly diagonalize it. Then, to leading order, sites 1 and 2 should be locked into the two-site ground state, with perturbative corrections coming from the Josephson couplings J_{13}, J_{14}, and J_{25}. This perturbation theory leads to an effective coupling through the sites 1 and 2. This alternative procedure is perhaps more appropriate, because it does not presuppose the harmonic approximation. In Figure 3.25, we assess how much of an error we make by adopting the harmonic approximation. Holding J fixed, we sweep $U = U_1 = U_2$ through J, comparing the RG with the alternative method outlined above. We see that the usual RG performs reasonably well, implying that the harmonic approximation is safe.

Finally, we consider another potentially dangerous scenario. We return to the lattice shown in Figure 3.24. Now, we assume that J_{12} is greater than all other Josephson couplings, but it too is much weaker than the charging energies U_1 and U_2. In particular, $\frac{J_{12}}{U_2} = 0.05$. Then, we sweep U_1 such that it passes through a regime where $|U_1 - U_2| < J_{12}$. The danger here is that the RG may ignore resonance effects associated with this region. Within the usual RG, sites 1 and 2 will be decimated in turn to give:

$$\tilde{J}_{35,\mathrm{RG}} \approx \frac{J_{13} J_{12} J_{25}}{U_1 U_2} \tag{3.75}$$

$$\tilde{J}_{34,\mathrm{RG}} \approx \frac{J_{13} J_{14}}{U_1} \tag{3.76}$$

where we ignore subleading corrections coming from potential applications of the sum rule, depending on the order of decimation of sites 1 and 2. To consider potential

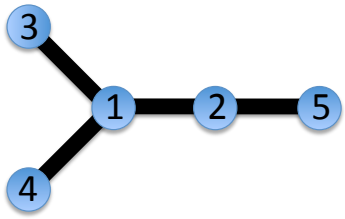

Figure 3.24: The graph structure for the tests reported in Figures 3.25 and 3.26. The links J_{13}, J_{14}, and J_{25} are assumed to be perturbatively weak. The charging energies U_1 and U_2 and the Josephson coupling J_{12} are varied to explore potentially troublesome scenarios.

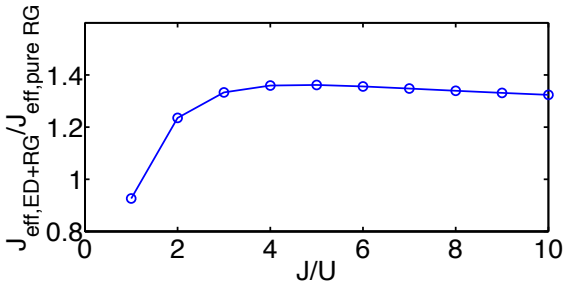

Figure 3.25: In this test, $J = J_{12}$ is assumed to be the strongest coupling in the system, but $U = U_1 = U_2$ may be of the same order. An effective coupling between sites 3 and 5 is calculated using two methods. One is the usual RG scheme used in this chapter. Another is a hybrid exact diagonalization and RG scheme: The two-site problem of sites 1 and 2 is exactly diagonalized. Then, the resulting cluster is decimated away, and perturbation theory is used to calculate an effective coupling between sites 3 and 5. The two candidate values for the effective coupling J_{35} are compared in the plot, as a function of $\frac{J}{U}$.

resonance effects, we can also implement the same hybrid exact diagonalization and RG procedure that we used above. In Figure 3.26, we compare the two methods and find excellent agreement.

3.D.4 SDRG vs. Exact Diagonalization

As a final test of the RG procedure, we now compare the RG to exact diagonalization of small systems. We truncate the possible number fluctuation on each site

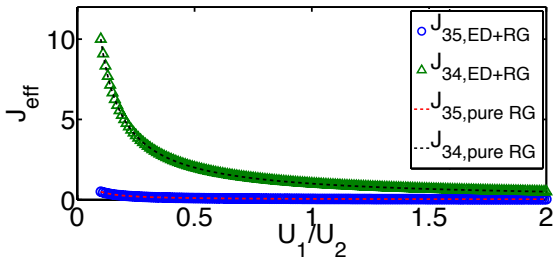

Figure 3.26: In this test, $J_{12} = 0.05$ and $U_2 = 1$. Hence, J_{12} is relatively quite weak. However, we vary U_1 such that it passes through a regime where $|U_1 - U_2| < J_{12}$, where there may be a danger of resonance effects. We calculate the effective couplings J_{34} and J_{35} using two schemes. One is the usual RG scheme used in this chapter. Another is a hybrid exact diagonalization and RG scheme: the two-site problem of sites 1 and 2 is exactly diagonalized. Then, the resulting cluster is decimated away, and perturbation theory is used to calculate an effective coupling between sites 3 and 4 and between sites 3 and 5. The effective couplings predicted by the two methods are compared in the plot, as a function of $\frac{U_1}{U_2}$. No resonance effects are observed.

to $n_j = -1, 0, 1$, interpret these three values as possible z-axis spin projections of a spin-one object, and in doing so, arrive at a "spin-one" model:

$$\hat{H}_{\text{rot-spin}} = -\sum_{\langle jk \rangle} \frac{J_{jk}}{2} (\hat{S}_j^+ \hat{S}_k^- + \hat{S}_j^- \hat{S}_k^+) + \sum_j U_j \hat{S}_{zj}^2 \tag{3.77}$$

The Hilbert space of this spin-one model grows with the size of the lattice as 3^{L^2}. The particle number conservation of the rotor model manifests here as total spin conservation along the z-axis. This means that we can partition the Hilbert space into different total spin sectors and diagonalize the sectors separately. For most ground state expectation values, we just need to diagonalize the total spin zero sector, and to calculate a charging gap, the only additional diagonalization needed is for the total spin one sector. Despite these simplifications, computational limitations restrict us to studying 3×3 lattices using CLAPACK. Testing the RG against exact diagonal-

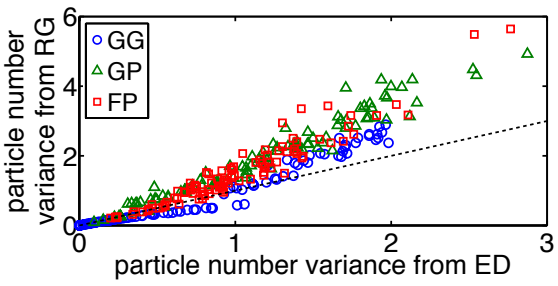

Figure 3.27: A sample-by-sample comparison of the particle number variance predictions from exact diagonalization and from the renormalization procedure. The disordered couplings are drawn from three different choices of the distributions, with 100 samples per distribution type. See the text of Appendix 3.D for details on the distribution choices GG, GP, and FP. Also pictured is the coincidence line $y = x$, along which the points would ideally fall for full quantitative agreement.

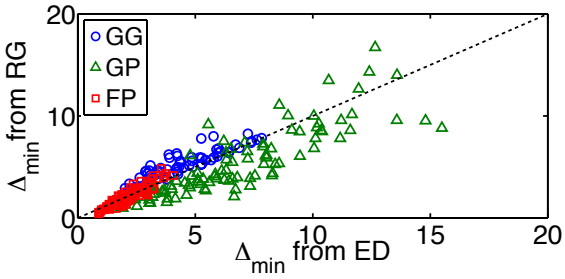

Figure 3.28: Same as Figure 3.27, but the quantity being compared is the charging gap Δ_{min}.

ization cannot directly tell us about the reliability of the RG at criticality, because exact diagonalization is limited to very small system sizes. However, a comparison with exact diagonalization *can* tell us how well the RG captures information about small patches of a larger lattice, and this information is potentially quite valuable for building confidence in the RG.

Another complication arises precisely due to the Hilbert space truncation: the spin-one model may not always approximate the full rotor model well. This is especially true when there is cluster formation, because then the rotor model has more of an opportunity to access particle number fluctuations that exceed 1 in magni-

tude. In other words, the SDRG and the exact diagonalization of the spin-one model constitute *different* approximations to the behavior of our dirty boson model. We cannot expect the two approximations to show full quantitative agreement, but we proceed with the comparison, despite its limitations, with the hope of at least seeing qualitative correspondence between the two methods.

Our general approach to comparing the RG with exact diagonalization will be to measure physical quantities, on individual 3×3 samples, using both methods and assess if there is a correlation between the predictions. The first quantity that we compare is the particle number variance V_n (3.61)[12]. We also compare the charging gap $\Delta_{\min}$, the minimum energy needed to add a particle or hole to the system. This quantity is typically estimated during site decimation. The logic behind site decimation is that the charging energy for some site X is greater than all other scales in the problem, and therefore, the site can be disconnected from the rest of the lattice to leading order. Then, the charging energy gives an estimate for the local charging gap on that site. During the RG, we find many such charging gaps from the various site decimations. The minimum among all of these gives an estimate for the charging gap for the whole system. This minimum is always given by the charging energy of the final remaining site, so an estimate of $\Delta_{\min}$ can be simply obtained by renormalizing down to a single site problem and measuring the charging energy of the remaining site. However, for the purposes of comparison to exact diagonalization, we find that we obtain better quantitative agreement between the two methods if we renormal-

[12]The interested reader may consult Appendix 3.B to see how this quantity is calculated during the renormalization procedure.

ize down to an effective two-site problem and then perform exact diagonalization on that system. Exactly diagonalizing the $n_{\text{tot}} = 0$ and $n_{\text{tot}} = 1$ sectors of this two-site problem then yields a charging gap for the system. In this exact diagonalization, we need not truncate the on-site number fluctuation to $n_X = -1, 0, 1$. Instead, in the numerics, we typically truncate to $n_X = -100 \ldots 100$.

In Figures 3.27 and 3.28, we present comparisons for three different data sets. In the first data set, we use take $P_i(U)$ and $P_i(J)$ to be Gaussian. We fix $U_0 = 10$ and $\sigma_U = 3$. Then, we randomly sample J_0 in the interval $(0, 5)$ and σ_J in the interval $(0, \frac{J_0}{3})$. The aim is to approximate some of the environments that the RG encounters in runs such as those reported in Figure 3.6. The second data set uses the distributions described in Appendix 3.C: $P_i(U)$ is Gaussian and $P_i(J) \propto J^{-1.6}$. We randomly sample $U_0 \in (6.5, 20)$. Here, the motivation is to look at the types of environments that the RG encounters when it approaches the unstable fixed point from above. In the final data set, we try to mimic 3×3 patches that the RG might encounter near criticality. To this end, the initial J distribution is fixed to a power law $P(J) \propto J^{-1.16}$ (see equation (3.18)) and the cutoffs are chosen so that the ratio of J_{min} to J_{max} is approximately that observed in panel (b) of Figure 3.13. The distribution $P_i(U)$ is flat with $U_{\text{max}} = J_{\text{max}}$ and with U_{min} randomly sampled such that the ratio of U_{min} to U_{max} lies in (e^{-2}, e^{-1}). Figures 3.27 and 3.28 identify these three data sets with the labels GG, GP, and FP respectively.

Both figures show that the predictions of the RG are clearly correlated with the predictions from exact diagonalization, although the level of quantitative correspon-

dence varies. For the particle number variance, quantitative agreement is lost at higher values of the variance, essentially corresponding to cases in which there is clustering. One potential source of error could be the Hilbert space truncation of the exact diagonalization, although the structure of the harmonic ground state (3.6) makes it unlikely that this could account for the entire discrepancy. Nevertheless, these comparisons suggest that the RG is retaining useful information about the system.

Chapter 4

Many-Body Localization in a Quasiperiodic System

4.1 Introduction

4.1.1 Motivation

As we noted in Chapter 1, the many-body localization (MBL) transition is a very unusual type of phase transition that, in many ways, lies outside the usual paradigm of quantum criticality [17]. Therefore, the original prediction by Basko, Aleiner, and Altshuler has motivated many numerical studies aimed at clarifying the existence and nature of this transition. Most of this work has focused on the disordered problem. On the other hand, we also saw in Chapter 1 that single-particle localization does not require disorder; the phenomenon also occurs in quasiperiodic models of the Aubry-André type (1.16) [13]. Furthermore, we noted that modern experiments directly realize the Aubry-André model and that the noninteracting transition has been observed in cold atoms and photonic waveguides. In both of these experimental contexts, it is possible to explore nonequilibrium dynamics and study the role of in-

teractions in transport [96, 90]. This means that many-body localization could very well become experimentally relevant, and consequently, it is imperative to determine whether quasiperiodic models can also exhibit an MBL transition.

This chapter is devoted to numerically addressing precisely that issue. More concretely, suppose we begin with a half-filled, one-dimensional system of fermions or hardcore bosons in a particular randomly chosen many-body Fock state, with some sites occupied and others empty. Such a configuration of particles is typically far from the ground state of the system. Instead, by sampling the initial configuration uniformly at random (i.e., without regard to its energy content), we are actually working in the so-called *infinite temperature limit*. If the particles are allowed to hop and interact for a sufficiently long time, the standard expectation is that the system should *thermalize*: that is, all microscopic states that are consistent with conservation laws should become equally likely and local observables should thereby assume some thermal distribution [117]. Can this expectation be violated in the presence of a quasiperiodic potential? In other words, can the system fail to serve as a good heat bath for itself? If so, can this be traced to the persistence of localization even in the presence of interactions?

4.1.2 Preview of the Results

Our work indicates that the answer to these questions is yes. We use numerical simulations of unitary evolution of a many-body quasiperiodic system to measure three kinds of observables in the limit of very late times: the correlation between the

initial and time-evolved particle density profiles, the many-body participation ratio, and the Rényi entropy. Our observations are consistent with the existence of two phases in the parameter space of our model that differ qualitatively in *ergodicity*. At finite interparticle interaction strength u and large hopping g, there exists a phase in which the usual assumptions of statistical mechanics hold. The initial state evolves into a superposition of a finite fraction of the total number of possible configurations, and consequently, local observables approximately assume their thermal values. This is the many-body ergodic phase. However, at small hopping g, there is a phase in which particle transport away from the initial configuration is not strongly enhanced by interactions. The system explores only an exponentially small fraction of configuration space, and local observables do not even approximately thermalize. This is the many-body localized phase. Figure 4.1 presents a schematic illustration of the proposed phase diagram. Although interactions induce an expansion of the ergodic regime, the localized phase survives at finite u, and consequently, there is evidence for a quasiperiodic MBL transition.

Both the many-body ergodic and localized phases differ *qualitatively* from their counterparts in the *noninteracting* AA model. The noninteracting extended phase is not ergodic, indicating that interactions are necessary for thermalization. Meanwhile, the many-body localized phase is expected to exhibit logarithmic growth of the bipartite entanglement entropy to an extensive value, albeit with subthermal entropy density. Such behavior is in fact consistent with the recent observations in the disordered problem [72, 15]. This growth is absent in the AA localized phase with-

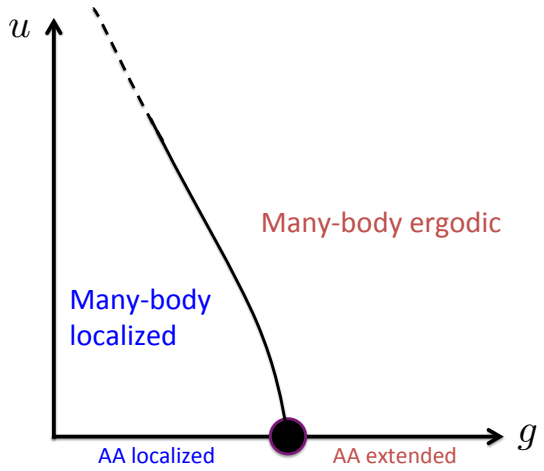

Figure 4.1: The proposed phase diagram of our interacting Aubry-André model at high energy density. Interactions convert the localized and extended phases of the AA model into many-body localized and ergodic phases and induce an expansion of the many-body ergodic phase. The phases of the interacting model differ qualitatively from their noninteracting counterparts. The differences are explained in Section 4.4 below.

out interactions. Despite this difference, the interacting and noninteracting localized phases are similar in their inability to thermalize the particle density.

There has certainly been substantial previous work on localization in many-body quasiperiodic systems. For instance, Vidal et al. [138] adapted the approach of Giamarchi and Schulz [58] to study the effects of a perturbative quasiperiodic potential on the low-energy physics of interacting fermions in one dimension. Very recently, He et al. [68] studied the ground state Bose glass to superfluid transition for hardcore bosons in a 1D quasiperiodic lattice. Our work differs fundamentally from these and many other studies precisely because it focuses on nonequilibrium behavior in the high-energy (infinite temperature) limit and argues that a localization transition can even occur in this regime.

4.1.3 Organization of the Chapter

We begin our study in Section 4.2 by introducing our interacting extension of the standard AA model. Since the MBL transition is a nonequilibrium phase transition, our goal is to follow the real time dynamics. To simplify this task, we describe a method of modifying the dynamics of our model, such that numerical integration of the new dynamics is somewhat easier than the original problem. In Section 4.3, we introduce the quantities that we measure in our simulations and present the numerical results. Then, in Section 4.4, we argue that our data points to the existence of many-body localized and many-body ergodic phases by proposing model late time states for each of these regimes and comparing to the numerical results from Section 4.3. Next, in Section 4.5, we extract estimates for the phase boundary from our data, motivating the phase diagram in Figure 4.1. Finally, we conclude in Section 4.6 by summarizing our results, drawing connections to theory and experiment, and suggesting avenues for future extensions of our work.

We relegate two exact diagonalization studies to Appendix 4.A. First, we examine the impact of our modified dynamics upon the single-particle and many-body problems. Second, we study the many-body level statistics of the interacting model. We find evidence for a crossover between Poisson and Wigner-Dyson statistics, consistent with the usual expectation in the presence of a localization transition [126].

4.2 Model and Methodology

In this section, we motivate and introduce our model and our numerical methodology for studying real time dynamics.

4.2.1 The "Parent" Model

We would like to consider one-dimensional lattice models of the following general form:

$$\hat{H}_{\text{pm}} = \sum_{j=0}^{L-1} \left[h_j \hat{n}_j + J(\hat{c}_j^\dagger \hat{c}_{j+1} + \hat{c}_{j+1}^\dagger \hat{c}_j) + V \hat{n}_j \hat{n}_{j+1} \right] \tag{4.1}$$

Here, $\hat{c}_j$ is a fermion annihilation operator, and $\hat{n}_j \equiv \hat{c}_j^\dagger \hat{c}_j$ is the corresponding fermion number operator. The three terms in the Hamiltonian (4.1) then correspond to an on-site potential, nearest-neighbor hopping, and nearest-neighbor interaction respectively. For now, we leave the boundary conditions unspecified. In 1D, the Hamiltonians (4.1) for hardcore bosons and fermions differ only in the matrix elements describing hopping over the boundary. With open boundary conditions, the Hamiltonians (and consequently all properties of the spectra) are identical.

If we set $V = 0$ in the Hamiltonian (4.1) and take h_j to be genuinely disordered, we recover the noninteracting Anderson Hamiltonian. If we then turn on a finite $V = J$, we obtain a model that is related to the spin models that have been studied in the context of MBL [111, 15]. Alternatively, suppose we set $V = 0$ again and take:

$$h_j = h \cos(2\pi k j + \delta) \tag{4.2}$$

With a generic irrational wavenumber k and an arbitrary offset δ, we obtain the noninteracting AA model that we first encountered in Section 1.3.2.2 of Chapter 1 [13]. For our purposes, we would like to use an incommensurate potential of the form (4.2), with $h = 1$ and $g \equiv \frac{J}{h}$ and $u \equiv \frac{V}{h}$ left as tuning parameters to explore different phases of the model (4.1).

4.2.2 Numerical Methodology and Modification of the Quantum Dynamics

Probing the MBL transition necessarily involves studying highly excited states of the system, and this precludes the application of much of the extensive machinery that has been developed for investigating low-energy physics. Consequently, several studies of MBL have resorted to exact diagonalization or other methods involving similar numerical cost [109, 111, 101]. In our work, we also use a numerical methodology that scales exponentially in the size of the system. However, in order to access longer evolution times in larger lattices, we introduce a modification of the quantum dynamics. This modification is inspired by a scheme used previously by Oganesyan, Pal, and Huse in a study of classical spin chains [110]. There, at any given time, either the even spins in the chain were allowed to evolve under the influence of the odd spins or vice versa. This provided access to late times that would have been too difficult to access by direct integration of the standard classical equations of motion.

By analogy, we propose allowing hopping on *each bond* in turn. At any given

time, the instantaneous Hamiltonian looks like:

$$\hat{H}_m = La_m J(\hat{c}_m^\dagger \hat{c}_{m+1} + \hat{c}_{m+1}^\dagger \hat{c}_m) + \sum_{j=0}^{L-1} [h_j \hat{n}_j + V \hat{n}_j \hat{n}_{j+1}] \tag{4.3}$$

We will specify the value of a_m in Section 4.2.C below, where we discuss our choice of boundary conditions. The state of the system is allowed to evolve under this Hamiltonian for a time $\frac{\Delta t}{L}$, and this evolution can be implemented by applying the unitary operator:

$$\hat{U}_m = \exp\left(-i\frac{\Delta t}{L}\hat{H}_m\right) \tag{4.4}$$

One full time-step of duration Δt consists of cycling through all the bonds:

$$\hat{U}(\Delta t) = \prod_{m=0}^{L-1} \hat{U}_m \tag{4.5}$$

Note that, in (4.3), the hopping is enhanced by L because the hopping on any given bond is activated only once per cycle, while the potential and interaction terms always act. Therefore, the factor of L ensures that the average Hamiltonian over a time Δt has the form (4.1). The advantage of employing the modified dynamics is that the $\hat{H}_m$ only couple pairs of configurations, so preparing the $\hat{U}_m$ reduces to exponentiating order V_H two-by-two matrices, where V_H is the size of the Hilbert space. This is generally a simpler task than exponentiating the original Hamiltonian (4.1). Our scheme only constitutes a polynomial speedup over exact diagonalization, but that speedup can increase the range of accessible lattice sizes by a few sites.

The modified dynamics raise several important issues that should be discussed [99].

The periodic time-dependence of the Hamiltonian induces so-called "multiphoton" (or "energy umklapp") transitions between states of the "parent" model (4.1) that differ in energy by $\omega_H = \frac{2\pi}{\Delta t}$, reducing energy conservation to *quasienergy* conservation modulo ω_H. We need to question whether this destroys the physics of interest: does the single-particle Aubry-André transition survive, or do the multiphoton processes destroy the localized phase?

We take up this question in Appendix 4.A, where we present a Floquet analysis of the single-particle and many-body problems. We find that, for sufficiently small Δt, the universal behavior of the single-particle AA model is preserved. At larger Δt, multiphoton processes can strongly mix eigenstates of the single-particle parent model, increasing the single-particle density-of-states and destroying the AA transition. In the spirit of the earlier referenced work on classical spin chains [110], our perspective in this chapter is to identify whether MBL can occur in a model *qualitatively* similar to our parent model (4.1). Therefore, to explore dynamics on long time scales, we avoid destroying the single-particle transition, but still choose Δt to be quite large within that constraint.

In Appendix 4.A, we also examine the consequences of our choice of Δt for the quasienergy spectrum of the many-body model. Our results suggest that multiphoton processes do not, in fact, strongly modify the parent model's spectrum for much of the parameter range that we explore in this chapter[1]. This means that partial energy

[1]There is an exception to this statement: multiphoton processes do seem to play an important role deep in the ergodic phase, where the energy content of the system is especially high. See Appendix 4.A and the discussion of the time-dependence of the autocorrelator χ in Section 4.3.A for more details.

conservation persists in our simulations despite the introduction of a time-dependent Hamiltonian, and we need to keep this fact in mind when we analyze our numerical data below.

Finally, we note in passing that several recent studies have focused on the localization properties of time-dependent models [80, 100, 44], including one on the quasiperiodic Harper model [81], but that the intricate details of this topic are somewhat peripheral to our main focus.

4.2.3 Details of the Numerical Calculations

In studies of the 1D AA model, it is conventional to approach the thermodynamic limit by choosing lattice sizes according to the Fibonacci series ($L = \ldots 5,\ 8,\ 13,\ 21,\ 34 \ldots$) and wavenumbers for the potential as ratios of successive terms in the series (e.g., $k = \frac{13}{21}$ on a lattice of size $L = 21$) [134]. These values of k respect periodic boundary conditions while converging to the inverse of the golden ratio $\frac{1}{\phi} = 0.618033\ldots$. For any finite lattice, the potential is only commensurate with the entire lattice (since successive terms in the Fibonacci series are mutually prime), and the duality mapping of the AA model is always exactly preserved. For our purposes however, this approach offers too few accessible system sizes and complicates matters by generating odd values of L.

Instead, we found empirically that finite size effects are least problematic if we use exclusively even L, always keep the wavenumber of the potential fixed at $k = \frac{1}{\phi}$,

and set:

$$a_m = 1 - \delta_{m,L-1} \tag{4.6}$$

in equation (4.3), thereby forbidding hopping over the boundary[2]. Note that, with these boundary conditions, our model describes hardcore bosons as well as fermions. The bosonic language maintains closer contact with cold atom experiments [50]; the fermionic language is more in keeping with the MBL literature [17, 109].

Using the approach described above, we have simulated systems up to size $L = 20$ at half-filling. Our simulations always begin with a randomly chosen configuration (or Fock) state so that the initial state has no entanglement across any spatial bond in the lattice (i.e., each site is occupied or empty with probability 1). Except in the exact diagonalization studies of Appendix 4.A, we always set $\Delta t = 1$. We integrate out to $t_f = 9999$ and ultimately average the evolution of measurable quantities over several *samples*, where a sample is specified by the choice of the initial configuration and offset phase to the potential (4.2). The sample counts used in the numerics are provided in Table 4.1.

4.3 Numerical Results

We now introduce the quantities that we measure to characterize the different regimes of our model. We also present the numerical data along with some quali-

[2]To appropriately realize open boundary conditions, we should also turn off interactions over the boundary. When exploring different options for the boundary conditions, we varied J over the boundary and neglected to vary V. This is unfortunate in that it makes the model somewhat stranger. However, our boundary conditions are chosen for convenience anyway, and the numerics suggest that the choice of boundary conditions does not impact the essential physics discussed in this chapter.

L	N	V_H	samples
8	4	70	500
10	5	252	500
12	6	924	500
14	7	3432	250
16	8	12870	250
18	9	48620	250
20	10	184756	50

Table 4.1: For the various simulated lattice sizes L, the particle number N, the configuration space size V_H, and the number of samples used in the numerics. Note that we always work at half-filling.

tative remarks about the observed behavior. However, we largely defer quantitative phenomenology and modeling of the data to Section 4.4.

4.3.1 Temporal Autocorrelation Function

One signature of localization is the system's retention of memory of its initial state. Because we simulate the reversible evolution of a closed system, the quantum state of the entire system retains full memory of its past. Nevertheless, we may still ask if the information needed to deduce the initial state is preserved locally or if it propagates to distant parts of the system. A diagnostic measure with which to pose this "local memory" question is the *temporal autocorrelator* of site j:

$$\chi_j(t) \equiv (2\langle\hat{n}_j\rangle(t) - 1)(2\langle\hat{n}_j\rangle(0) - 1) \tag{4.7}$$

Here, the angular brackets refer to an expectation value in the quantum state. This single-site autocorrelator may be averaged over sites and then over samples (as defined

in Section 4.2.C) to obtain:

$$\chi(t;L) \equiv \left[\frac{1}{L}\sum_{j=0}^{L-1}\chi_j(t)\right] \tag{4.8}$$

The sample average is indicated here with the large square brackets. Typically, to reduce the effects of noise, we also average over a few time steps within each sample (i.e., perform *time binning*) before taking the sample average.

We can discriminate three qualitatively different behaviors of χ vs. t in our interacting model. Figure 4.2 shows examples of each of these behaviors at interaction strength $u = 0.32$. Panel (a) is characteristic of the low g regime, where the autocorrelator stays invariant over several orders of magnitude of time, and there is no statistically significant difference between time series for different L. At higher g, as in panel (b), the time series show approximately power law decay culminating in saturation to a late time asymptote. For the largest systems, the power law is roughly consistent with the diffusive expectation of $t^{-\frac{1}{2}}$ decay. The late time asymptote decays with L (as expected from energy conservation[3]) suggesting that the power law decay may continue indefinitely in the thermodynamic limit. Surprisingly, at still larger g,

[3]The statistical fluctuation of the total energy of the randomly chosen initial configuration is of order $\sqrt{L}$. Suppose the total energy is conserved by the dynamics. We can write $E/\sqrt{L} = x_0 + hA_0\cos\theta_0 = x_\infty + hA_\infty\cos\theta_\infty$. Here, the subscripts 0 and ∞ refer to the initial and late time states, x_0 and x_∞ are bounded random numbers capturing the expectation value of interactions (and hopping at late times), h is the non-random amplitude of the quasiperiodic potential, and A_0 and A_∞ are positive bounded amplitudes of the Fourier components at the wavevector k of the quasiperiodic potential. This ansatz implies a finite correlation between the random phases θ_0 and θ_∞. Therefore, one of the Fourier modes of χ remains correlated as $L \to \infty$, and we expect $\chi \sim \frac{1}{L}$ in the ergodic phase. Note that this argument truly applies only to the energy-conserving parent model. In fact, in our numerics, there is only partial energy conservation, and energy non-conserving events become more prevalent as u, g, or L is raised. This means that χ will generically decay faster than $\frac{1}{L}$ at large L in the ergodic phase.

there is a third behavior, exemplified by panel (c). For the largest lattice sizes, the power law era is not followed by saturation but by an extremely rapid decay. The rapid decay is most evident in the large g, large u regime, where the energy density of the parent model (4.1) is relatively large. This implies that this behavior might be tied to the multiphoton processes induced by periodic modulation of the Hamiltonian; correspondingly, it also implies that, for fixed g and u, we might be able to induce the appearance of the rapid decay by increasing Δt. We have tested this numerically, and the results support the connection to the energy non-conserving multiphoton processes. This suggests that there are only two distinct regimes of the parent model represented in Figure 4.2, differentiated by the L dependence of the asymptotic value of the autocorrelator. We will proceed under this working assumption.

The difference between these two regimes is brought out more clearly in Figure 4.3. We focus on a late time $t = t_{\text{test}}$ and probe $\chi(t_{\text{test}}; L)$ as a function of g for different lattice sizes. Panels (a)-(c) show data for $u = 0$, 0.04, and 0.64 respectively. All the panels show a "splaying" point of the χ vs. L curves, separating a high g regime in which $\chi(t_{\text{test}}; L)$ decays with L from a low g regime in which it does not. The value of g at this feature decreases monotonically with u. Most importantly, in each case, this value is robust to changing t_{test}; if we halve t_{test} from the value that appears in Figure 4.3, the feature appears at approximately the same value of g. This property of the data is very fortunate: in Section 4.4.C below, we will use the splaying feature in these plots to put a numerical lower bound on the transition value of g for different interaction strengths. Since time scales get very long near the

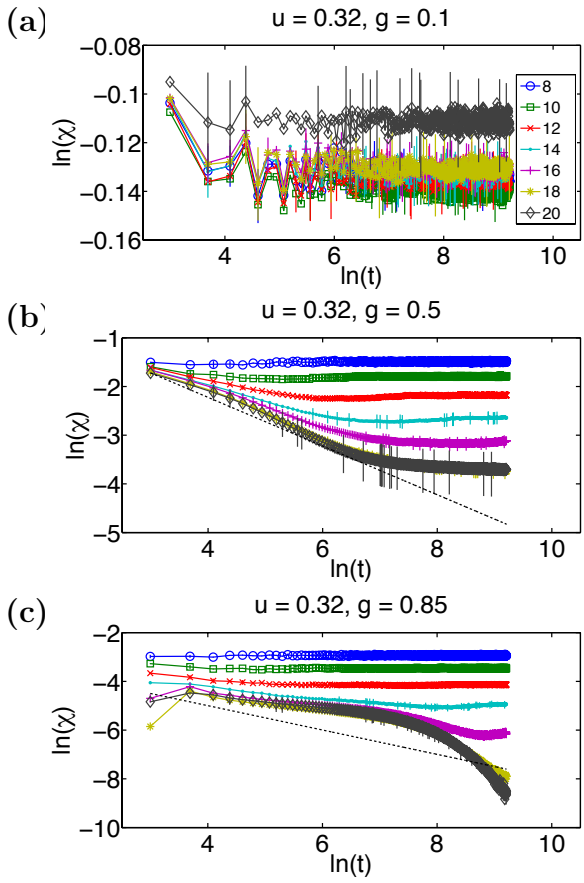

Figure 4.2: Three characteristic time series for the temporal autocorrelator with $u = 0.32$ and $\Delta t = 1$. In each panel, we show time series for a particular value of the hopping g. Only a few representative error bars are displayed in each time series. The legend refers to different lattice sizes L. The reference lines in panels (b) and (c) show diffusive $t^{-\frac{1}{2}}$ decay.

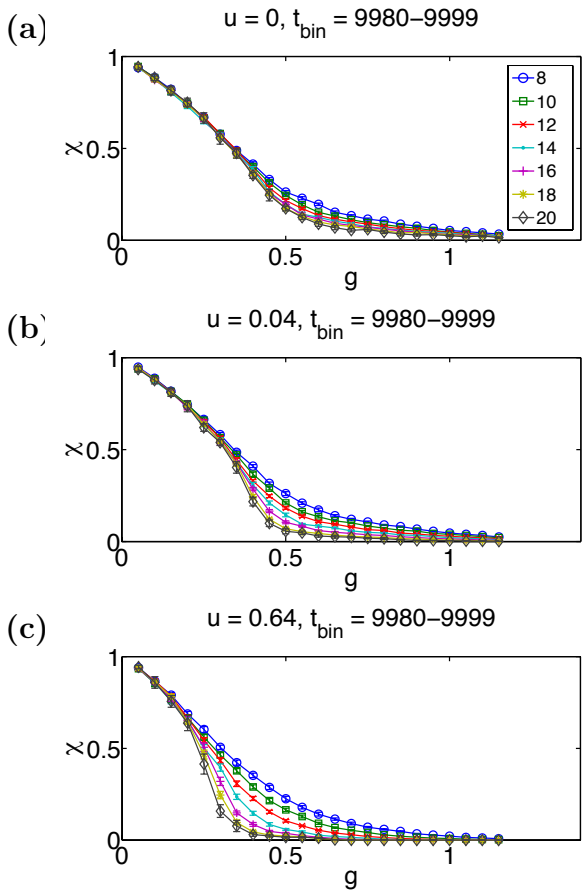

Figure 4.3: The value of χ in the latest time bin ($t = 9980 \ldots 9999$) plotted against g. In panels (a)-(c), $u = 0$, 0.04, and 0.64 respectively. The legend refers to different lattice sizes L.

transition, it is difficult to simulate out to convergence in this regime. Nevertheless, the fact that the value of g at the splaying feature remains fixed in time implies that we can deduce the phase structure from our finite-time observations.

4.3.2 Normalized Participation Ratio

One of the commonly used diagnostics for studying single-particle localization is the inverse participation ratio (IPR). This quantity is intended to probe whether quantum states explore the entire volume of the system and is often defined as the

sum over sites of the amplitude of the state to the fourth power: $\sum_j |\psi_j|^4$. Typically, the IPR is inversely proportional to the localization volume ξ^d in a single-particle localized phase and decays to zero as the inverse of the system volume in an extended phase.

We now describe how this quantity can be fruitfully exploited in the many-body context. Let c denote some specific configuration of N particles in L sites. Then, we can write the state of the system in the configuration basis as:

$$|\Psi(t)\rangle = \sum_{\{c\}} \psi_c(t)|c\rangle \tag{4.9}$$

The configuration-basis IPR is simply:

$$P(t;L) \equiv \left[\sum_c |\psi_c(t)|^4\right] \tag{4.10}$$

where the square brackets, as usual, denote a sample average. Interpreting $P(t;L)$ as the inverse of the number of configurations on which $|\Psi(t)\rangle$ has support, we now define the normalized participation ratio (NPR):

$$\eta(t;L) \equiv \frac{1}{P(t;L)V_H} \tag{4.11}$$

The quantity $\eta(t;L)$ then represents the fraction of configuration space that the system explores. We expect $\eta(t;L)$ to be independent of L at late times in the ergodic phase. In the many-body localized phase, we expect $\eta(t;L)$ to decay exponentially

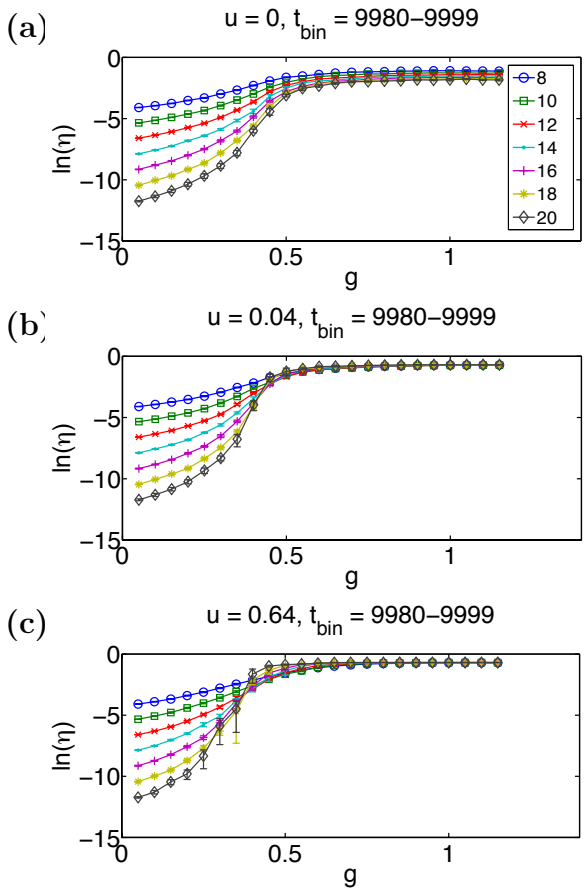

Figure 4.4: The value of η in the latest time bin ($t = 9980 \ldots 9999$) plotted against g. In panels (a)-(c), $u = 0$, 0.04, and 0.64 respectively. The legend refers to different lattice sizes L. See equation (4.11) for the definition of η. In the ergodic phase $\eta \approx 0.5$.

with L.

In Figure 4.4, we plot $\eta(t_{\text{test}}; L)$ vs. g for $u = 0$, 0.04, and 0.64. The figure reveals an important difference between the noninteracting and interacting models. At low g, both with and without interactions, η decays exponentially with L:

$$\eta \propto \exp(-\kappa L) \tag{4.12}$$

with $\kappa > 0$. More surprisingly, η also decays with L at *large* g in the noninteracting

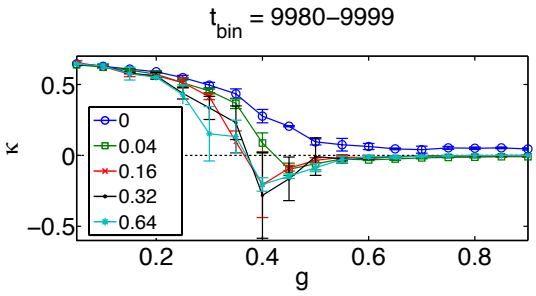

Figure 4.5: Estimates of κ from a fit of $\eta \propto e^{-\kappa L}$ in the latest time bin ($t_{\text{bin}} = 9980 - 9999$). The legend refers to different values of the interaction strength u.

case; all that happens is that κ becomes essentially independent of g. With even small interactions, however, η becomes system-size independent in the large g regime, following our ansatz for an ergodic phase. We bring out this point more clearly in Figure 4.5, in which we extract estimates for the decay coefficient κ for various values of the interaction strength. Thus, the extended phase of the noninteracting AA model appears to be a special, non-ergodic limit.

Before proceeding, we should caution that, in panels (b) and (c) of Figure 4.4, the collapse at high g looks very appealing because of the use of a semilog plot and would not be so striking on a normal scale. The axes have been chosen to highlight the exponential scaling at low g, which would not be as apparent if we simply plotted η vs. g. However, regarding the absence of perfect collapse at high g, note that the raw data for the IPR differ by several orders of magnitude for different values of the lattice size L. Given this, the coincidence of the order of magnitude of η for different values of L is already a good indication of the proposed scaling, and some corrections to this scaling should be expected given the modest accessible system sizes.

4.3.3 Rényi Entanglement Entropy

Unlike the normalized participation ratio, which provides a global characterization of the time-evolved state, bipartite entanglement is arguably a better proxy for whether a part of the system can act as a good heat bath for the rest. In the many-body ergodic phase, we expect the bipartite entanglement entropy to be a faithful reflection of the thermodynamic entropy. This implies an extensive entropy, pinned to its thermal infinite temperature value throughout the phase[4]. In contrast, in the many-body localized phase, we expect an extensive but subthermal entanglement entropy. This expectation is consistent with the results of three recent papers that focus on the behavior of entanglement measures in the many-body localized phase of the disordered problem [72, 15, 142]. These papers also study the time dependence of the entropy beginning from an unentangled product state. In the many-body localized phase, this growth is found to be slow, generically logarithmic in time. Because our model lacks disorder altogether, it may be interesting to explore the entanglement dynamics here as well. In what follows, we comment on the dynamics, but we primarily use the late time entanglement entropy as yet another tool to help distinguish between the many-body localized and ergodic phases.

Let subsystem A refer to lattice sites $0, 1, \ldots \frac{L}{2} - 1$, and let subsystem B refer to the remaining sites in the chain. We can compute the reduced density matrix

[4]This statement should be interpreted with some care. Quantum entanglement entropy measures, such as the Rényi entropy that we define in equation (4.14), carry information about the off-diagonal elements in the reduced density matrix. These terms have no classical analog and would not be considered in a thermodynamic calculation. This difference can result in discrepancies in the subleading behavior. For instance, consider our calculation of the bipartite Rényi entropy of the model state $|\Phi\rangle$ in Section 4.4.A: the quantum Rényi entropy is one bit lower than the Rényi entropy calculated by classical counting of configurations. A more precise analog of the classical entropy would thus be a "diagonal" entropy in which all off-diagonal elements of the reduced density matrix were neglected.

of subsystem A by beginning with the full density matrix $\hat{\rho}(t) = |\Psi(t)\rangle\langle\Psi(t)|$ and tracing out the degrees of freedom associated with subsystem B:

$$\hat{\rho}_A(t) \equiv Tr_B\{\hat{\rho}(t)\} \tag{4.13}$$

The sample-averaged order-2 Rényi entropy of subsystem A is then given by:

$$S_2(t; L) \equiv \left[-\log_2\left(Tr_A\{\hat{\rho}_A(t)^2\}\right)\right] \tag{4.14}$$

Both S_2 and the standard von Neumann entropy are expected to attain the same values in the ergodic phase; we choose to focus on the former to save on the computational cost of diagonalizing the reduced density matrix (4.13).

Our first task is to examine whether the putative localized phase of our model exhibits the same behavior that was observed with tDMRG [72, 15]. In panel (a) of Figure 4.6, we focus on a low value of g and plot S_2 vs. $\ln(t)$ for $L = 10$ lattices. At very early times, the time series all tend to coincide, reflecting the formation of short-range entanglement at the cut between the subsystems. Afterwards, the noninteracting time series saturates for several orders of magnitude of time, while the interacting time series show behavior that is consistent with logarithmic growth. In order to clearly establish the saturation that follows the slow growth, we have had to focus on small lattices. Panel (b) of Figure 4.6 shows data for large g. Here, the most striking difference between the noninteracting and interacting models lies in the saturation value of the entropy: the interacting model is substantially more

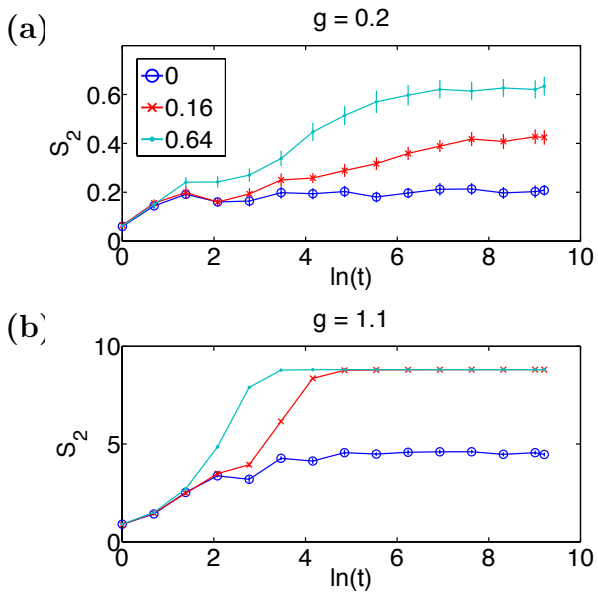

Figure 4.6: Example time series of the Rényi entropy for two values of the tuning parameter g. The legend refers to different values of the interaction strength u. Panel (a) shows data for $L = 10$ lattices at $g = 0.2$. Panel (b) shows data for $L = 20$ lattices at $g = 1.1$. In the localized regime, we need to use smaller lattices to see convergence Renyi entropy.

entangled, but the saturation value does not appear to depend on the value of u. We will see below that this is another indication that thermalization only occurs in the *interacting*, large g regime.

Figure 4.7 shows late time values of the Rényi entropy *density* plotted against the tuning parameter g. We first focus on the high g regime. In panel (a), $u = 0$, and $S_2(t_{\text{test}}; L) \propto L$ for large g. However, the entropy density is less than $\frac{1}{2}$, which is the thermal result when the system has ergodic access to all configurations consistent with particle number conservation. The situation is dramatically different in panels (b) and (c), where $u = 0.04$ and 0.64 respectively. At high g, the entropy actually looks *superextensive*. This is just a finite size effect, because the entropy is well fit to

a linear growth of the form:

$$S_2(t_{\text{test}}; L) = mL - S_{\text{def}} \tag{4.15}$$

where S_{def} is a constant deficit, typically around $1.15 - 1.3$. In Figure 4.8, we show that the slope $m \approx \frac{1}{2}$ at large g in the interacting problem. This implies that the entropy is thermal in the $L \to \infty$ limit, where the deficit S_{def} is negligible.

Now, we turn to the low g regime. Without interactions, the off-diagonal elements in the reduced density matrix (4.13) typically contain only a few frequencies originating from localized single particle orbitals immediately adjacent to the cut. The number of relevant orbitals is constant in L. As a result, the off-diagonal elements cannot fully vanish, and the reduced density matrix never thermalizes. The resulting entanglement entropy is independent of L as shown in the inset of panel (a). In the interacting problem, while the orbitals immediately adjacent to the cut still have roughly the same frequencies, the "spectral drift" (i.e., the spread of these lines due to sensitivity to the configuration of distant particles) allows for a much larger number of distinct and mutually incoherent contributions to offdiagonal elements of the reduced density matrix. These off-diagonal elements can dephase more efficiently, leading to a partial thermalization. This is the mechanism that likely underlies the extensive but subthermal entropy observed by Bardarson et al. [15]. For small L, our numerical results in the low g regime agree well with this expectation. For larger L, the slow dynamics of the entropy formation makes it difficult to observe saturation, both in our work and in the tDMRG study of Bardarson et al.

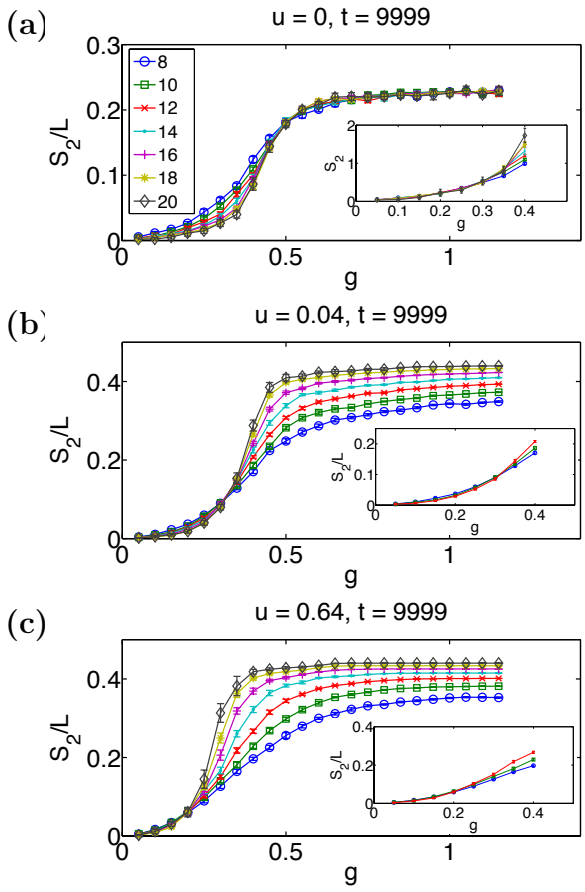

Figure 4.7: The value of $\frac{S_2}{L}$ at $t = 9999$ plotted against g. In panels (a)-(c), $u = 0$, 0.04, and 0.64 respectively. The legend refers to different lattice sizes L. In panel (a), the inset plot shows S_2 vs. g in the low g regime. In panels (b) and (c), the insets show $\frac{S_2}{L}$ vs. g for low L in the low g regime.

If the entropy eventually becomes extensive for all L, then the "crossing" feature that is present in panels (b) and (c) of Figure 4.7 would become a "splaying" feature, with the entropy *density* independent of L at small g. In any case, an interesting property of the data is that the values of g at the crossing features of the $S_2(t_{\text{test}}; L)$ vs. g plots are consistent with the locations of the splaying features in the corresponding $\chi(t_{\text{test}}; L)$ vs. g plots of Figure 4.3. This seems to be the case for all u. Thus, these features may be useful in locating the transition.

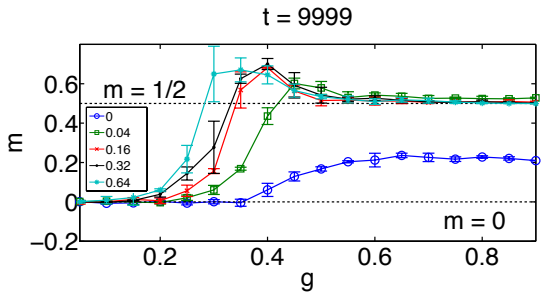

Figure 4.8: The estimated slope of S_2 vs. L at late times as a function of g. The legend refers to different values of the interaction strength u.

4.4 Modeling the Many-Body Ergodic and Localized Phases

Above, we presented numerical evidence that our interacting AA model contains two regimes that show qualitatively distinct behavior of the autocorrelator, normalized participation ratio, and Rényi entropy. Next, we will propose and characterize model quantum states that qualitatively (and sometimes quantitatively) reproduce the numerically observed late time behavior in the two regimes. These model states expose more clearly why the two regimes of our model are appropriately identified as many-body ergodic and localized phases.

4.4.1 The Many-Body Ergodic Phase

To model the behavior of the putative ergodic phase, we begin by writing down a generic model state in the configuration basis:

$$|\Phi\rangle = \sum_{\{c\}} \phi_c |c\rangle = \sum_{n=0}^{\frac{L}{2}} \sum_{\{c_A, c_B\}} \phi_{AB}^{(n)} |c_A^{(n)}, c_B^{(n)}\rangle \tag{4.16}$$

Here, the c refer to configurations of the full chain, whereas the c_A and c_B refer to configurations of the subsystems A and B, as defined in Section 4.3.C above. The superscripts on the configurations and expansion coefficients refer to the number of particles in subsystem A. Writing the state in terms of the subsystem configurations will be useful shortly, but for now we focus on the statistical properties of the amplitude ϕ_c. We assume that this amplitude is distributed as a complex Gaussian random variable:

$$p(\phi) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{|\phi|^2}{2\sigma^2}\right) \tag{4.17}$$

Within this distribution, $\langle|\phi|^2\rangle = 2\sigma^2$ and $\langle|\phi|^4\rangle = 8\sigma^4$. From these average values, it is possible to deduce that:

$$\sigma = \frac{1}{\sqrt{2V_H}} \tag{4.18}$$

for normalization and that the IPR is $P_\Phi = \frac{2}{V_H}$. This, in turn, implies:

$$\eta_\Phi = \frac{1}{2} \tag{4.19}$$

This result is reproduced *quantitatively* in the numerics in Figure 4.4.

Next, suppose we compute the reduced density matrix of subsystem A in the state $|\Phi\rangle$:

$$\hat{\rho}_A = \sum_n \sum_{\{c_A, c_{A'}, c_B\}} \phi^{*(n)}_{AB} \phi^{(n)}_{A'B} |c^{(n)}_A\rangle\langle c^{(n)}_{A'}| \tag{4.20}$$

To find the Rényi entropy, we need to compute the trace of the square of this operator:

$$Tr_A\{\hat{\rho}_A^2\} = \sum_n \sum_{\{c_A, c_{A'}, c_B, c_{B'}\}} \phi_{AB}^{*(n)} \phi_{A'B}^{(n)} \phi_{AB'}^{*(n)} \phi_{A'B'}^{(n)} \tag{4.21}$$

When we average over our distribution of amplitudes (4.17), only the coherent terms survive[5]:

$$\begin{aligned} Tr_A\{\hat{\rho}_A^2\} \approx & \sum_n \sum_{\{c_A, c_B, c_{B'}\}} \langle |\phi_{AB}^{(n)}|^2 |\phi_{AB'}^{(n)}|^2 \rangle \\ & + \sum_n \sum_{\{c_A, c_{A'}, c_B\}} \langle |\phi_{AB}^{(n)}|^2 |\phi_{A'B}^{(n)}|^2 \rangle \\ & - \sum_n \sum_{\{c_A, c_B\}} \langle |\phi_{AB}^{(n)}|^4 \rangle \end{aligned} \tag{4.22}$$

The final term accounts for the double counting of terms where $c_A = c_{A'}$ and $c_B = c_{B'}$ simultaneously. We now introduce the notation:

$$\gamma(P, Q) = \frac{P!}{Q!(P-Q)!} \tag{4.23}$$

and evaluate the expectation values in equation (4.21) to obtain:

$$Tr_A\{\hat{\rho}_A^2\} \approx \frac{2}{V_H^2} \sum_n \gamma \left(\frac{L}{2}, n \right)^3 \tag{4.24}$$

Finally, using a Stirling approximation to the combination function and a saddle-point

[5]Only the first term on the right-hand side of equation (4.22) would appear in a "classical counting" derivation of the thermodynamic entropy. The other two terms account for off-diagonal elements in the reduced density matrix (4.20). Please see footnote 53 for more details.

approximation for the sum, we find the entropy:

$$S_{2,\Phi} \approx \frac{L}{2} - \log_2\left(\frac{4}{\sqrt{3}}\right) \approx \frac{L}{2} - 1.2 \tag{4.25}$$

This is the same form observed in the numerics (4.15), and the deficit S_{def} lies in the observed range. Asymptotically in L, the entropy (4.25) is *maximal*, and this is exactly the expected behavior when the particle number thermalizes.

There is an important caveat to note here: we have argued above that, if multi-photon processes do not completely destroy energy conservation, then this can lead to relic autocorrelations at late times. This implies that the assumption of independent random amplitudes cannot be exactly correct on a finite lattice. However, the numerically observed relic autocorrelations decay with L, suggesting that our assumptions get better as the system size grows. Therefore, in the thermodynamic limit, this phase is truly thermal.

4.4.2 The Many-Body Localized Phase

Our model for the time-evolved state in the localized regime is founded upon the following intuition: there exists a length scale ξ, which is analogous to the single-particle localization length and beyond which particles are unlikely to stray from their positions in the initial state. Then, if we partition our lattice of length L into blocks of size ξ, exchange of particles between blocks is less important than rearrangements of the particles within each block. Consequently, the total number of configurations accessed by the state of the full system is approximately the product of the number

of configurations accessed within each block. This multiplicative assumption should be very safe in a localized phase. We additionally assume that, within each block, the dynamics completely scramble the particle configuration. If a certain block of length ξ contains n particles in the initial state, then the time-evolved state contains equal amplitude for each of the possible ways of arranging n particles in those ξ sites. In keeping with our numerical protocol, we randomly select the initial state from the space of all possible Fock states of a certain global particle number. Then, a block of ξ sites contains n particles with probability:

$$w(\xi, n) = \frac{\gamma(\xi, n)}{2^{\xi}} \left[1 + O\left(\frac{\xi^2}{L}\right)\right] \tag{4.26}$$

We will consider the limit $L \gg \xi \gg 1$, where we can approximate the probability by the first term. The assumptions proposed above motivate writing down a state of the form:

$$|\Lambda\rangle = \frac{1}{\sqrt{M}} \sum_{\{c_1, \ldots c_{\frac{L}{\xi}}\}}^{\sim} z\left(c_1, \ldots c_{\frac{L}{\xi}}\right) |c_1, \ldots c_{\frac{L}{\xi}}\rangle \tag{4.27}$$

where the tilde on the sum indicates that it should only run over configurations that are consistent with the initial distribution of particles among the blocks. The factors z are complex phases which depend upon the configuration, and M is a normalization which is equal to the total number of configurations represented in the state $|\Lambda\rangle$.

Before beginning our analysis of the state $|\Lambda\rangle$, we should note that, in contrast to our calculations in the ergodic phase, our goal in the localized regime will be to qualitatively tie the numerically observed large L behavior to the existence of the

length scale ξ. Unfortunately, we cannot achieve the quantitative accuracy of the ergodic model state $|\Phi\rangle$ with the localized toy-model described above.

We begin by estimating the autocorrelator between the initial state and the model time-evolved state $|\Lambda\rangle$. A nonzero autocorrelator emerges, because each block is only at half-filling *on average*. Fluctuations away from half-filling (in either direction) yield a positive typical value of the autocorrelator within a block. Indicating an average over the distribution (4.26) with an overline, we find the block value $\overline{\chi_{\text{block}}} \approx \frac{1}{L}$. This is also the average value for the whole system when $L \gg \xi$:

$$\chi_\Lambda \approx \frac{1}{\xi} \tag{4.28}$$

Next, to estimate the IPR, we need to compute the normalization factor M. We begin by estimating the number of explored configurations in each block. The average of the logarithm of the number of explored configurations within a block is:

$$\overline{\ln(M_{\text{block}})} \approx \ln\left(\sqrt{\frac{2}{\pi\xi}}2^\xi\right) - \frac{1}{2} \tag{4.29}$$

Then, using $\overline{\ln M} \approx \frac{L}{\xi}\overline{\ln M_{\text{block}}}$, we can estimate M itself as:

$$M \approx e^{\overline{\ln M}} \approx 2^L \left(\frac{\pi e \xi}{2}\right)^{-\frac{L}{2\xi}} \tag{4.30}$$

Using this normalization, we can estimate the NPR η_Λ:

$$\ln \eta_\Lambda \approx -\frac{L}{2\xi} \ln\left(\frac{\pi e \xi}{2}\right) + \frac{1}{2} \ln L + \frac{1}{2} \ln\left(\frac{\pi}{2}\right) \tag{4.31}$$

This qualitatively agrees with the numerically observed behavior (4.12) up to subleading corrections, and in the large-L limit:

$$\kappa \approx \frac{1}{2\xi} \ln\left(\frac{\pi e \xi}{2}\right) \tag{4.32}$$

Note that equations (4.28) and (4.32) imply a relationship between the scaling behaviors of χ and κ in the localized regime. This relationship is *not* reflected in our numerical data, in part because we cannot truly attain the limit $L \gg \xi \gg 1$. The numerically computed value of κ, for example, can contain finite size corrections of order $\frac{\ln(L)}{L}$ or $\frac{\xi^2}{L}$. Also, we must keep in mind that the state $|\Lambda\rangle$ is just a toy model that does not capture fine details of the time-evolved states in this regime. Thus, we must be content with reproducing the qualitative behavior of each measurable quantity individually, without expecting the relationships between these quantities in $|\Lambda\rangle$ to be exactly reproduced in the data.

We now turn to the Rényi entropy, the quantity which most strikingly distinguishes between the noninteracting and interacting localized phases. To examine this quantity, we revert to partitioning the system in half, instead of into blocks of size ξ. As long as $\xi \ll \frac{L}{2}$, the assumptions that we made above about the blocks of size ξ hold even better for the subsystems A and B. For example, we can assume that there are

"explored sets" of M_A configurations in subsystem A and M_B configurations in subsystem B respectively, with $M = M_A M_B$. We consider computing the reduced density matrix $\hat{\rho}_A$, exactly as in equation (4.20) above. If the off-diagonal elements of this density matrix remain perfectly phase-coherent, it can easily be shown that $S^{\rm coh}_{2,\Lambda} = 0$. In reality, there will be a local contribution to the entropy from particles straying over the cut between subsystems A and B. This mimics the situation in noninteracting localized phases. Alternatively, suppose that dephasing is sufficiently strong that we can proceed by analogy with the ergodic phase, beginning with equation (4.21) and keeping only coherent terms as in equation (4.22). Thereafter, the calculation for the model localized state $|\Lambda\rangle$ differs from the calculation for $|\Phi\rangle$. We need to consider the statistics of the configuration probabilities $|\lambda_{AB}|^2$. For $|\lambda_{AB}|^2 \neq 0$, we need the configurations on both subsystems to lie within their respective explored sets; this occurs in subsystem A, for example, with probability $\frac{M_A}{\gamma(\frac{L}{2},n)}$. This reasoning leads to the "dephased" entropy:

$$\begin{aligned} S^{\rm dp}_{2,\Lambda} &\approx -\log_2\left(\frac{1}{M_A}+\frac{1}{M_B}-\frac{1}{M_A M_B}\right) \\ &\approx -\log_2\left(\frac{2}{\sqrt{M}}-\frac{1}{M}\right) \\ &\approx \frac{1}{2}\left[1-\frac{1}{2\xi}\log_2\left(\frac{\pi e\xi}{2}\right)\right]L-1 \end{aligned} \tag{4.33}$$

where we have additionally made the approximation that typically $M_A \approx M_B \approx \sqrt{M}$. With only partial loss of coherence, the entropy would lie between these two limiting cases: $S^{\rm coh}_{2,\Lambda} \leq S_{2,\Lambda} \leq S^{\rm dp}_{2,\Lambda}$. Thus, dephasing alone, without additional particle transport, can induce an extensive entropy.

Indeed, our numerics support the view that the main difference between the non-interacting and many-body localized phases is the amount of dephasing. There does not seem to be a qualitative difference in particle transport. The particle configuration stays trapped near its initial state, even with interactions, and the system does not thermalize.

4.5 Tracing the Phase Boundary

In this section, we use the data from Section 4.3 to extract estimates of the phase boundary between the localized and ergodic phases. Estimating the location of the MBL transition is extremely challenging. Given the numerically accessible lattice sizes, satisfying finite size scaling analyses are difficult to perform. Nevertheless, rough estimates have been made in the disordered problem [109, 101, 111, 39], and we will now attempt to extract an approximate phase boundary for our model.

We first consider the autocorrelator. Above, we noted the "splaying" feature in the late time plots of the autocorrelator vs. g. The value of g at this feature can be taken as a lower bound for the transition. For g slightly greater than this value, it is possible that χ only decays with L because $\xi > L$ for accessible lattice sizes. Considering two lattice sizes ($L = 16$ and 20) and finding when their values of χ deviate, we find the values reported in the first column of Table 4.2.

Next, we consider the fitting parameter κ in equation (4.12). In Figure 4.5, we see that there is a region where $\kappa < 0$ for finite interaction strength. Because $\eta \leq 1$, finite size effects are clearly dominating the estimate in this region. We can use the

u	χ	κ	m
0.04	0.35	0.45	0.45
0.16	0.30	0.40	0.40
0.32	0.25	0.40	0.40
0.64	0.25	0.40	0.35

Table 4.2: Bounds or estimates of the transition value of g_c at various values of u and based on various measured quantities. The column titled χ gives a lower bound on the transition value of g based on the autocorrellator. The remaining two columns give estimates of g_c based on κ and m, as defined in Sections 4.3.B and 4.3.C respectively. See Section 4.5 for the reasoning behind the estimates. All values carry implicit error bars of ± 0.05, as that is the discretization of our simulated values of g. This error bar should be interpreted, for instance, as the error on our estimate of the location of the maximum value of m. The error on our estimate of g_c is, of course, much larger.

value of g where κ is minimal to track how this region moves as u is varied. This yields the second column of the table.

Finally, a similar approach can be applied to extract estimates of g_c from the fits (4.15). There exists a region where $m > \frac{1}{2}$, but this is mathematically inconsistent in the thermodynamic limit. Therefore, if we find the value of g that maximizes m, we can again estimate the location of the region dominated by finite size effects, yielding the final column of Table 4.2.

The estimates of the transition value g_c in Table 4.2 were obtained using data for the latest time that we simulated (the time bin $t_{\text{bin}} = 9980 \ldots 9999$ for χ and κ and $t = 9999$ for m). However, we have also estimated g_c for data obtained at a half and a quarter of this integration time, finding consistent results. Thus, the general phase structure of the model is invariant to changing the observation time, even though not all measurable quantities have converged to their asymptotic values. Consolidating the information from the estimates in Table 4.2, we propose that the phase diagram qualitatively resembles Figure 4.1.

Before proceeding, it is worth noting that our rough estimates of the phase boundary do not make assumptions regarding the character of the MBL transition (i.e., whether it is continuous or first order). In fact, some of our plots (e.g., panel (c) of Figure 4.7) hint at the possibility of a discontinuous change in S_2 as a function of g in the thermodynamic limit. We are not aware of any results that rule out a first-order MBL transition, and we must keep this possibility in mind.

4.6 Conclusion

Recently, evidence has accumulated that Anderson localization can survive the introduction of sufficiently weak interparticle interactions, giving rise to a many-body localization transition in disordered systems [17, 19, 109, 111, 39]. The MBL transition appears to be a thermalization transition: in the proposed many-body localized phase, the fundamental assumption of statistical mechanics breaks down, and the system fails to serve as its own heat bath [109, 111]. We have presented numerical evidence that this type of transition can also occur in systems lacking true disorder if they instead exhibit "pseudodisorder" in the form of a quasiperiodic potential.

From one perspective, this may be an unsurprising claim. For $g < \frac{1}{2}$ the localized single-particle eigenstates of the quasiperiodic Aubry-André model have the same qualitative structure as those of the Anderson model, and the effects of introducing interactions ought to be similar. By this reasoning, perhaps it is even possible to guess the phase structure of an interacting AA model using knowledge of an interacting

Anderson model: we simply match lines of the two phase diagrams that correspond to the same noninteracting, single-particle localization length.

However, this perspective misses important effects in all regions of the phase diagram. Most obviously, the AA model has a transition at $u = 0$, and it is interesting to see how this transition gets modified as it presumably evolves into the MBL transition at finite u. It is also important to remember that quasiperiodic potentials are completely spatially correlated. This means that the AA model lacks Griffiths effects, and this may have subtle consequences for the dynamics. Finally, the AA model contains a phase that is absent in the one-dimensional Anderson model, the $g > \frac{1}{2}$ extended phase, and we have seen above that interactions have a profound effect upon this regime.

As we saw in Section 1.5 of Chapter 1, understanding MBL in the quasiperiodic context is especially pertinent given the current experimental situation. Some experiments that probe localization physics in cold atom systems use quasiperiodic potentials, constructed from the superposition of incommensurate optical lattices, in place of genuine disorder. The group of Inguscio, in particular, has recently explored particle transport for interacting bosons within this setup [50, 96]. Meanwhile, the AA model has also been realized in photonic waveguides, and experimentalists have studied the effects of weak interactions on light propagation through these systems. They have also begin to investigate "quantum walks" of two interacting photons in disordered waveguides [90, 88]. This protocol resembles the one we have implemented numerically, so similar physics may arise.

Given the current experimental relevance of localization phenomena in quasiperiodic systems, we hope that our study will motivate further attempts to understand these issues. Unfortunately, our ability to definitively identify and analyze the MBL transition is limited by the modest lattice sizes and evolution times that we can simulate. Vosk and Altman recently developed a strong-disorder renormalization group for dynamics in the disordered problem [142], but the reliability of such an approach in the quasiperiodic context is unclear. A time-dependent density matrix renormalization (tDMRG) group study of this problem would be a valuable next step. Tezuka and García-García have published tDMRG results on localization in an interacting AA model, but their focus was not on the thermalization questions of many-body localization [132]. It would be worthwhile to pose these questions using a methodology that allows access to much larger lattices. However, even tDMRG may have difficulty capturing the highly-entangled ergodic phase [72, 15]; thus, an effective numerical approach for definitively characterizing the transition remains elusive.

4.A Exact Diagonalization

This appendix collects exact diagonalization results that supplement the real time dynamics study in the main body of the chapter.

4.A.1 Floquet Analysis of the Modified Dynamics

The goal of the first part of this appendix is to examine the consequences of the modifications to the quantum dynamics described in Section 4.2.B above. We

first verify that the AA transition survives by diagonalizing the single-particle AA Hamiltonian (i.e., the Hamiltonian (4.1) with $u = \frac{V}{h} = 0$) and the single-particle unitary evolution operators (4.5) for various choices of the time step Δt. Subsequently, we employ the same approach to examine how varying Δt impacts the quasienergy spectrum of the interacting, many-body model.

4.A.1.1 Robustness of the Single-Particle Aubry-André Transition

To study the single-particle transition, we focus on the inverse participation ratio:

$$P_{\mathrm{sp}}(g; L) = \left(\sum_{j=0}^{L-1} |\psi_j|^4 \right) \tag{4.34}$$

Here, ψ_j denotes the amplitude of the wave function at site j of an L site lattice. We enclose the sum in equation (4.34) in parentheses to indicate important differences in the averaging procedure with respect to the many-body inverse participation ratio (4.10). In the many-body case, we computed the IPR as a sum over configurations in the quantum state at a particular time in the real time evolution. Then, we averaged over samples, where a sample was specified by a choice of the offset phase to the potential (4.2) and an initial configuration. Throughout this appendix, we instead specify a "sample" solely by the offset phase δ, and we average over *eigenstates* within each sample before averaging over samples.

As noted previously, the usual AA model has a transition that must occur, by duality, at $g_c = \frac{1}{2}$. Near the transition, the localization length is known to diverge with exponent $\nu = 1$ [134]. Our exact diagonalization results indicate that, at the

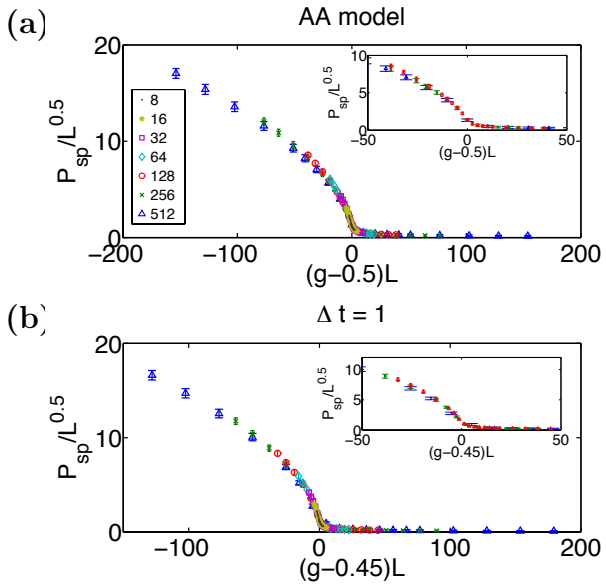

Figure 4.9: Collapse of single-particle IPR vs. g, using the scaling hypothesis (4.35). The legend refers to different lattice sizes L. In panel (a), we show data for the usual AA Hamiltonian (4.1). In panel (b), we show data obtained from diagonalizing the unitary evolution operator for one time step in the modified dynamics (4.5). We use potential wavenumber $k = \frac{1}{\phi}$ and 50 samples for all lattice sizes. The insets show magnified views of the curves for the three largest lattice sizes in the vicinity of the transition.

transition, $P_{\rm sp}(g_c, L) \sim L^{-\frac{1}{2}}$. Hence, we can make the following scaling hypothesis for the IPR:

$$P_{\rm sp} = L^{-\frac{1}{2}} f((g - g_c)L) \tag{4.35}$$

In panel (a) of Figure 4.9, we show that we can use this scaling hypothesis to collapse data for the standard AA model. We show data for $L = 8$ to $L = 512$, with potential wavenumber $k = \frac{1}{\phi}$ and open boundary conditions. For all lattice sizes, we average over 50 samples.

To establish the stability of the AA transition to the modified dynamics, we now must ask if the IPR obtained from diagonalizing the unitary evolution operators (4.5) can be described using the scaling hypothesis (4.35). Panel (b) of Figure 4.9 shows that this is indeed the case for $\Delta t = 1$. The only parameter that needs to be changed is g_c, which decreases slightly as Δt is raised. This implies that there is a transition in the Floquet spectrum of the system that can be tuned by varying Δt. It would be a worthwhile exercise to map out the phase diagram of this single-particle problem in the $(g, \Delta t)$ plane. We leave this for future work.

4.A.1.2 Properties of the Many-Body Quasienergy Spectrum

We now turn our attention back to the effects of the modified dynamics upon the full, many-body model. In Section 4.2.B above, we emphasized that our time-dependent model lacks energy conservation, with multiphoton processes inducing transitions between states of the parent model (4.1) that differ in energy by $\omega_H = \frac{2\pi}{\Delta t}$. In this part of the appendix, we will examine how varying Δt impacts the quasienergy

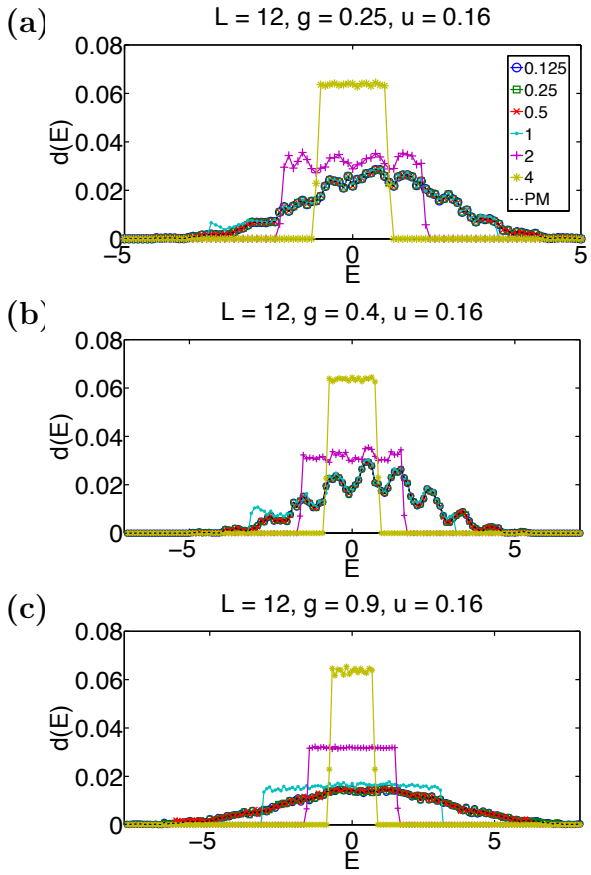

Figure 4.10: The density-of-states vs. quasienergy for $L = 12$ systems at half-filling with interaction strength $u = 0.16$. The legend refers to different values of Δt; the time-independent, parent model is referred to as "PM." In panels (a)-(c), $g = 0.25$, 0.4, and 0.9 respectively.

spectrum of the time-dependent model, using the approach that we applied to the single-particle case above: we diagonalize the time-independent Hamiltonian as well as the unitary evolution operator for one time step of the time-dependent model.

In Figure 4.10, we plot the density-of-states $d(\Delta t, E)$ in quasienergy space of the parent model and time-dependent models for different values of Δt. We focus on $L = 12$ systems at half-filling with fermions (or, because we continue to use the boundary conditions described in Section 4.2.C, hardcore bosons). We fix the interaction strength to $u = 0.16$ and tune g to explore different regimes of the model. In panels (a)-(c), we plot data for $g = 0.25$, 0.4, and 0.9. According to Table 4.2, these values of g put the system in the localized phase, near the transition, and in the ergodic phase respectively.

We first consider the consequences of varying Δt while holding the other parameters fixed. For sufficiently small Δt, the quasienergy spectrum faithfully reproduces all the structure of the energy spectrum of the parent model. This is unsurprising, because if ω_H is greater than the bandwidth of the parent model's spectrum, direct multiphoton processes will not take place. If we now tune ω_H so that it is less than this bandwidth, the quasienergy spectrum begins to deviate from the parent model's spectrum at its edges. This effect can be seen, for instance, by examining the trace for $\Delta t = 1$ in panels (a) or (b). For even higher values of Δt (i.e., lower values of ω_H), multiphoton processes strongly mix the states of the parent model, resulting in a flat quasienergy spectrum.

The effect of multiphoton processes can also be enhanced by broadening the parent

model's spectrum, which can be achieved by raising g or u. In panel (c) of Figure 4.10 for instance, multiphoton processes have significantly flattened the spectrum for $\Delta t = 1$, and deviations from the parent model are even visible for $\Delta t = 0.5$. Because we always use $\Delta t = 1$ in our real time dynamics simulations, it is perhaps fortunate that $g = 0.9$ is well within the proposed ergodic phase for $u = 0.16$ and that, near the critical point (i.e., in panel (b)), the quasienergy spectrum for $\Delta t = 1$ still retains much of the structure of the parent model's spectrum.

However, there is one more caveat to keep in mind: the energy content of the system also grows with L. At fixed g, u, and Δt, the properties of the parent and time-dependent models deviate from one another as the system size grows. If we truly want to faithfully reproduce the dynamics of the parent model with the modified dynamics, it may be necessary to scale Δt down as we raise L. However, recall that our goal is simply to find MBL in a model qualitatively similar to the parent model (4.1). Even with this more modest goal in mind, there is still the danger that, on sufficiently large lattices, multiphoton processes might couple a very large number of localized states and thereby destroy the many-body localized phase of the parent model. Our numerical observations indicate that this does not happen for the system sizes that we can simulate. We can keep Δt fixed at unity for $L \leq 20$ without issues, accepting the possibility that the sequence of models that we would *in principle* simulate on still larger lattices may require progressively smaller values of Δt.

4.A.2 Level Statistics of the Many-Body Parent Model

Localization transitions are often characterized by transitions in the level statistics of the energy spectrum [126]. Two of us previously looked at the level statistics of the disordered problem and identified a crossover from Poisson statistics in the many-body localized phase to Wigner-Dyson statistics in the many-body ergodic phase [109]. The intuition that underlies this crossover is the following: in a localized phase, particle configurations that have similar potential energy are too far apart in configuration space to be efficiently mixed by the kinetic energy term in the Hamiltonian. Therefore, level repulsion is strongly suppressed, and Poisson statistics hold. Conversely, in an ergodic phase, there is strong level repulsion which lifts degeneracies, leading to Wigner-Dyson (i.e., random matrix) statistics.

Along the lines of the aforementioned study of the disordered problem, we focus on the gaps between successive eigenstates of the spectrum of the many-body parent model (4.1):

$$\delta_n \equiv E_{n+1} - E_n \tag{4.36}$$

and a dimensionless parameter that captures the correlations between successive gaps in the spectrum:

$$r_n \equiv \frac{\min(\delta_n, \delta_{n+1})}{\max(\delta_n, \delta_{n+1})} \tag{4.37}$$

For a Poisson spectrum, the r_n are distributed as $\frac{2}{(1+r)^2}$ with mean $2\ln(2) - 1 \approx 0.386$; meanwhile, when random matrix statistics hold, the mean value of r has been numerically determined to be approximately 0.5295 ± 0.0006 [109].

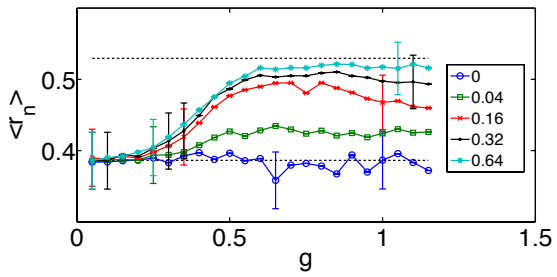

Figure 4.11: The mean of the ratio between adjacent gaps in the spectrum, defined in (4.37). This data was obtained by exact diagonalization of the parent model (4.1) for $L = 12$ systems. All data points have been averaged over 50 samples, and the legend refers to different values of the interaction strength u. The mean value of $\langle r_n \rangle$ shows a crossover from Poisson statistics (indicated by the bottom reference line) to Wigner-Dyson statistics (indicated by the top reference line), for the largest values of u. Representative error bars have been included in the plot; the absent error bars have roughly the same size.

In Figure 4.11, we present exact diagonalization results for $L = 12$ lattices at half-filling with potential wavenumber $k = \frac{1}{\phi}$ and the boundary conditions described in Section 4.2.C above. We show data for the same parameter range examined in the body of this chapter and average over 50 samples for each value of g and u. For the largest value of u, the mean value of r_n interpolates between the expected values as g is raised, consistent with the existence of a localization transition. We have also checked that the distributions of r_n have the expected forms in the small and large g limits in this regime. For smaller values of u, we can speculate that $\langle r_n \rangle$ grows with L at large g and approaches the expected value for very large L. To argue for a MBL transition on the basis of exact diagonalization, we would need to study this sharpening of the crossover as L is raised. This would indeed be an interesting avenue for future work. For our present purposes, however, we only want to check consistency with our real time dynamics data, as we have done in Figure 4.11.

Bibliography

[1] A. G. Abanov, J. C. Talstra, and P. B. Wiegmann. *Physical Review Letters*, 81:2112, 1998.

[2] E. Abrahams, P. Anderson, D. Licciardello, and T. Ramakrishnan. *Physical Review Letters*, 42(10):673, 1979.

[3] D. Akahoshi, M. Uchida, Y. Tomioka, T. Arima, Y. Matsui, and Y. Tokura. *Physical Review Letters*, 90(17):177203, 2003.

[4] M. Albert and P. Leboeuf. *Physical Review A*, 81(1):013614, 2010.

[5] A. Allain, Z. Han, and V. Bouchiat. *Nature Materials*, 11(7):590, 2012.

[6] A. Altland and B. Simons. *Condensed Matter Field Theory*. Cambridge University Press, 2006.

[7] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael. *Physical Review Letters*, 93(15):150402, 2004.

[8] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael. *Physical Review Letters*, 100(17):170402, 2008.

[9] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael. *Physical Review B*, 81(17):174528, 2010.

[10] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. *Science*, 269(5221):198, 1995.

[11] P. Anderson. *Physical Review*, 109(5):1492, 1958.

[12] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Brooks/Cole, 1976.

[13] S. Aubry and G. André. *Ann. Israel Phys. Soc*, 3(133):1, 1980.

[14] K. G. Balabanyan, N. V. Prokofiev, and B. Svistunov. *Physical Review Letters*, 95(5):055701, 2005.

[15] J. H. Bardarson, F. Pollmann, and J. E. Moore. *Physical Review Letters*, 109(1):17202, 2012.

[16] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. *Physical Review*, 108(5):1175, 1957.

[17] D. Basko, I. Aleiner, and B. Altshuler. *Annals of Physics*, 321(5):1126, 2006.

[18] D. Basko, I. Aleiner, and B. Altshuler. *Physical Review B*, 76(5):052203, 2007.

[19] D. Basko, I. Aleiner, and B. Altshuler. *Problems of Condensed Matter Physics*, 1(9):50, 2007.

[20] J. Bellissard, R. Lima, and D. Testard. *Communications in Mathematical Physics*, 88(2):207, 1983.

[21] T. Berkelbach and D. Reichman. *Physical Review B*, 81(22):224429, 2010.

[22] R. Bhatt and P. Lee. *Physical Review Letters*, 48(5):344, 1982.

[23] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, and A. Aspect. *Nature*, 453(7197):891, 2008.

[24] G. Biroli, A. Ribeiro-Teixeira, and M. Tarzia. *ArXiv preprint arXiv:1211.7334*, 2012.

[25] U. Bissbort and W. Hofstetter. *EPL*, 86:50007, 2009.

[26] N. Bogoliubov. *J. Phys. USSR*, 11(23):4, 1947.

[27] N. Bray-Ali, J. Moore, T. Senthil, and A. Vishwanath. *Physical Review B*, 73(6):064417, 2006.

[28] C. Bruder, R. Fazio, and G. Schön. *Annalen der Physik*, 14(9-10):566, 2005.

[29] E. Canovi, D. Rossini, R. Fazio, G. Santoro, and A. Silva. *Physical Review B*, 83(9):094431, 2011.

[30] J. Chaves and I. Satija. *Physical Review B*, 55(21):14076, 1997.

[31] Y. P. Chen, J. Hitchcock, D. Dries, M. Junker, C. Welford, and R. Hulet. *Physical Review A*, 77(3):033632, 2008.

[32] D. N. Christodoulides, F. Lederer, and Y. Silberberg. *Nature*, 424(6950):817, 2003.

[33] L. N. Cooper. *Physical Review*, 104(4):1189, 1956.

[34] R. Crane, N. Armitage, A. Johansson, G. Sambandamurthy, D. Shahar, and G. Grüner. *Physical Review B*, 75(18):184530, 2007.

[35] B. Crooker, B. Hebral, E. Smith, Y. Takano, and J. Reppy. *Physical Review Letters*, 51(8):666, 1983.

[36] E. Dagotto. *New Journal of Physics*, 7(1):67, 2005.

[37] C. Dasgupta and S. Ma. *Physical Review B*, 22(3):1305, 1980.

[38] K. Davis, M.-O. Mewes, M. v. Andrews, N. Van Druten, D. Durfee, D. Kurn, and W. Ketterle. *Physical Review Letters*, 75(22):3969, 1995.

[39] A. De Luca and A. Scardicchio. *EPL*, 101(3):37003, 2013.

[40] R. Desbuquois, L. Chomaz, T. Yefsah, J. Léonard, J. Beugnon, C. Weitenberg, and J. Dalibard. *Nature Physics*, 2012.

[41] J. Deutsch. *Physical Review A*, 43(4):2046, 1991.

[42] V. Dobrosavljevic. *Conductor-Insulator Quantum Phase Transitions*, page 1, 2012.

[43] F. J. Dyson. *Physical Review*, 92(6):1331, 1953.

[44] L. DAlessio and A. Polkovnikov. *Annals of Physics*, 333, 2013.

[45] J. Edwards and D. Thouless. *Journal of Physics C: Solid State Physics*, 5(8):807, 1972.

[46] T. Eggarter and R. Riedinger. *Physical Review B*, 18(2):569, 1978.

[47] A. Eilmes, U. Grimm, R. Roemer, and M. Schreiber. *The European Physical Journal B-Condensed Matter and Complex Systems*, 8(4):547, 1999.

[48] A. Einstein. *Quantentheorie des Einatomigen Idealen Gases*. Akademie der Wissenshaften, in Kommission bei W. de Gruyter, 1924.

[49] S. N. Evangelou and D. E. Katsanos. *Physical Review B*, 56:12797, 1997.

[50] L. Fallani, J. Lye, V. Guarrera, C. Fort, and M. Inguscio. *Physical Review Letters*, 98(13):130404, 2007.

[51] R. Fernandes and J. Schmalian. *Physical Review Letters*, 106(6):67004, 2011.

[52] D. Fisher. *Physical Review B*, 50(6):3799, 1994.

[53] D. Fisher. *Physica A: Statistical Mechanics and its Applications*, 263(1-4):222, 1999.

[54] D. S. Fisher. *Physical Review B*, 51(10):6411, 1995.

[55] M. Fisher, P. Weichman, G. Grinstein, and D. Fisher. *Physical Review B*, 40(1):546, 1989.

[56] A. Frydman, O. Naaman, and R. Dynes. *Physical Review B*, 66(5):052509, 2002.

[57] T. Giamarchi, P. Le Doussal, and E. Orignac. *Physical Review B*, 64(24):245119, 2001.

[58] T. Giamarchi and H. Schulz. *Physical Review B*, 37(1):325, 1988.

[59] A. Gottlob and M. Hasenbusch. *Physica A: Statistical Mechanics and its Applications*, 201(4):593, 1993.

[60] M. Greiner, O. Mandel, T. Esslinger, T. Hänsch, and I. Bloch. *Nature*, 415(6867):39, 2002.

[61] R. B. Griffiths. *Physical Review Letters*, 23(1):17, 1969.

[62] T. Guénault. *Basic Superfluids*. CRC Press, 2002.

[63] V. Gurarie, L. Pollet, N. Prokofiev, B. Svistunov, and M. Troyer. *Physical Review B*, 80(21):214519, 2009.

[64] F. Haldane. *Physical Review Letters*, 47(25):1840, 1981.

[65] P. Harper. *Proceedings of the Physical Society. Section A*, 68:874, 1955.

[66] A. Harris. *Journal of Physics C: Solid State Physics*, 7:1671, 1974.

[67] Y. Hashimoto, K. Niizeki, and Y. Okabe. *Journal of Physics A: Mathematical and General*, 25(20):5211, 1992.

[68] K. He, I. I. Satija, C. W. Clark, A. M. Rey, and M. Rigol. *Physical Review A*, 85(1):013617, 2012.

[69] S. Hikami, A. I. Larkin, and Y. Nagaoka. *Progress of Theoretical Physics*, 63(2):707, 1980.

[70] F. Hrahsheh and T. Vojta. *Physical Review Letters*, 109(26):265303, 2012.

[71] D. Huse. Private communication. 2011.

[72] M. Žnidarič, T. c. v. Prosen, and P. Prelovšek. *Physical Review B*, 77:064426, 2008.

[73] S. Iyer, V. Oganesyan, G. Refael, and D. A. Huse. *Physical Review B*, 87:134202, 2013.

[74] S. Iyer, D. Pekker, and G. Refael. *In preparation.*

[75] S. Iyer, D. Pekker, and G. Refael. *Physical Review B*, 85(9):094202, 2012.

[76] S. Jitomirskaya. *Annals of Mathematics-Second Series*, 150(3):1159, 1999.

[77] S. Jitomirskaya and B. Simon. *Communications in Mathematical Physics*, 165(1):201, 1994.

[78] A. Karahalios, A. Metavitsiadis, X. Zotos, A. Gorczyca, and P. Prelovšek. *Physical Review B*, 79:024425, 2009.

[79] E. Khatami, M. Rigol, A. Relaño, and A. Garcia-Garcia. *Physical Review E*, 85(5):050102, 2012.

[80] T. Kitagawa, T. Oka, and E. Demler. *Annals of Physics*, 327, 2012.

[81] A. R. Kolovsky and G. Mantica. *Physical Review B*, 86(5):054306, 2012.

[82] S. Kondov, W. McGehee, J. Zirbel, and B. DeMarco. *Science*, 334(6052):66, 2011.

[83] J. M. Kosterlitz and D. J. Thouless. *Journal of Physics C: Solid State Physics*, 6(7):1181, 1973.

[84] I. Kovács and F. Iglói. *Physical Review B*, 82(5):054437, 2010.

[85] Y. E. Kraus, Y. Lahini, Z. Ringel, M. Verbin, and O. Zilberberg. *Physical Review Letters*, 109(10):106402, 2012.

[86] F. Krüger, S. Hong, and P. Phillips. *Physical Review B*, 84(11):115118, 2011.

[87] Y. Lahini, A. Avidan, F. Pozzi, M. Sorel, R. Morandotti, D. N. Christodoulides, and Y. Silberberg. *Physical Review Letters*, 100(1):013906, 2008.

[88] Y. Lahini, Y. Bromberg, D. Christodoulides, and Y. Silberberg. *Physical Review Letters*, 105(16):163905, 2010.

[89] Y. Lahini, Y. Bromberg, Y. Shechtman, A. Szameit, D. Christodoulides, R. Morandotti, and Y. Silberberg. *Physical Review A*, 84(4):041806, 2011.

[90] Y. Lahini, R. Pugatch, F. Pozzi, M. Sorel, R. Morandotti, N. Davidson, and Y. Silberberg. *Physical Review Letters*, 103(1):13901, 2009.

[91] A. Leggett. *Physical Review Letters*, 29(18):1227, 1972.

[92] D. Licciardello and D. Thouless. *Journal of Physics C: Solid State Physics*, 8(24):4157, 1975.

[93] E. H. Lieb and F. Wu. *Physical Review Letters*, 20(25):1445, 1968.

[94] E. Loh, J. Gubernatis, R. Scalettar, S. White, D. Scalapino, and R. Sugar. *Physical Review B*, 41(13):9301, 1990.

[95] F. London. *Nature*, 141(3571):643, 1938.

[96] E. Lucioni, B. Deissler, L. Tanzi, G. Roati, M. Zaccanti, M. Modugno, M. Larcher, F. Dalfovo, M. Inguscio, and G. Modugno. *Physical Review Letters*, 106(23):230403, 2011.

[97] J. Luttinger. *Journal of Mathematical Physics*, 4:1154, 1963.

[98] J. Lye, L. Fallani, M. Modugno, D. Wiersma, C. Fort, and M. Inguscio. *Physical Review Letters*, 95(7):070401, 2005.

[99] M. Maricq. *Physical Review B*, 25:6622, 1982.

[100] D. Martinez and R. Molina. *The European Physical Journal B-Condensed Matter and Complex Systems*, 52(2):281, 2006.

[101] C. Monthus and T. Garel. *Physical Review B*, 81(13):134202, 2010.

[102] J. Moore. Private communication. 2011.

[103] O. Motrunich, K. Damle, and D. A. Huse. *Physical Review B*, 65(6):064206, 2002.

[104] O. Motrunich, S. Mau, D. Huse, and D. Fisher. *Physical Review B*, 61(2):1160, 2000.

[105] N. Mott. *Proceedings of the Physical Society. Section A*, 62:416, 1949.

[106] N. Mott. *Philosophical Magazine*, 19(160):835, 1969.

[107] N. Mott and W. Twose. *Advances in Physics*, 10(38):107, 1961.

[108] N. F. Mott. *Metal-insulator Transitions*. Taylor & Francis London, 1990.

[109] V. Oganesyan and D. Huse. *Physical Review B*, 75(15):155111, 2007.

[110] V. Oganesyan, A. Pal, and D. A. Huse. *Physical Review B*, 80:115104, 2009.

[111] A. Pal and D. Huse. *Physical Review B*, 82(17):174411, 2010.

[112] M. Pasienski, D. McKay, M. White, and B. DeMarco. *Nature Physics*, 2010.

[113] L. Pollet, N. Prokofiev, B. Svistunov, and M. Troyer. *Physical Review Letters*, 103(14):140402, 2009.

[114] N. Prokofiev and B. Svistunov. *Physical Review Letters*, 92(1):15703, 2004.

[115] J. Reppy. *Physica B+C*, 126(1-3):335, 1984.

[116] P. Ribeiro, M. Haque, and A. Lazarides. *ArXiv preprint arXiv:1211.6012*, 2012.

[117] M. Rigol, V. Dunjko, and M. Olshanii. *Nature*, 452(7189):854, 2008.

[118] G. Roati, C. DErrico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, and M. Inguscio. *Nature*, 453(7197):895, 2008.

[119] T. Roscilde and S. Haas. *Physical Review Letters*, 99(4):47205, 2007.

[120] S. Sachdev. *Quantum Phase Transitions*. Wiley Online Library, 2007.

[121] A. Sandvik. *Physical Review Letters*, 96(20):207201, 2006.

[122] T. Schulte, S. Drenkelforth, J. Kruse, W. Ertmer, J. Arlt, K. Sacha, J. Zakrzewski, and M. Lewenstein. *Physical Review Letters*, 95(17):170411, 2005.

[123] T. Schwartz, G. Bartal, S. Fishman, and M. Segev. *Nature*, 446(7131):52, 2007.

[124] P. Sengupta and S. Haas. *Physical Review Letters*, 99(5):50403, 2007.

[125] D. Sherman, G. Kopnov, D. Shahar, and A. Frydman. *Physical Review Letters*, 108(17):177006, 2012.

[126] B. Shklovskii, B. Shapiro, B. Sears, P. Lambrianides, and H. Shore. *Physical Review B*, 47(17):11487, 1993.

[127] A. Siebesma and L. Pietronero. *EPL (Europhysics Letters)*, 4(5):597, 1987.

[128] B. Simon. *Advances in Applied Mathematics*, 3(4):463, 1982.

[129] M. Srednicki. *Physical Review E*, 50(2):888, 1994.

[130] D. Stauffer and A. Aharony. *Introduction to Percolation Theory*. Taylor & Francis, London, 1994.

[131] N. Teichmann, D. Hinrichs, M. Holthaus, and A. Eckardt. *Physical Review B*, 79(10):100503, 2009.

[132] M. Tezuka and A. M. García-García. *Physical Review A*, 85:031602, 2012.

[133] G. Theodorou and M. H. Cohen. *Physical Review B*, 13(10):4597, 1976.

[134] D. Thouless and Q. Niu. *Journal of Physics A: Mathematical and General*, 16:1911, 1983.

[135] D. J. Thouless. *Physics Reports*, 13(3):93, 1974.

[136] S.-I. Tomonaga. *Progress of Theoretical Physics*, 5(4):544, 1950.

[137] M. Verbin, O. Zilberberg, Y. E. Kraus, Y. Lahini, and Y. Silberberg. *Physical Review Letters*, 110(7):076403, 2013.

[138] J. Vidal, D. Mouhanna, and T. Giamarchi. *Physical Review Letters*, 83(19):3908, 1999.

[139] J. Voit. *arXiv preprint cond-mat/0005114*, 2000.

[140] T. Vojta. *Journal of Low Temperature Physics*, 161:299, 2010.

[141] R. Vosk and E. Altman. *Physical Review B*, 85(2):024531, 2012.

[142] R. Vosk and E. Altman. *Physical Review Letters*, 110(6):067204, 2013.

[143] L. Wang and A. Sandvik. *Physical Review Letters*, 97(11):117204, 2006.

[144] L. Wang and A. Sandvik. *Physical Review B*, 81(5):054417, 2010.

[145] Y. Wang, W. Guo, and A. Sandvik. *Arxiv preprint arXiv:1110.3213*, 2011.

[146] P. Weichman. *Arxiv Preprint arXiv:0810.3263*, 2008.

[147] M. White, M. Pasienski, D. McKay, S. Zhou, D. Ceperley, and B. DeMarco. *Physical Review Letters*, 102(5):55301, 2009.

[148] R. Yu, L. Yin, N. S. Sullivan, J. Xia, C. Huan, A. Paduan-Filho, N. F. Oliveira Jr, S. Haas, A. Steppke, C. F. Miclea, et al. *Nature*, 489(7416):379, 2012.

[149] Y. Zou. *Quantum Phases and Phase Transitions in Disordered Low-Dimensional Systems*. PhD thesis, California Institute of Technology, 2010.

[150] Y. Zou, R. Barnett, and G. Refael. *Physical Review B*, 82(22):224205, 2010.