THE EFFECTS OF DISORDER IN STRONGLY INTERACTING QUANTUM SYSTEMS

Steven J. Thomson

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

2016

The Effects of Disorder in Strongly Interacting Quantum Systems

Steven J. Thomson

University of
St Andrews

This thesis is submitted in partial fulfilment for the degree
of PhD at the University of St Andrews

24 August 2016

1. Candidate's declarations:

I, Steven Thomson, hereby certify that this thesis, which is approximately 43,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in August 2012 and as a candidate for the degree of PhD in August 2012; the higher study for which this is a record was carried out in the University of St Andrews between 2012 and 2016.

Date ……………..   Signature of candidate ………………………………………………

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of ……… in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date ……………..   Signature of supervisor ………………………………………………

3. Permission for publication: (*to be signed by both candidate and supervisor*)

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY
a)   No embargo on print copy

ELECTRONIC COPY
a)   No embargo on electronic copy
b)

Date ………………..   Signature of candidate …………………………………..

                                    Signature of supervisor …………………………………

For my parents, whose unwavering support and encouragement has enabled me to get this far. I know you probably won't understand a word of this thesis, but I hope it makes you proud regardless.

Abstract

This thesis contains four studies of the effects of disorder and randomness on strongly correlated quantum phases of matter. Starting with an itinerant ferromagnet, I first use an order-by-disorder approach to show that adding quenched charged disorder to the model generates new quantum fluctuations in the vicinity of the quantum critical point which lead to the formation of a novel magnetic phase known as a helical glass.

Switching to bosons, I then employ a momentum-shell renormalisation group analysis of disordered lattice gases of bosons where I show that disorder breaks ergodicity in a non-trivial way, leading to unexpected glassy freezing effects. This work was carried out in the context of ultracold atomic gases, however the same physics can be realised in dimerised quantum antiferromagnets. By mapping the antiferromagnetic model onto a hard-core lattice gas of bosons, I go on to show the importance of the non-ergodic effects to the thermodynamics of the model and find evidence for an unusual glassy phase known as a Mott glass not previously thought to exist in this model.

Finally, I use a mean-field numerical approach to simulate current generation quantum gas microscopes and demonstrate the feasibility of a novel measurement scheme designed to measure the Edwards-Anderson order parameter, a quantity which describes the degree of ergodicity breaking and which has never before been experimentally measured in any strongly correlated quantum system.

Together, these works show that the addition of disorder into strongly interacting quantum systems can lead to qualitatively new behaviour, triggering the formation of new phases and new physics, rather than simply leading to small quantitative changes to the physics of the clean system. They provide new insights into the underlying physics of the models and make direct connection with experimental systems which can be used to test the results presented here.

ii

Acknowledgements

First and foremost, I'd like to thank my supervisors, **Frank Krüger** and **Chris Hooley** for taking me on as a PhD student and for all their advice and support over the last four years. I'd also like to thank the Condensed Matter Doctoral Training Centre and in particular acknowledge the eternal debt that all its students owe to **Christine Edwards**, **Julie Massey** and **Wendy Clark** for everything they've done on behalf of all of the students and staff members in the Centre.

I've had the tremendous good fortune to share an office with some excellent officemates and friends. I'd like to thank **Calum Lithgow** for teaching me more about politics and grammar than I ever wanted to know, but also for educating me about whisky, and **Helen Cammack** for never letting me forget that there can be more to life than equations. I endeavour to one day figure out just what she's talking about. I'd also like to thank **Scott Pearson** for his hard-won advice and wisdom that's been so much help to me over the years.

All of the staff and students of the School of Physics & Astronomy deserve more thanks than I can possibily articulate. In particular, thanks to **Jonathan Keeling** for his help and advice, to **Graham Bruce** for trusting my crazy ideas and to **Liam Walker** for choosing to work with us for his Masters project, as well as to our collaborator **Tiffany Harte** for her invaluable contributions to that work. During the last few years, I've also benefited greatly from opportunities to talk with many others from outwith St Andrews and in particular I'd like to thank **Andrew Green** and **Chris Pedder** for useful conversations and advice. I particularly wish to thank my examiners, **Brendon Lovett** and **Joe Bhaseen** for investing their time in reading my thesis and offering useful comments, as well as for making the viva one of the definite highlights of my PhD.

I can't in good conscience not take this opportunity to acknowledge in writing the debt I owe to my school teachers. Specifically, I want to thank my physics teachers **Stuart Dudley** and **John Sneddon**, who taught me to trust myself and question everything; my technological studies teacher **Gary Plimer** who taught me how to think like a problem-solver, and my chemistry teacher **Lesley Shaw** who taught me not to give up when things get difficult. I'd like to specially thank **David Parker**, the finest maths teacher I could ever imagine, for helping me discover the beauty of mathematics.

And as I have learned my teachers, so too have I learned from my students. To all of you whom I've taught over the years - I suspect I've learned at least as much from you as you've learned from me and I'd like to thank you all for being the most fantastic bunch of students I could have hoped for. I hope my tutoring hasn't left you with any lasting trauma.

To all my friends outside the School, both in St Andrews and elsewhere, without whom I'd have never managed to emerge from the PhD as anything resembling a functioning human being. To the staff of the local bottle shops and bars, without whom I may not have coped. To whoever invented coffee. To my family for believing in me. To the ghosts and the poets, to the lost and the free.

Were I not the type to use five hundred and eighty five words where nineteen would suffice, I could have summed up all of the above in a single sentence:

To all those who have helped me out over the years: you know who you are, and thank you.

"It's the questions we can't answer that teach us the most. They teach us how to think...The harder the question, the harder we hunt. The harder we hunt, the more we learn. An impossible question..."

The Wise Man's Fear, Patrick Rothfuss

Contents

Chapter 1

Introduction

Nothing in life is perfect and what's true in life is true for quantum materials. Randomness and disorder are present to some degree in every substance in the world, despite every effort to create the cleanest, purest materials possible. Even small concentrations of impurities in quantum materials can dramatically alter their properties, for better or for worse.

Since the time of the Crusades, blacksmiths have known that introducing controlled carbon impurities into swords could vastly improve the properties of the blade, leading to the development of the famed Damascus steel. Stronger than any contemporary materials, Damascus steel was later found to contain carbon nanotubes [1], making it arguably the first commercial application of nanotechnology. We now know that doping materials can lead to the emergence of dramatic properties such as superconductivity [2]. Understanding the effects of disorder is crucial to the continued development of materials and technology but is also important for our understanding of the underlying physics in strongly interacting quantum systems. Disorder and impurities often act as excellent probes of quantum behaviour, changing the properties of the system in question in measurable ways and giving us valuable information as to the underlying physics.

In this chapter, we will review the theory of phase transitions in clean systems (classical and quantum) and discuss the problem of disorder. Stage thus set, in the later chapters we will examine my original contributions to the theory of phase transitions in disordered systems.

1.1 Classical Phase Transitions

A phase transition is the process by which a thermodynamically large number of interacting particles collectively change their state of matter, e.g. from liquid to solid, liquid to gas, magnetic to non-magnetic or a vast number of other possibilities.

Classical phase transitions occur at finite temperature, where the transformation between phases is driven by thermal fluctuations. Consequently, they are also known as 'thermal transitions'. At a classical phase transition, the dynamic and static behaviour of the system decouple from each other and the transition can be described by a time-independent theory.

Many phase transitions are accompanied by a breaking of some fundamental symmetry of the system and a corresponding change in how ordered the material is. Transitions are driven by the competition between energy and entropy as the system tries to minimise its Helmholtz free energy $\mathcal{F} = E - TS$, where E is the internal energy of the system, T is temperature and S is entropy. At low temperatures, the entropic term has a negligible effect and so the system adopts a configuration that minimises its internal energy, which tend to be highly ordered states with broken symmetries. At higher temperatures, the system will transition into a state where it instead maximises its entropy S, which tend to be more isotropic and symmetrical configurations. For example, in the case of liquid water, as the temperature is lowered the isotropic, symmetric liquid (large entropy) freezes into a highly ordered (low entropy) sixfold-symmetric crystal lattice where the rotational symmetry of the liquid is broken in favour of the hexagonal symmetry of ice[1].

This picture is largely generic, though there are some phase transitions where the nature of the symmetry breaking present (if any) remains ambiguous [3–5] and others which don't break any symmetries at all, such as liquid water into water vapour or the Kosterlitz-Thouless transition [6] seen in the two dimensional XY model.

Phase transitions typically fall into two main classes: the discontinuous first-order transitions and the continuous second-order transitions. First order transitions carry with them discontinuities in entropy and a latent heat, whereas second order transitions do not. In general, an n^{th} order transition is characterised by a discontinuity in the n^{th} derivative of the free energy, however this is not considered a rigorous definition [7] as some derivatives of the free energy diverge on approach to a phase transition, e.g. specific heat in ferromagnets.

The nature of a transition is not immutable and can change for different values of external parameters such as temperature, pressure, applied magnetic

---

[1]There are at least sixteen other distinct types of ice crystal existing under a variety of conditions which display different forms of ordering, but the hexagonal crystalline ice 1_h is the most commonly encountered variety here on Earth.
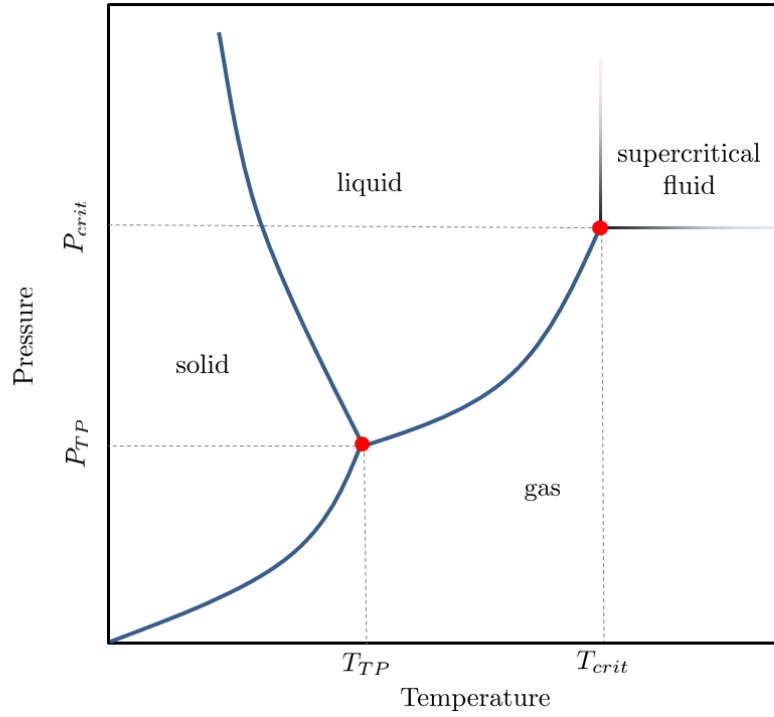
Figure 1.1 *A schematic phase diagram of water in the temperature-pressure plane showing the approximate regions where each phase occurs. The blue lines represent phase transitions. The lower red dot is the triple point at which all three phases can co-exist. The upper red dot is the critical point where the first-order liquid to gas transition terminates and becomes second order. Above this point there is no distinct liquid-to-gas transition and instead a so-called supercritical fluid exists in this parameter regime.*

field etc. For example, the water-to-vapour transition is first-order up until a critical point, at which the transition becomes second order (shown schematically in Fig. 1.1). Above this critical point, the distinction between liquid and gas largely disappears and the phase is known as a 'supercritical fluid'. In this region there is no clear change seen in any parameter that can define whether the system is in a liquid phase or a gaseous phase. Below the critical point, however, crossing the transition from liquid-to-gas results in a discontinuous jump in density. The density is an example of an order parameter, a quantity that changes when the material crosses from one phase to another, and they play an important role in the theory of phase transitions.

1.2 Landau Theory

Almost every phase transition has an order parameter that defines where the transition takes place. This is some quantity that has a thermodynamic average of zero in the disordered phase and takes on some finite value in the ordered phase. Near a second-order or weakly first-order phase transition, we can assume that the order parameter is small.

The key idea behind Landau theory is that the free energy near a transition can be Taylor expanded in powers of the order parameter. Implicitly, this assumes the free energy to be analytic in the order parameter. For this description to be physical, the free energy must obey the same symmetries as the Hamiltonian. In practice, this typically leads to only even powers of the order parameter featuring in the expansion. More generally, any term in the expansion that explicitly breaks any symmetry in the initial Hamiltonian will be absent.

As an illustrative example, let us consider expanding some generic free energy $\mathcal{F}$ in even powers of some order parameter m like so:

$$\mathcal{F} = \alpha m^2 + \beta m^4 + \gamma m^6 + ... \tag{1.1}$$

where the coefficients α, β and γ are all functions of the microscopic parameters of the system in question. The coefficient of the final term in the expansion must always be positive or else the free energy becomes unbounded from below, which is unphysical.

If α, β and γ are all positive then the free energy is minimised by $m = 0$. If $\alpha < 0$ and $\beta > 0$ then γ is not needed and we see that the free energy is minimised by a non-zero $m = \pm\sqrt{|\alpha|/2\beta}$. The transition is given by $\alpha = 0$, i.e. the line along which a non-zero m becomes favoured. As α turns from positive to negative, $\mathcal{F}$ develops two minima which emerge smoothly from $m = 0$, shown in Fig. 1.2. The ground state magnetisation continuously increases from zero and this is a second-order phase transition.

This formalism can also describe first-order phase transitions. Consider the case of $\alpha > 0$, $\gamma > 0$ but $\beta < 0$, also shown in Fig. 1.2. We see that the free energy now acquires two local minima at some $m \neq 0$. If β is sufficiently negative (specifically, if $\beta < 0$ and $\beta^2 \geq 4\alpha\gamma$) then these minima will drop below the $m = 0$ minimum, meaning that the ground state abruptly jumps from $m = 0$ to some $m \neq 0$. This is therefore a first-order transition.
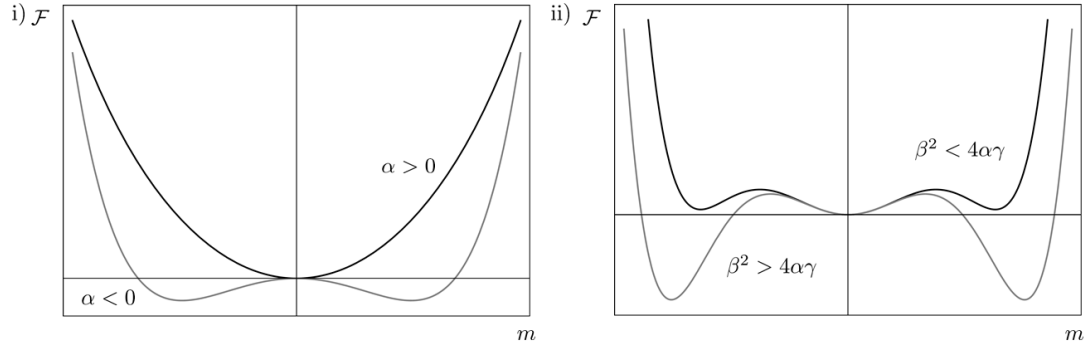
Figure 1.2 *Schematic free energy diagrams. i) A second-order transition as α goes from positive to negative, with minima located at $m \neq 0$ smoothly moving out from the centre. ii) A first-order transition where β becomes sufficiently negative that the two local minima shown in black dip below the $m = 0$ minimum and the magnetisation abruptly jumps from $m = 0$ to some non-zero m.*

1.3   Correlation Length

Approaching the critical point at which a second-order phase transition occurs, fluctuations of the order parameter on larger and larger length scales become increasingly important. These are known as critical fluctuations. In the vicinity of a critical point there is only one physically relevant length scale, known as the correlation length, which is the largest length scale on which the order parameter fluctuations occur and takes the following form:

$$\xi \propto \left| \frac{T - T_c}{T_c} \right|^{-\nu} , \tag{1.2}$$

where T_c is the critical temperature at which the phase transition occurs and ν is known as the correlation length exponent.

Close to the phase transition, the correlation length diverges. It becomes the largest length scale in the system and consequently the only relevant one for determining the properties of the system. The system effectively self-averages over large volumes, in other words, and small microscopic fluctuations become unimportant as only large-scale properties such as dimensionality and lattice geometry affect the overall behaviour.

This leads to the so-called universal behaviour of second-order phase transitions. In all but a few cases, the nature of the transition is independent of the microscopic details, since the diverging correlation length essentially overrules all other properties of the system. This is largely why phenomenological theories of second-order phase transitions have enjoyed so much success: since no length scale other than ξ matters, our effective theories don't need to take into account the

detailed microscopics of the system as they can only cause quantitative changes.

The diverging correlation length close to a critical point is also responsible for the phenomenon of critical opalescence, where a normally transparent liquid will turn opaque when the correlation length becomes comparable to the wavelength of light, leading to the scattering of light and causing the liquid to take on a cloudy white appearance. This is a dramatic and very visual demonstration of the diverging correlation length that occurs near a classical critical point.

1.4 Quantum Phase Transitions

Classical phase transitions are driven solely by thermal fluctuations. If we cool down our system to $T = 0$, then, does that mean we end up in a fluctuationless ground state that exhibits no phase transitions? Classical theory alone would have us believe so, but quantum theory tells us that such a fluctuationless state is impossible. Driven by the Heisenberg uncertainty principle, even a system in its ground state at $T = 0$ will exhibit quantum mechanical fluctuations that, in combination with the tuning of some non-thermal physical parameter, can lead to so-called quantum phase transitions. For $T > 0$, even if the phases themselves require a quantum mechanical description (such as superfluidity or superconductivity), the transition into and out of the phase is still a classical phase transition in which quantum mechanics plays no role.

Quantum phase transitions [8] differ from the classical phase transitions described earlier in that the kinetic and potential parts of the Hamiltonian do not commute, meaning the partition function does not factorise. Static and dynamic behaviour no longer decouple and must be treated on an equal footing, leading to not only a divergent correlation length in space, but also a divergent correlation time.

At zero temperature, the correlation length can be expressed in a similar form as Eq. 1.2 but with some non-thermal quantity g taking the place of temperature:

$$\xi \propto \left| \frac{g - g_c}{g_c} \right|^{-\nu},$$ (1.3)

where g is some parameter such as pressure or applied magnetic field and g_c is the critical value of this parameter at which the phase transition occurs. The correlation time is given by:

$$t_C \propto \left| \frac{g - g_c}{g_c} \right|^{-d_z \nu}.$$ (1.4)

where d_z is the dynamical critical exponent. It is also commonly denoted z, however in this thesis we shall use d_z for the dynamical critical exponent and define z as the lattice co-ordination number.

6

The necessity of treating space and time on an equal footing introduces additional dimensions to what was previously a time-independent problem and means that quantum phase transitions in d dimensions may often be mapped onto classical phase transitions in $d + d_z$ dimensions.

If $d_z = 1$ then the relevant correlation functions scale the same in each spatial dimension as they do in time. This is often not true for quantum critical points, and for $d_z \neq 1$ we end up with a theory where space and time scale differently, leading to quantum critical points exhibiting much richer behaviour than simply being a classical critical point in 'more dimensions'.

A divergent correlation time is not the only extra consideration we need make in describing a quantum phase transition. The real distinguishing feature of a quantum phase transition is that it occurs at zero temperature, but it turns out that we can incorporate both real time and temperature together in a single move.

Imaginary time

In classical statistical mechanics, Boltzmann weightings are given by terms of the form $\mathrm{e}^{-\beta \hat{\mathcal{H}}}$ where $\hat{\mathcal{H}}$ is the Hamiltonian and $\beta = 1/k_B T$ is the inverse temperature. In quantum mechanics, evolution in real time is given by $\mathrm{e}^{-i\hat{\mathcal{H}}t/\hbar}$. In order to perform quantum statistical mechanics, we can make a Wick rotation from real time to imaginary time $t \to -i\tau$ where $\tau \in [0, \hbar\beta]$ such that we can describe these thermal weightings in terms of imaginary time. This allows us to construct partition functions as path integrals over an imaginary time dimension of length $\hbar\beta$. The mapping takes the form:

$$\mathcal{Z} = \mathrm{Tr}\left[\mathrm{e}^{-\beta\mathcal{H}}\right] \tag{1.5}$$

$$= \int \mathcal{D}[\phi^*, \phi] \exp\left[\int_0^{\hbar\beta} \mathrm{d}\tau \, (\phi^* \partial_\tau \phi - \mathcal{H})\right], \tag{1.6}$$

where the rewriting in terms of a path integral over bosonic coherent states ϕ is as shown in Ref. [9]. This representation allows us to treat both real-time dynamics and quantum statistics at the same time. At zero temperature, the length of the imaginary time dimension is infinite and its contribution can be treated on an equal footing with the spatial dimensions. At non-zero temperatures, the length of the imaginary time dimension is finite. The quantum-to-classical mapping can then lead to classical systems with odd geometries, being infinite in extent in d dimensions and finite in d_z.

To obtain any information on real-time dynamics, one must perform an analytic continuation from imaginary time to real time, a process which is not well-controlled and can often lead to unphysical results. For example, the analytic continuation can obscure long-range correlations in imaginary time [10]. The mapping to the classical system can also generate complex Boltzmann weights known as Berry phases [11] which have no classical analogue and lead to new types

of critical phenomena not seen in classical systems. For fermions, the mapping can even generate negative Boltzmann weights - this is known as the fermion sign problem, and is one of the major problems facing modern condensed matter physics. (It is in the complexity class NP-hard [12], which is featured in the Clay Mathematics Institute's Millennium Prize Problems with a \$1 million prize for a solution [13].)

Given that the length of the imaginary time dimension is infinite only at $T = 0$ and that the quantum phase transition itself only occurs at zero temperature, one might wonder whether the imaginary time dependence can be neglected whenever $\hbar\beta$ is finite and a purely classical description used for all but the extreme $T = 0$ case. This turns out not to be the case: although the quantum phase transition itself is restricted to $T = 0$, there is a large region at finite temperatures where quantum fluctuations originating at zero-temperature quantum critical points turn out to be extremely important.

1.5   Quantum Criticality

Despite quantum phase transitions strictly only taking place at $T = 0$, their effects may still be felt at finite[2] temperatures. From the correlation time t_c it follows that we can define a corresponding quantum fluctuation frequency ω. Consider a system at temperature $k_B T$ with some generic quantum fluctuations of strength $\hbar\omega$. When $\hbar\omega \gg k_B T$ then the system will undergo a quantum phase transition. Conversely, for $k_B T \gg \hbar\omega$, the system will undergo a classical phase transition. This criterion will always be satisfied sufficiently close to the transition for any $T > 0$ and the transition will always be classical in nature, even if the phases themselves require a quantum mechanical description.

There is, however, a region of the phase diagram where the quantum and thermal fluctuations are approximately equal in size. Plotting the boundaries defined by $\hbar\omega \approx k_B T$ defines the so-called quantum critical region: a roughly conical region extending out from the $T = 0$ quantum critical point. A schematic phase diagram is shown in Fig. 1.3. Within this region, the effects of quantum fluctuations are strongly felt and the behaviour of the system will not be classical. For example, Fermi liquid theory [9] will not apply in this region even for an 'ideal' metal due to the presence of these strong quantum fluctuations [14].

These quantum fluctuations can lead to all sorts of interesting behaviour. Experimentally, second-order phase transitions are commonly seen to turn first-order in the vicinity of QCPs in a variety of materials [15–19], and often entirely new phases are found. Examples of this include the superconductivity seen in UGe_2 [20–23], the modulated nematic phase seen in $Sr_3Ru_2O_7$ [24–27], the partially-ordered phase seen in MnSi [28, 29] and the modulated spiral seen in

---

[2]Following Ref. [8], I engage in the same 'almost standard abuse of language' and refer to non-zero temperatures as finite.
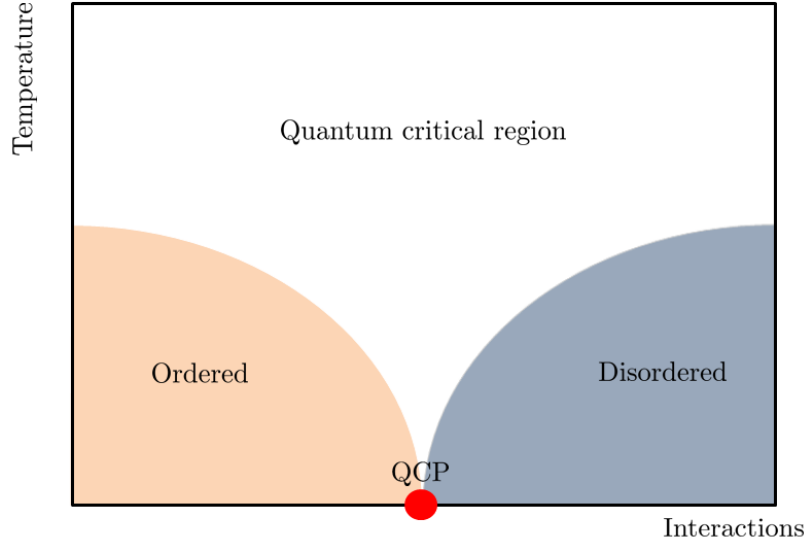
Figure 1.3 *Schematic phase diagram showing quantum ordered, quantum disordered and quantum critical regions. The quantum critical point (QCP) is indicated in red. Above it, the roughly-conical region where critical fluctuations strongly influence the behaviour of the material extends up to non-zero temperatures.*

PrPtAl [30], to name but a few.

It is this fluctuation-dominated region that makes the study of quantum phase transitions both useful and interesting. The extended effects of quantum fluctuations beyond the strictly zero-temperature limit means that understanding the role of quantum fluctuations at zero temperature can still tell us useful information about the finite-temperature response of the material.

The hope is that by developing our understanding of quantum criticality we will learn to control these fluctuations and eventually learn to design new materials capable of exhibiting macroscopic quantum phenomena all the way up to room temperature, stabilising phases such as high-temperature superconductivity. Some coupled light-matter systems already exhibit macroscopic quantum behaviour at high temperatures, such as polariton condensates where the low quasiparticle mass leads to an abnormally high condensation temperature as compared with solid-state systems [31], however these are the exception rather than the rule. Control of the quantum mechanical properties of solid state systems remains one of the outstanding challenges in the development of quantum technologies.

1.6 Universal Behaviour

Due to the diverging correlation length and the power-law behaviour assumed by various thermodynamic quantities (such as the specific heat, correlation length, order parameter scaling and susceptibility, to name but a few), the behaviour of a system in the vicinity of a second-order phase transition can be completely described by a set of numbers known as critical exponents. Similar to the correlation length exponent ν, each of these thermodynamic quantities has a corresponding critical exponent. These are shown in Table 1.1.

Common critical exponents	
specific heat	$C \sim |\Delta|^{\alpha}$
magnetisation	$M \sim |\Delta|^{\beta}$
magnetic susceptibility	$\chi \sim |\Delta|^{-\gamma}$
field dependence of M	$M \sim |(h - h_c)/h_c|^{1/\delta}$
spatial correlation length	$\xi \sim |\Delta|^{-\nu}$
temporal correlation length	$\tau \sim \xi^{d_z} = |\Delta|^{-d_z \nu}$

Table 1.1 *A list of common thermodynamic quantities and how they scale with their critical exponents, where $\Delta = (T - T_c)/T_c$ in the classical case and $\Delta = (g - g_c)/g_c$ in the quantum case. In the fourth line, h is the external magnetic field. Information from Ref. [9].*

The last entry of Table 1.1 shows the role that the dynamical critical exponent d_z plays in the relation between the temporal correlation length τ and the spatial correlation length ξ and illustrates the extra effective dimensionality of $T = 0$ quantum systems.

The different critical exponents are related to one another through the scaling relations shown in Table 1.2, where most of the exponents are as shown in Table 1.1 and η is the anomalous dimension of the correlation function, defined by the scaling form:

$$ C(\mathbf{r}) \sim \begin{cases} \frac{1}{|\mathbf{r}|^{d-2+\eta}} & |\mathbf{r}| \ll \xi \\ \exp(-|\mathbf{r}|/\xi) & |\mathbf{r}| \gg \xi \end{cases} . \tag{1.7}$$

This expression reflects that the correlation function changes from exponential to power-law scaling at the length scale ξ given by the correlation length. In other words, when ξ diverges at the critical point, we move from short-ranged exponential decay to long-range power-law correlations. This is only strictly true at the critical point: even in the ordered phase, the connected correlation function need not display power-law decay. Its precise form depends on the excitation spectrum of the system in question.

These relations mean that it is not necessary to measure or calculate each critical exponent individually - knowledge of any two of the critical exponents is sufficient to uniquely determine the others and fully describe the scaling of a

Scaling laws	
Fisher	$\nu(2 - \eta) = \gamma$
Rushbrooke	$\alpha + 2\beta + \gamma = 2$
Widom	$\beta(\delta - 1) = \gamma$
Josephson	$2 - \alpha = \nu d$

Table 1.2 *Scaling laws. Information from Ref. [9]. The final relation is the only one to involve the spatial dimension and is known as a hyperscaling relation. Hyperscaling does not hold in $d > 4$ when mean-field theory becomes exact and the correlation length exponent no longer has any dependence on spatial dimensionality.*

Table of critical exponents					
	α	β	γ	δ	ν
$d = 2$ Ising	0	1/8	1/4	15	1
$d = 3$ Ising	0.110(1)	0.6265(2)	1.2372(5)	4.789(2)	0.6301(4)
$d = 3$ XY	0.0146(8)	0.3485(2)	1.3177(5)	4.780(2)	0.67155(27)
$d = 3$ Heisenberg	-0.122(9)	0.3662(25)	1.390(5)	4.796(1)	0.7073(30)
$d = 4$ Ising	0	1/2	1	3	1/2

Table 1.3 *An illustrative and non-exhaustive list of calculated critical exponents for a variety of different universality classes. The numerical values were sourced from Refs. [32–34] and the numbers in brackets indicate the uncertainty in the final digit. The $d = 4$ Ising results are exactly given by mean-field theory. Wherever any particular exponents were not given in the references they were calculated from the ones given using the scaling relations listed in Table 1.2 and the uncertainties propagated in the standard manner.*

given system. To illustrate this, Table 1.3 shows a selection of critical exponents for a variety of different models in different dimensions. Within the uncertainties, it is possible to use the scaling relations in Table 1.2 to calculate any exponent from any other two. (The uncertainties listed are in most cases the result of direct calculation from the sources referenced. Consequently they may be smaller than the resulting uncertainties obtained if one calculates the exponent from others listed and propagates their errors accordingly.)

Since the critical exponents are determined from fundamental underlying physics rather than microscopic details, often very different physical systems display the same set of critical exponents. These sets of critical exponents are called universality classes.

Being able to group phase transitions into universality classes regardless of the specific material or system in which they are realised is an immensely powerful tool and one of the main reasons why many areas of theoretical condensed matter physics work at all. If universal behaviour did not exist and the behaviour of every

physical system depended critically on its specific microscopic structure, much of the predictive power of condensed matter theory would be lost. The ability to instead pick out the most relevant pieces of the underlying physics and discard the irrelevant degrees of freedom is what allows our models to be so general, not to mention one of the main reasons they are mathematically tractable to begin with.

1.7 Renormalisation Group

Arising directly from the idea of focusing only on the most relevant underlying physics, the main technique we will employ in Chapters 3 and 4 is the renormalisation group, a powerful analytical technique allowing us to 'zoom out' from our microscopic picture of individual quantum particles and instead examine the global behaviour of a thermodynamically large system. Renormalisation works because of the diverging correlation length in the vicinity of quantum critical points, allowing us to eliminate the short wavelength non-universal degrees of freedom in order to obtain a description of the system entirely in terms of the long-wavelength universal modes which determine its bulk behaviour. Initially employed in the context of quantum electrodynamics [35], the renormalisation group in its condensed matter form was developed in Refs. [36, 37] and has since become one of the most important tools used in condensed matter theory.

Essentially, renormalisation is a technique used to move away from a microscopic picture to a macroscopic one in a procedure of 'controllably forgetting' about the high energy, short wavelength degrees of freedom. Formally, we typically define a cutoff length or energy scale and integrate out the modes which we are not interested in, thereby leaving an action ostensibly in terms of (for example) the long wavelength, low energy modes but that has still taken the higher energy fluctuations into account. See Refs. [38, 39] for a full and comprehensive review of renormalisation group techniques.

In a renormalisable system, variables fall into three main categories. A 'relevant' variable is one which grows in magnitude as we zoom out, and therefore is important to the long-wavelength behaviour of the system. Conversely, an 'irrelevant' variable is one that renormalises to zero, indicating that it has no bearing on the long-wavelength behaviour which we are interested in (though it can contribute quantitative corrections). In between these two extremes, a variable that does not scale at all under the renormalisation group process is called 'marginal' since it isn't clear from a simple scaling argument whether or not it is important to the physics we are interested in. By tracking how these variables evolve under the renormalisation group process (the renormalisation group flow), we can calculate the long-wavelength behaviour of the model.

A fourth class exists for variables which are themselves irrelevant but have a significant knock-on effect on other variables, e.g. if a variable γ depends on

the inverse power of a variable Λ then although Λ may itself be irrelevant, it must be retained in the theory in order to properly account for the behaviour of γ. These variables are known as 'dangerously irrelevant' and contribute to the leading scaling behaviour despite being themselves irrelevant.

Fig. 1.4 shows a cartoon representation of the renormalisation group process. Panel i) shows an initial field, panel ii) shows the integration over the high frequency components (here performed with a low-pass filter) and panel iii) shows the final rescaling such that the final 'renormalised' field resembles the initial field except that it now contains fewer degrees of freedom.

A key assumption of renormalisation group is that one is free to integrate out the short-wavelength, high energy modes independently of the long wavelength universal behaviour of the system. In a disordered system, this assumption is not necessarily valid [41], as we shall see later. In strongly disordered systems, the short-wavelength behaviour can implicitly depend on the long-wavelength behaviour and care must be taken to implement the renormalisation group procedure properly.

The 'zooming out' procedure may equivalently be performed in real space or in momentum-space [40], and renormalisation more generally can also be applied in other frameworks such as density matrix renormalisation group [42, 43]. In this thesis, we focus on momentum-shell renormalisation group where we integrate out infinitesimal shells of highest momenta modes and compute the effect they have on the remaining lower-momenta modes. With that in mind, we shall specialise to momentum-shell renormalisation group and look at an example calculation.

Momentum-Shell Renormalisation Group

Beginning from a partition function defined in terms of some field ψ in the following way:

$$\mathcal{Z} = \int \mathcal{D}[\psi] \mathrm{e}^{-S[\psi]}, \tag{1.8}$$

we can split the field into 'fast' high momentum fields $\psi_>$ and 'slow' low-momentum fields $\psi_<$ such that

$$\psi = \psi_< + \psi_>. \tag{1.9}$$

We may then integrate out the fast fields by writing the partition function in the following form:

$$\mathcal{Z} = \int \mathcal{D}[\psi_<]\mathcal{D}[\psi_>] \mathrm{e}^{-S[\psi_<,\psi_>]}. \tag{1.10}$$

To obtain an expression for a free energy that only depends on the slow fields, we need to perform the integral over the fast fields. If there is no cross-coupling

i)

ii)

iii)
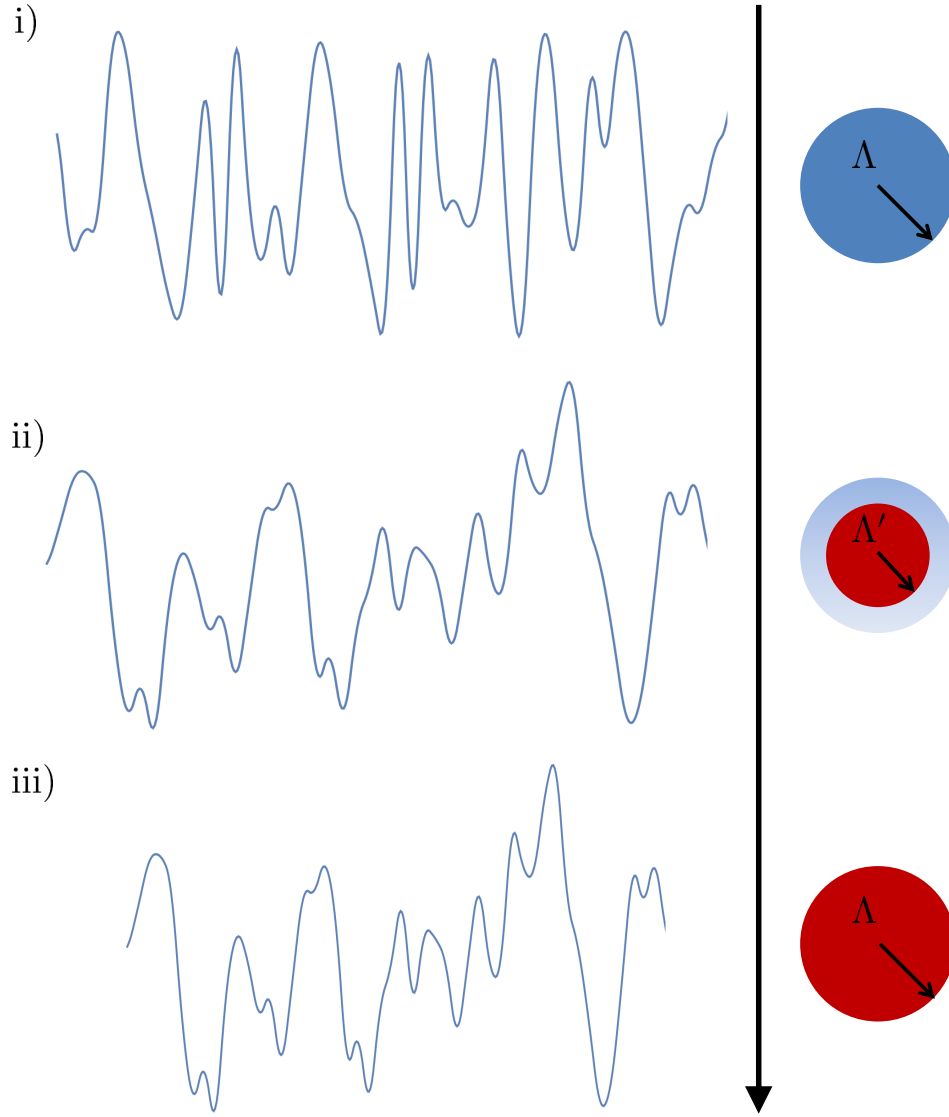
Figure 1.4 *A cartoon representation of the renormalisation process. a) The initial field and its original spherical Brillouin zone of radius Λ. b) The field after removal of the high-frequency modes and its new smaller Brillouin zone of radius Λ'. c) The final field after rescaling the length such that the oscillations look like those of the initial field and its Brillouin zone with radius Λ in the rescaled units. (Figure based on a similar schematic representation in Ref. [40]. To produce this cartoon, the elimination of the higher frequencies was performed with a low-pass filter.)*

between the fast and the slow fields, they can be integrated out independently of one another but in general the fast and slow fields will be mixed. We can write our action as a 'fast' action, a 'slow' action and a contribution which mixes the fast and slow fields:

$$S[\psi_<, \psi_>] = S_0[\psi_<] + S_1[\psi_>] + \delta S[\psi_<, \psi_>], \tag{1.11}$$

such that we can write:

$$A\mathrm{e}^{-\tilde{S}[\psi_<]} = \mathrm{e}^{-S_0[\psi_<]} \int \mathcal{D}[\psi_>] \mathrm{e}^{-S_1[\psi_>]} \mathrm{e}^{-\delta S} \tag{1.12}$$

$$= \mathrm{e}^{-S_0[\psi_<]} \langle \mathrm{e}^{-\delta S} \rangle_{S_1[\psi_>]}, \tag{1.13}$$

where A is some unimportant constant and we compute the average with respect to $S_1[\psi_>]$. Treating this term perturbatively, we may write:

$$\langle \mathrm{e}^{-\delta S[\psi]} \rangle_{S_1[\psi_>]} \approx \left\langle 1 - \delta S + \frac{1}{2}(\delta S)^2 + ... \right\rangle \tag{1.14}$$

$$= 1 - \langle \delta S \rangle + \frac{1}{2} \langle \delta S^2 \rangle \tag{1.15}$$

$$= \exp\left(\ln\left(1 - \langle \delta S \rangle + \frac{1}{2} \langle \delta S^2 \rangle \right) \right) \tag{1.16}$$

$$\approx \exp\left(-\langle \delta S \rangle + \frac{1}{2} (\langle \delta S^2 \rangle - \langle \delta S \rangle^2) \right). \tag{1.17}$$

All that remains is to calculate the averages. Following the established paradigm of momentum-shell renormalisation group, we do this using Wick's theorem. This requires that the 'bare' action be quadratic in the fields ψ and we treat any higher-order terms perturbatively.

After evaluating the averages and re-exponentiating the remaining 'slow' field terms, we arrive at some new effective action $\tilde{S}[\psi_<]$. The system is renormalisable if $\tilde{S}[\psi_<]$ takes the same functional form as the initial action. It is remarkable that so many physical systems turn out to be renormalisable in this sense.

We then rescale to get back the 'resolution' lost by integrating out this momentum shell (shown schematically in Fig. 1.4 going from panel b) to panel c)), and this allows us to derive a set of RG equations showing us how the coefficients of the terms in the action change as we successively integrate out more infinitesimally thin shells. By studying these equations, we can find the scale-invariant fixed points that determine the collective behaviour of the system.

Gaussian Model

As a relatively simple example, consider a model with Gaussian action (as described in Ref. [40]). This example can be solved exactly as there is no coupling

of the fast and slow fields (i.e. $\delta S = 0$) and all the integrals are Gaussian.

$$S = \int_0^1 (ck^2 + r)\bar{\psi}\psi \quad \mathrm{d}k. \tag{1.18}$$

We can split the fields into fast and slow as described previously and integrate out the fast fields exactly to obtain an action of the following form:

$$\tilde{S} = \int_0^{\mathrm{e}^{-\mathrm{d}l}} (ck^2 + r)\bar{\psi}\psi \quad \mathrm{d}k. \tag{1.19}$$

We then want to rescale so the integration range is restored from 0 to 1, and we want to make the rescaling such that the k^2 prefactor remains unchanged. This corresponds to the transformations

$$k \to k\mathrm{e}^{\mathrm{d}l}, \tag{1.20}$$

$$|\psi|^2 \to \mathrm{e}^{\lambda \mathrm{d}l}|\psi|^2, \tag{1.21}$$

and we then choose λ such that the prefactor of k^2 remains unchanged. Under this rescaling the mass term r becomes:

$$r \to r\mathrm{e}^{2\mathrm{d}l} \approx r(1 + 2\mathrm{d}l), \tag{1.22}$$

and we can finally write down the renormalisation group equation for r describing how it scales with l:

$$\frac{\mathrm{d}r}{\mathrm{d}l} = 2r. \tag{1.23}$$

The term $2r$ is known as the 'tree-level scaling' term, and it describes the change of the bare coefficient r due to the loss of resolution and subsequent rescaling, i.e. with each rescaling step $r \to 2r$. Without any other variables in the problem to contribute to the scaling of r, we arrive at the result that r is a relevant variable which increases without bound under the renormalisation group process. This $r \to \infty$ limit corresponds to the particles in the system becoming completely uncorrelated with each other. This is because the Gaussian model does not take into account interactions between particles. Only the (unstable) point $r = 0$ is immune to this divergence. This is known as a fixed point, and it is the scale-invariant fixed points which are of critical importance in determining the properties of the system in question.

Fixed Points

The Gaussian model in the preceding section is exactly solvable but contains no interesting behaviour. In general, there will be cross-couplings between fast and slow fields and an RG equation will take the form $F'(l) = \alpha F(l) + \beta$ where the α refers to the tree level scaling and the β is the renormalisation due to the coupling between fast and slow fields. If $\alpha > 0$ then the variable $F(l)$ will initially increase under RG, i.e. it is a relevant variable. Similarly, $\alpha < 0$ defines an irrelevant variable which will renormalise to 0. In the case of $\alpha = 0$ the variable is marginal and does not scale under the RG - the higher-order vertex contributions (i.e. the β terms) must then be computed to determine whether the variable has any effect at all on the macroscopic behaviour of the system. The β terms have a more important role than this, however - they can also give rise to non-trivial fixed points where $F'(l) = 0$ which allow us to determine the behaviour of the system.

Once we have a set of renormalisation group equations for all running variables in the problem, we can solve them simultaneously for all the points $F'(l) = 0$ to obtain the points in parameter space which are stable under the renormalisation group flow. These are known as fixed points, as if the running variables take precisely these values they will not change under the renormalisation group process. In practice, the variables will never take precisely the fixed-point values but will approach them asymptotically. As the variables flow closer to the fixed point values, the system begins to look the same at all length scales, a property known as scale invariance. The most famous examples of scale invariance are fractals, as shown in Fig. 1.5, where any arbitrarily small sub-region contains the full information of the entire system.

In Section 1.3 we saw that we can always define a correlation length of a system. The existence of a well-defined dominant length scale does not naturally fit with the idea of scale-invariance exhibited at the renormalisation group fixed points. This implies that at the fixed point, either the correlation length is zero, or it is infinite. In the vicinity of a second-order phase transition the correlation length will diverge, and therefore we can identify the renormalisation group fixed points as potentially indicative of second-order phase transitions.

Following the relevant/irrelevant/marginal classification of variables under the renormalisation group process, the fixed points come in three different types:

- **Stable fixed points** where all variables are either irrelevant or marginal. These fixed points are attractive and correspond to stable phases.

- **Unstable fixed points** where all variables are relevant. The fixed point is repulsive and all variables flow away from it. Strictly at the fixed point, the system is scale invariant and nothing will flow, however any infinitesimal perturbation away from the fixed point values will result in the variables irreparably flowing away from it. These are solutions of the RG equations in
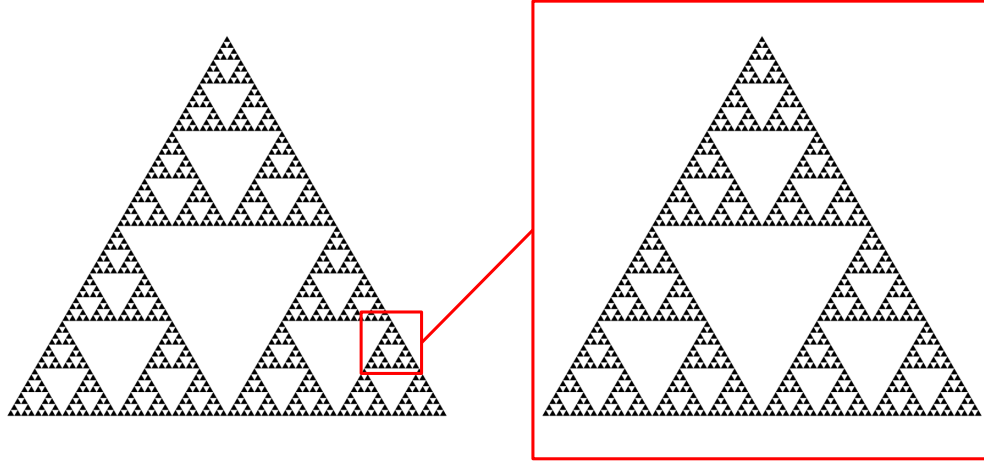
Figure 1.5   *The Sierpinski triangle, a fractal which repeats indefinitely as one zooms in. Each small subsection contains the complete information about the full structure. This is a dramatic example of scale-invariance (or 'self-similarity').*

the same sense that one can balance a pencil on its sharp end - technically valid but physically unrealisable.

- **Generic fixed points** with a mixture of relevant and irrelevant variables. These correspond to transitions between phases and act to drive the system towards stable fixed points corresponding to stable phases.

Solving the set of renormalisation group equations for all the variables in the problem thus allows us to find out what possible phases the model can adopt and describe the transitions into and out of those phases. This is the goal of a renormalisation group study of any system - the establishment and investigation of the fixed points.

The principles of scale invariance and universality lead to the renormalisation group being an enormously powerful method with which to investigate strongly correlated quantum systems. But what happens when we are no longer dealing with pristine, pure materials but instead consider the addition of disorder - how can we hope to systematically investigate randomness?

1.8   Disordered Systems

The preceding sections have dealt with clean systems, assuming there to be no impurities or randomness. In any real material, there will be anything from trace

impurities through to significant inhomogeneities that can drastically change the properties of the material. Experimentally, even the purest of samples can display important disorder-induced effects, and that's not even to mention the materials which are deliberately doped or disordered in an effort to modify their properties.

Even more so than classical systems, quantum materials can exhibit extreme sensitivity to the addition of disorder. Sometimes this disorder can be advantageous, as in the case of the doped cuprates [2, 44–47] and pnictides [48–51] where doping can induce superconductivity, but other times it can destroy the very features we're interested in, such as in Sr_2RuO_4 [52] where disorder rapidly suppresses the superconducting critical temperature.

The way a material responds to disorder can tell us something about the physics of the clean system, such as in the case of Sr_2RuO_4 [52] where the suppression of T_c tells us that the material has a non-s-wave order parameter[3]. Also, discontinuous changes of system parameters known as quenches can lead to unusual behaviour unique to disordered systems [55] and act as a probe of the underlying physics.

The study of disorder in condensed matter physics has a long history. A series of informal articles putting the problem of disorder in spin glass systems in historical context may be found in Refs. [56–62] which outline the development of the tools used to treat classical spin glass problems. Understanding localisation in strongly interacting disordered systems is a long-standing theoretical problem, and one that has become more relevant in recent years with the rise of complex many-body localisation effects both in [63] and out [64] of equilibrium. In this thesis we shall restrict ourselves to ostensibly equilibrium phases.

The systematic study of randomness is a significant challenge. The addition of disorder into a material immediately destroys translational invariance, rendering many common techniques unsuitable. In numerical simulations on small systems, care must be taken to ensure that the results do not depend on the particular realisations of disorder studied. In Chapter 5 we shall look at the local effects of disorder but in the majority of this thesis we are interested in extracting the bulk behaviour in the thermodynamic limit and so averages over different disorder realisations must be performed. Numerically, this boils down to re-running the simulations many times for many different disorder distributions, but analytical methods of disorder averaging are not so straightforward.

There are a variety of methods one may use to analytically study disordered systems. For example, disorder can be averaged out using the cavity method [65] in certain systems or handled with Keldysh techniques [66] in others. The method which we shall use in this thesis is one of the most general and widely-used techniques for disorder averaging and is known as the replica trick.

---

[3]An s-wave order parameter would be robust against disorder-induced scattering, whereas a p- or d-wave would be subject to strong destructive interference, as seen experimentally. Sr_2RuO_4 is believed to be a p-wave superconductor though the evidence remains circumstantial [52–54].

The Replica Trick

Most thermodynamic parameters we are interested in will involve calculations of the logarithm of the partition function, e.g. the free energy is given by $\mathcal{F} = -T \ln \mathcal{Z}$. This can lead to problems when trying to calculate disorder averages, as the average of the logarithm of $\mathcal{Z}$ is often extremely difficult to compute directly. We could consider performing some perturbative expansion of $\mathcal{Z}$, but an expansion of $\mathcal{Z}$ itself to any given order does not necessarily determine the expansion of $\ln \mathcal{Z}$ to the same order.

The replica trick [67–69] allows us to circumvent the difficulties in calculating the disorder average of $\ln \mathcal{Z}$ by rewriting the logarithm using the identity:

$$\overline{\ln \mathcal{Z}} = \lim_{n \to 0} \frac{\overline{\mathcal{Z}^n} - 1}{n}. \tag{1.24}$$

Rewriting the logarithm in this manner allows us to transform the difficult problem of taking the disorder average of the logarithm of the partition function into the much simpler (both conceptually and mathematically) task of taking the disorder average of $\mathcal{Z}^n$. The precise form of the cumulant expansion used in this thesis is shown in Appendix B.

In the clean case, Eq. 1.24 is exact but in the disordered case it is not necessarily so. In the clean case, n is not required to be an integer and the $n \to 0$ limit is well-defined. In the presence of disorder, Eq. 1.24 corresponds to making n discrete copies of the system with the same disorder realisation and averaging across all of these replica systems, then in the end taking the limit $n \to 0$ to recover the disorder-averaged $\overline{\ln \mathcal{Z}}$. Each replica of the system has the same disorder realisation, but is allowed to be in a different ground state.

The reason this method is known as the 'replica trick' is that the process of taking the limit $n \to 0$ is poorly defined. When taking the disorder average we assume we create n replica systems where n is an integer greater than zero. When taking the limit, we then send n continuously to zero. This discrepancy means we must be extremely careful in taking the replica limit. This step is the source of most of the controversy relating to the replica trick: the analytic continuation involved is often poorly controlled, and used without care can lead to unphysical results such as negative entropy in even the most straightforward spin glass systems [70]. We shall see later that the way to properly account for the subtleties in implementing the replica trick is to introduce additional order parameters and allow for a property known as replica symmetry breaking, which lifts the implicit assumption that the system has a single ground state. In the disordered systems we are most interested in, this is often not the case.

Spin Glasses

The main focus of this thesis is on transitions to and from glassy phases, which are different from the phase transitions mentioned in the previous section. Often, glassy phases do not break any symmetries and so there is no convenient order parameter such as magnetisation.

The term 'glass' generically refers to an amorphous, disordered system lacking any form of long-range order. Window glass is the obvious example, as the molecules do not form a crystal lattice structure. The non-crystalline nature of silicate window glass is in fact widely acknowledged in popular culture through the erroneous assertion that old windows are thicker at the bottom because the glass is really just a 'slow liquid'. This is not true [71], but it illustrates an important point about glasses in general that makes them distinct from typical phases: glasses are not strictly equilibrium states. They are the result of a system 'freezing' into a given configuration which may or may not be its ground state, and given enough time the system may still attempt to fall back to its ground state.

The most commonly encountered glasses in condensed matter physics are spin glass phases, a variation of which we shall look at in Chapter 2. Spin glasses are frozen disordered magnets, where the magnetic moments of the electrons are frozen in a random arrangement. They are 'glassy' in the sense that their lack of magnetic order is analogous to the lack of crystalline order in an amorphous solid such as window glass. The random arrangement of electron spins differs from a spin liquid or paramagnetic state in that the spins in a spin glass are pinned and are not free to move. Above some critical temperature the spin glass phase will 'melt' and the material will become paramagnetic again, but when cooled below this critical temperature the spins will freeze into some random arrangement.

Spin glasses are characterised by a highly degenerate ground state with a macroscopic number of different possible configurations. A simple example of how this degeneracy can arise in a spin system (here via geometric frustration rather than disorder) is that of antiferromagnetically coupled Ising spins on a triangular lattice (Fig. 1.6). It is impossible for all of the spins to satisfy all of their bonds. At least one bond must be in an energetically unfavourable configuration, a phenomenon known as frustration. Since there are multiple different ways to frustrate a single bond, there is a large ground state degeneracy.

Over a long enough time scale, this system will tunnel from one configuration into another. In a true spin glass, where the degeneracy is caused by disorder rather than frustration alone, the energy barrier between the different configurations is sufficiently large that, in the thermodynamic limit, tunnelling is completely suppressed. Below the spin glass critical temperature, the available configuration space becomes fragmented into 'valleys' separated by large energy barriers which prevent the system from changing from one configuration to another.
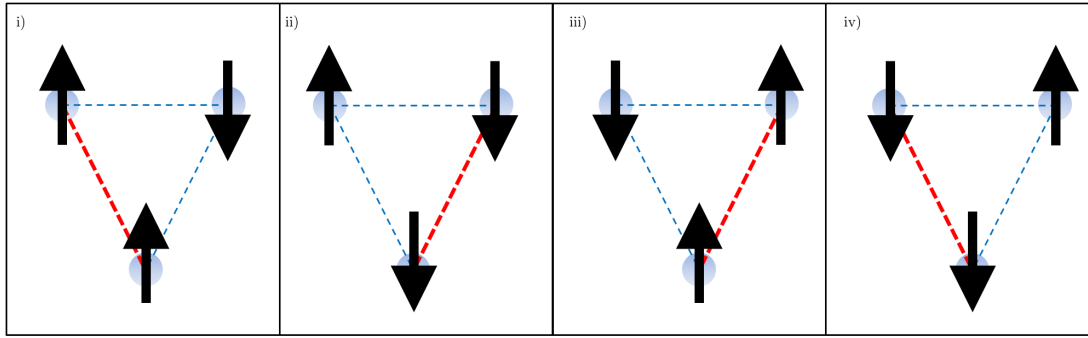
Figure 1.6 *Four of the possible arrangements of three Ising spins on a triangular lattice with antiferromagnetic exchange interactions. The red lines indicate frustrated bonds. Panels i) and ii) show that if the top-left and top-right spins are up and down respectively, then the bottom spin cannot be up without frustrating the leftmost bond) and cannot be down without frustrating the rightmost bond. Likewise, panels iii) and iv) show that if the top two spins are swapped, the same problem occurs in reverse.*

Ergodicity

A system is ergodic if it is capable of exploring the entire configuration space available to it, i.e. every microstate with the same energy is equally probable over infinitely long timescales. A non-ergodic system is one in which this statement does not hold and the system is restricted to a certain region of its phase space. The ergodic hypothesis [72] states that typical physical systems are ergodic, which is essentially the statement that time averages are equivalent to thermal averages.

Though typically thought of as a characteristic of strongly disordered systems which 'freeze' the system into a small subset of all available states, the breaking of ergodicity is actually quite common [73]. Take an Ising ferromagnet, for example. Below the Curie temperature at which the material becomes ferromagnetically ordered, the magnetisation will spontaneously choose to be 'up' or 'down' but once the choice is made the system cannot move from one configuration to the other.

An ordered ferromagnet is non-ergodic in the sense that, if one computes expectation values of magnetisation according to the ergodic hypothesis over the all available configurations (i.e. over both up and down magnetisations) the result is zero. For any system in the thermodynamic limit this does not reflect its true behaviour[4], and so expectation values must instead be computed over a restricted volume of phase space, either the up configuration or the down configuration. The

---

[4]Strictly, an ordered ferromagnet is only non-ergodic in the theromodynamic limit where the energy barrier between 'up' and 'down' configurations becomes formally infinite and the probability of the system to tunnel from one state to another is zero. In any finite-size system, there will remain a finite probability for the system to switch from one state to another. This probability will decrease with increasing system size.
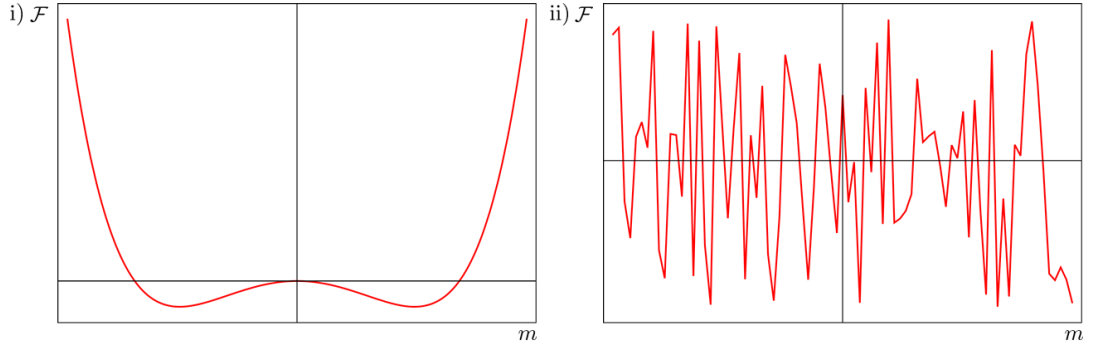
Figure 1.7 *Schematic free energy landscape of a ferromagnetic system (left, with well defined minima related by a symmetry of the Hamiltonian) versus a glassy system (right, with lots of metastable local minima unrelated by symmetries and with large energy barriers between them). Figure based on a similar figure in Ref. [74].*

violation of the ergodic hypothesis in this way is characteristic of spontaneous symmetry breaking.

In the Ising example, the two phase space regions are related by a symmetry of the Hamiltonian describing the system. Disordered systems realise a much richer and more complex form of ergodicity breaking, displaying a macroscopic number of well-separated phase space regions not linked by any fundamental symmetries of the Hamiltonian. A schematic of this is shown in Fig. 1.7 where we compare the free energy landscape of an ordered ferromagnet (non-ergodic, with the global minima related by the $\mathbb{Z}_2$ symmetry of the Hamiltonian) with the free energy landscape of a spin glass which has a large number of minima unrelated by any discernible symmetry of the Hamiltonian.

Order Parameters

Given that glassy phases are, by definition, lacking in order, defining a quantity that can serve as an order parameter for the breaking of ergodicity in a glassy phase is challenging. The most commonly used quantity is the Edwards-Anderson order parameter, first coined in the context of spin glasses [68, 70, 75] where it took the form:

$$q_{EA} = \lim_{t \to \infty} \lim_{N \to \infty} \overline{\langle \hat{S}_i(0) \hat{S}_i(t) \rangle}, \tag{1.25}$$

where $\hat{S}_i$ is the spin operator on site i, the angled brackets indicate a thermal average and the overline indicates a disorder average. The Edwards-Anderson order parameter in this form will be non-zero only if the system is trapped in a single region of its configuration space. By taking the limit of $N \to \infty$ before $t \to \infty$, the energy barriers between free energy regions become infinite and win

out over the infinite timescale. The system becomes trapped in its valley, leading to a non-zero value of q_{EA}.

We can also define a purely thermally-averaged 'equilibrium' order parameter with no time-dependence:

$$q = \overline{\langle \hat{S}_i \rangle^2}, \tag{1.26}$$

which differs from q_{EA} in having contributions from all different configuration space 'valleys', e.g. the average is over all possible minima in the free energy landscape.

The forms of the order parameters above were introduced with the sole motivation of distinguishing spin glasses from paramagnetic phases, but a ferromagnet is also non-ergodic and will also have a non-zero q_{EA} as defined above. The following more modern definition is trivially zero in the paramagnet, identically zero in the ferromagnet and non-zero only in the spin glass. It takes the following form:

$$q_{EA} = \lim_{t \to \infty} \lim_{N \to \infty} \left[\overline{\left\langle \hat{S}_i(0) \hat{S}_i(t) \right\rangle} - \overline{\left\langle \hat{S}_i(0) \right\rangle}^2 \right]. \tag{1.27}$$

The pure thermally-averaged order parameter q becomes:

$$q = \left[\overline{\langle \hat{S}_i \rangle \langle \hat{S}_i \rangle} - \overline{\langle \hat{S}_i \rangle} \, \overline{\langle \hat{S}_i \rangle} \right]. \tag{1.28}$$

In replica language, q_{EA} measures long-time correlations within a replica and q is the equal-time correlation function between multiple replicas:

$$q_{EA} = \lim_{t \to \infty} \lim_{N \to \infty} \sum_{\alpha} \left[\overline{\left\langle \hat{S}_{i,\alpha}(0) \hat{S}_{i,\alpha}(t) \right\rangle} - \overline{\left\langle \hat{S}_{i,\alpha}(0) \right\rangle}^2 \right], \tag{1.29}$$

$$q = \sum_{\alpha\beta} \left[\overline{\langle \hat{S}_i^\alpha \rangle \langle \hat{S}_i^\beta \rangle} - \overline{\langle \hat{S}_i^\alpha \rangle} \, \overline{\langle \hat{S}_i^\beta \rangle} \right], \tag{1.30}$$

where α and β label the different replica systems. This form of q makes it easier to see where the intervalley contributions come from, and how q physically differs from q_{EA} - if the ground states of replicas α and β are separated by a large energy barrier, their contribution to q will be much smaller than if their ground states exist within the same free energy valley. This gives rise to a phenomenon known as replica symmetry breaking.

Replica Symmetry Breaking

In a glassy system both q and q_{EA} will take non-zero values. If $q = q_{EA}$ we say the system is replica symmetric. If $q \neq q_{EA}$, we say the system exhibits replica symmetry breaking. Equivalently, if we define $q = \sum_{\alpha,\beta} q_{\alpha\beta}$ then replica
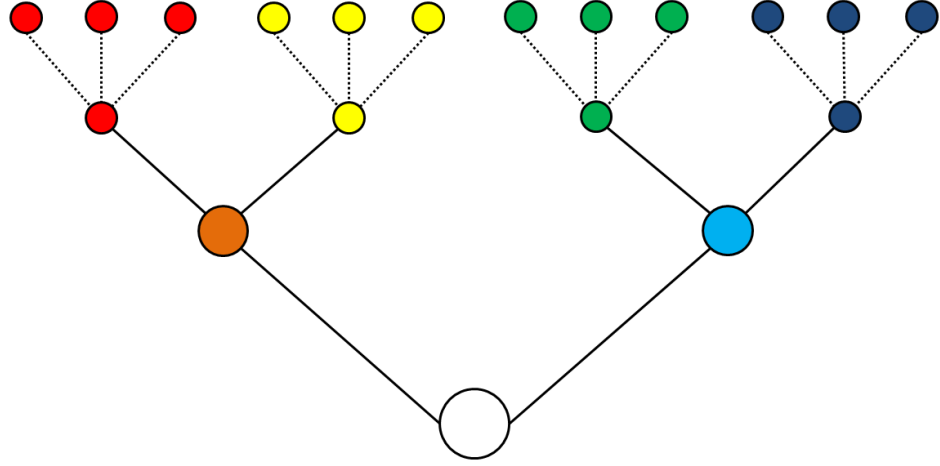
Figure 1.8 *A cartoon of the hierarchical free energy landscape of a spin glass. The entire configuration space (white) first splits into two valleys (orange and cyan), which then split into further valleys (red, yellow, green and blue) each of which has its own subset of possible states. Each colour in the top row (red, yellow, green, blue) corresponds to a different replica index.*

symmetry breaking is the same as saying that $q_{\alpha\beta}$ takes different values depending on which replicas α and β are being considered.

Replica symmetry breaking (RSB) is the solution to the unphysical behaviour that is sometimes caused by naive application of the replica trick. If a system exhibits RSB, it means there is a macroscopic number of degenerate or nearly-degenerate states in the configuration space, each separated by large energy barriers. If a system freezes into one glassy state, it may be inhibited from tunnelling into any other states with the same energy. Typically, the free energy landscape takes on a so-called hierarchical form, with local minima within local minima within local minima, and so on. An illustration of this is shown in Fig. 1.8 where the different replicas are signified by different colours.

The Edwards-Anderson order parameters test for this. Because q_{EA} measures long-time correlations within a single replica, it contains information about how likely it is for the system to tunnel from one replica configuration into another. The other order parameter q measures the overlap between different replicas, and specifically $q_{\alpha\beta}$ measures the overlap between specific pairs of replicas. If all replicas are equivalent to each other, $q_{EA} = q$ and the system is replica symmetric. If the disorder has led to the fractionalisation of the configuration space, with a large number of potential replica configurations which are unable to tunnel into one another, then $q_{\alpha\beta}$ will vary with α and β, leading to $q_{EA} \neq q$ and we know the system exhibits replica symmetry breaking.

It is precisely this idea that prevents blind application of renormalisation group techniques from working on disordered systems. If the systems exhibits replica symmetry breaking, the short-wavelength and seemingly irrelevant degrees of freedom are in fact completely determined by which replica configuration the system is in, and therefore are implicitly dependent upon the long-wavelength properties of the system. The interlinked nature of the short- and long-wavelength modes in disordered systems can be taken into account by incorporating replica symmetry breaking into the renormalisation group scheme [41].

In the mean-field theory of spin glasses, RSB arises through q_{EA} and q in a quite natural manner. In renormalisation, one must break replica symmetry by hand, look for corresponding RSB fixed points and test the susceptibility of the system to replica symmetry breaking perturbations. RSB does not arise in as elegant a manner in these systems, but it is important to allow for it in order to ensure the RG produces physical results.

Although the Edwards-Anderson order parameters were developed for classical spin glasses, they are relevant to any 'glassy' system, classical or quantum. Once we know (or at least suspect) we have a glassy phase, these order parameters can be used to help quantify it. To know whether or not we suspect disorder to cause glassiness, we need to look at the properties of the clean system and see how stable it is to the addition of disorder.

1.9 Phase Transitions in Disordered Materials

Once we begin looking at the effects of disorder on materials, the natural question is to ask what effect disorder has (if any) on the nature of the quantum phase transition itself and the nature of the phases either side of it. It could be that disorder destabilises one phase in favour of another, or that the phase boundary stays the same but critical exponents are modified. In the case of Griffiths phases, disorder causes entirely new 'mixed' phases to appear in between the phases of the clean system. Our first port of call in determining the relevance of disorder is known as the Harris criterion - the following discussion is based upon Refs [76] and [77].

Harris Criterion

The Harris criterion states that disorder is a relevant perturbation to a renormalisation group fixed point in the clean system if $\nu d < 2$ where ν is the correlation length exponent and d is the spatial dimension. In the case of quenched disorder this criterion is equally valid for quantum and classical phase transitions due to the perfect correlation of disorder in time (i.e. d is not replaced with $d + d_z$).

Critical points with quenched disorder can be classified by the behaviour of the average disorder strength [78]. For systems which fulfil the Harris criterion (i.e. for which $\nu d > 2$), the disorder renormalises to zero on large length scales and the critical behaviour is that of the clean system in the absence of disorder. If the Harris criterion is violated ($\nu d < 2$), there are two possibilities:

i) The average disorder strength remains finite at all length scales and the transition is controlled by a finite-disorder critical point, where the scaling is of conventional power-law form but the critical exponents have been modified from the clean system values.

ii) The average disorder strength increases under renormalisation, becoming more and more relevant at longer and longer length scales. The transition is therefore controlled by an exotic infinite-randomness fixed point with unconventional scaling relations.

The Harris criterion and the classifications arising from it deal solely with the behaviour of the average disorder strength. This is not sufficient for a full description of disordered systems - randomness which leads to fluctuations on finite length scales can give rise to new physics not captured by the behaviour of the average disorder strength.

The presence of these relevant fluctuations on finite length scales lead to situations where disorder can cause certain rare regions of the system to be locally in one phase even when (based on the average disorder strength) the bulk system is predicted to be in another. These rare regions lead to nonanalyticities in the free energy known as Griffiths singularities in the vicinity of the phase transition and they can dramatically change the behaviour of the system.

Transitions in the presence of rare regions

The presence of rare regions can lead to three different situations, governed by the relation between the effective dimensionality of the rare regions d_{RR} (which includes imaginary time) and the lower critical dimension of the ordering transition d_c^-.

i) $d_{RR} < d_c^-$: When the dimension of the rare region is lower than the lower critical dimension of the ordering transition, the rare regions cannot order independently of the rest of the system. Their contribution is at most power-law in the volume of the system, whereas the density of the rare regions decreases exponentially. The fixed point is a conventional finite-disorder type with power-law correlations. The rare regions thus lead to exponentially small corrections to the thermodynamic properties.

ii) $d_{RR} = d_c^-$: When the rare regions are precisely at the lower critical dimension, they still cannot order independently of the bulk, however their contribution to the thermodynamics is exponential in the system volume. The
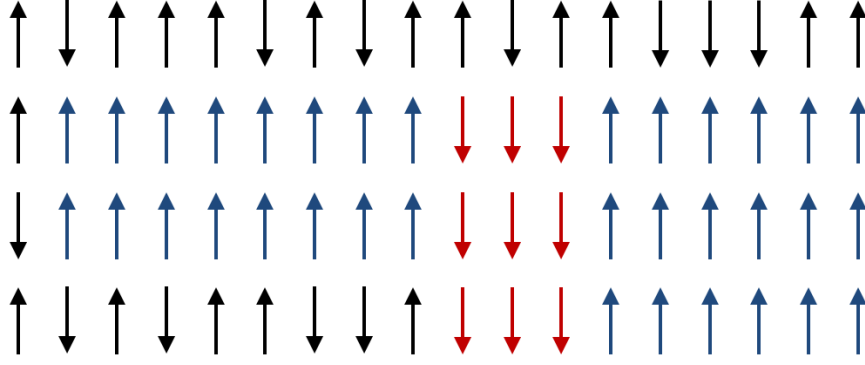
Figure 1.9   *An illustration of an Ising spin system with different regions of the system locally in different phases. The black regions are paramagnetic and exhibit no long-range order, while the red and blue regions refer to patches of locally ordered ferromagnetism.*

effects of the rare regions can therefore overpower the low density, leading to Griffiths singularities in the free energy and non-power-law scaling. The renormalisation group fixed point associated with this type of transition is the so-called 'infinite randomness' type, where the effects of disorder increase as the short-wavelength modes are integrated out and the disorder strength diverges.

iii) $d_{RR} > d_c^-$: In this case, the rare regions *are* able to order independently of the bulk system. Their contribution to the thermodynamics is such that not only does it change the behaviour, but in fact it destroys the sharp transition entirely, smearing it out over a wide region of phase space and leading to the formation of a new 'mixed' phase where rare ordered regions exist within an otherwise disordered system.

The case in which the transition is smeared out leads to a region known as a Griffiths region, or Griffiths phase. In this region, exponentially rare pockets of one phase exist within a bulk system which is in another phase. It is this type of phase that will be the main focus of Chapters 3, 4 and 5.

1.10   Outlook

Having now taken a whistle-stop tour of phase transitions, disordered systems and the main theoretical tools which we will be using in this thesis, we are now equipped to move on to look at the original research I have conducted during my PhD. In this thesis, we will examine a number of different types of physical systems subject to a variety of types of disorder.

- **Chapter 2** - We begin with an investigation into the behaviour of how disorder can couple to quantum fluctuations at finite temperatures. In this chapter, we examine fluctuation-induced phase reconstruction near an itinerant ferromagnetic quantum critical point. Using the fermionic quantum order-by-disorder mechanism, we recover the finding from previous work that fluctuations near the quantum critical point favour the formation of an incommensurate spiral phase in the absence of disorder. Upon the addition of quenched charge disorder and application of the replica trick, the incommensurate phase becomes stabilised over a slightly larger region of the phase diagram but it no longer displays long-range order, instead adopting short-range correlations and a strongly anisotropic correlation length. The disorder causes local changes in the pitch of the spiral magnetism, resulting in a novel phase we term a 'helical glass'.

- **Chapter 3** - Having looked perturbatively at quantum fluctuations at finite temperatures, we now move on to a zero-temperature renormalisation group study of the Bose glass, a rare-region Griffiths phase found in the disordered Bose-Hubbard model. Restricting ourselves to random mass disorder only, we recap previous work done on this model in the context of a replica symmetric field theory and for the first time consider the effects of replica symmetry breaking on the renormalisation group equations. Allowing for replica symmetry breaking in the most general Parisi form, we find the Mott insulator to Bose glass transition is governed by a one-step replica symmetry broken fixed point, signifying a breakdown of ergodicity, and derive Edwards-Anderson order parameters to quantify this.

- **Chapter 4** - Motivated by the need to find experimental systems where replica symmetry breaking could be experimentally detected, we conduct a related renormalisation group study of a dimerised quantum antiferromagnet with random intra-dimer bond disorder, which maps onto a hard-core Bose-Hubbard model. This model contains magnetic analogues of the Mott insulating and superfluid phases of the conventional Bose-Hubbard model as well as fractionally-filled 'checkerboard' phases. When disorder is added, a magnetic analogue of the Bose glass phase intervenes between the Mott insulating and superfluid phases. Using the full replica symmetry breaking renormalisation group technique developed in Chapter 3, we show that without replica symmetry breaking the renormalisation group procedure does not correctly capture the physics of the Bose glass. In addition, we find a strong suppression of the compressibility near the tip of the Mott lobes which we identify with a Mott glass, an incompressible rare-region phase never before analytically predicted to exist in a three-dimensional bosonic system.

- **Chapter 5** - Switching to a fully numerical study for the first time in this thesis, we examine the capabilities of current-generation experimental setups to probe the local properties of glassy phases. Specifically, we consider the case of the quantum gas microscope. Capable of single-atom resolved fluorescence imaging of both bosons and fermions, quantum

gas microscopes can directly image exotic glassy phases. Using a Bose-Hubbard model with mean-field numerics, we model the experimental system, reproduce existing experimental results in the clean case and show that the addition of disorder leads to changes that are well within current experimental resolution, even at finite experimental temperatures. We show for the first time that quantum gas microscopes can measure the thermally-averaged Edwards-Anderson order parameter, serving as a proof-of-principle that quantum gas microscopes are ideal for investigating the properties of disordered strongly correlated phases. With such local probes, many of the theoretical predictions in the rest of the thesis can now be experimentally tested and an exciting new experimental toolbox is opened up.

In all of the following, I work in units where all natural constants are equal to unity, e.g. $\hbar = 1$, $k_B = 1$ and $c = 1$.

Chapter 2

Disordered Itinerant Ferromagnetic Quantum Critical Points

While the majority of this thesis is concerned with the effects of disorder in bosonic systems, the earliest work I did as part of my PhD was a study of the effects of quenched charge disorder on phase reconstruction near an itinerant ferromagnetic quantum critical point. This work was my introduction to disordered systems before I moved on to the bosonic work that forms the bulk of this thesis.

In this work, I look at how disorder can generate new quantum fluctuations in the vicinity of a ferromagnetic quantum critical point and show that this leads to a novel glassy phase known as a helical glass which persists up until a tricritical point at non-zero temperature.

The following work was published in "*Helical glasses near ferromagnetic quantum criticality*", S. J. Thomson, F. Krüger and A. G. Green, *Physical Review B* **87**, 224203 (2013) [79].

2.1 Background

The conventional theory of itinerant ferromagnetic quantum criticality is the Hertz-Millis (or Hertz-Millis-Moriya) theory [80–82], which works by determining an effective action for a conducting fermionic system in terms of dynamical fluctuations of a bosonic order parameter and from there allowing calculations of the free energy and other thermodynamic properties. Although Hertz-Millis theory is successful at describing many aspects of itinerant ferromagnetic quantum critical points, it fails to predict some of the more striking phenomena observed in experiments such as fluctuation-induced first order behaviour and the emergence of new phases near quantum critical points.

The fermionic order-by-disorder mechanism [83] is one of the proposed ways to extend and repair the conventional Hertz-Millis theory by including low-energy particle-hole fluctuations which couple to the order parameter and result in additional non-analytic corrections to the free energy. These fluctuation corrections reproduce the experimentally seen first-order behaviour in the vicinity of quantum critical points as well as the stabilisation of incommensurately ordered magnetic phases.

The name 'order-by-disorder' refers to the competition between internal energy and entropy rather than to any impurities or randomness. Previous work dealt purely with clean systems. My contribution to this mechanism, motivated by Ref. [84], was to consider the addition of quenched charge disorder into the system. We found that new fluctuation-corrections arise due to the presence of disorder in the system and these change the nature of the incommensurate phase from a long-range-ordered spiral ferromagnet to a new fluctuation-induced short-range ordered helical glass phase unlike anything predicted before in the literature.

2.2 Stoner Mean-Field Theory

The conventional mean-field theory of itinerant ferromagnetism in a fermionic Hubbard model is known as Stoner ferromagnetism [85, 86]. Before going on to more sophisticated extensions, we'll take a look at the mean-field description of the paramagnet to ferromagnet transition.

We start from the fermionic Hubbard Hamiltonian:

$$\hat{\mathcal{H}} = \sum_{\mathbf{k},\sigma=\pm} \left[\varepsilon(k) - \mu\right]\hat{n}_{\mathbf{k},\sigma} + g\sum_i \hat{n}_{i,+}\hat{n}_{i,-}, \tag{2.1}$$

where $\hat{n}$ is the fermionic number operator, $\sigma = \pm$ refers to the spin index, $\varepsilon(k) = k^2/2$ is the free-electron dispersion, g is the usual Hubbard contact interaction and we work in units where both the electron mass and $\hbar$ are equal to unity.

As a first step, we employ a mean-field decoupling to rewrite the interaction term $\bar\psi_+\psi_+\bar\psi_-\psi_- \approx M(\bar\psi_+\psi_+ - \bar\psi_-\psi_-)$, assuming without loss of generality that $\mathbf{M} = M\hat{z}$, and that M is uniform across the system. Taking the continuum limit then leads to the mean-field Hamiltonian:

$$\hat{\mathcal{H}}_{MF} = \sum_{\sigma=\pm}\int_{\mathbf{k}}(\varepsilon(k) - \sigma g M - \mu)\hat{n}_{\mathbf{k},\sigma} + g\int \mathrm{d}^3\mathbf{r}\, M^2. \qquad (2.2)$$

This is diagonal in the fermionic operators. One may then express the partition function in the conventional manner as $\mathcal{Z} = \mathrm{Tr}\left[\exp(-\beta\hat{\mathcal{H}}_{MF})\right]$, then use $\mathcal{F} = -T\ln\mathcal{Z}$ to obtain a mean-field free energy density in the following form:

$$\mathcal{F} = gM^2 - T\sum_\sigma\int_{\mathbf{k}}\ln\left(\mathrm{e}^{-\beta(\varepsilon_\sigma(k)-\mu)} + 1\right), \qquad (2.3)$$

where $\varepsilon_\sigma(k) = \varepsilon(k) - \sigma g M$ is the mean-field dispersion in the presence of uniform ferromagnetism and $\beta = 1/T$ is the inverse temperature. This free energy can be expanded in the usual Landau manner as a series expansion in even powers of the magnetisation M:

$$\mathcal{F} = a_2 M^2 + a_4 M^4 + a_6 M^6... \qquad (2.4)$$

where the coefficients a_2, a_4, a_6 etc are all determined in terms of the microscopic parameters and are given by:

$$a_2 = \left.\frac{\partial\mathcal{F}}{\partial M^2}\right|_{M=0} = g + g^2\int_k f^{(1)}(\varepsilon(k)), \qquad (2.5)$$

$$a_4 = \left.\frac{\partial^2\mathcal{F}}{\partial(M^2)^2}\right|_{M=0} = \frac{1}{12}g^4\int_k f^{(3)}(\varepsilon(k)), \qquad (2.6)$$

$$a_6 = \left.\frac{\partial^3\mathcal{F}}{\partial(M^2)^3}\right|_{M=0} = \frac{2}{6!}g^6\int_k f^{(5)}(\varepsilon(k)) \qquad (2.7)$$

where $f(\varepsilon(k)) = [\exp(\beta(\varepsilon(k)-\mu))+1]^{-1}$ is the Fermi-Dirac distribution function and the superscript indicates the number of derivatives with respect to $\varepsilon(k)$. The coefficients can be calculated analytically in the $T = 0$ limit and at low temperatures using the Sommerfeld expansion [87]. The integrals can also be evaluated numerically, however derivatives of the Fermi function become sharply peaked at low temperatures and it can be difficult to ensure convergence.

The second order paramagnet-to-ferromagnet transition is given by the $a_2 = 0$ line, shown in Fig. 2.1. At mean-field level the quartic coefficient a_4 is always greater than zero, ensuring the free energy is bounded from below and there is no first-order transition. Though this means we do not yet require the sixth-order coefficient a_6, we will end up needing this later once we start including the effects of quantum fluctuations that go beyond mean-field.

Using Eq. 2.5 the condition for ferromagnetism $a_2 < 0$ can also be expressed as $g\rho(\varepsilon_F) > 1$, where $\rho(\varepsilon_F) = -\int_k f'(\varepsilon(k))$ is the density of states at the Fermi
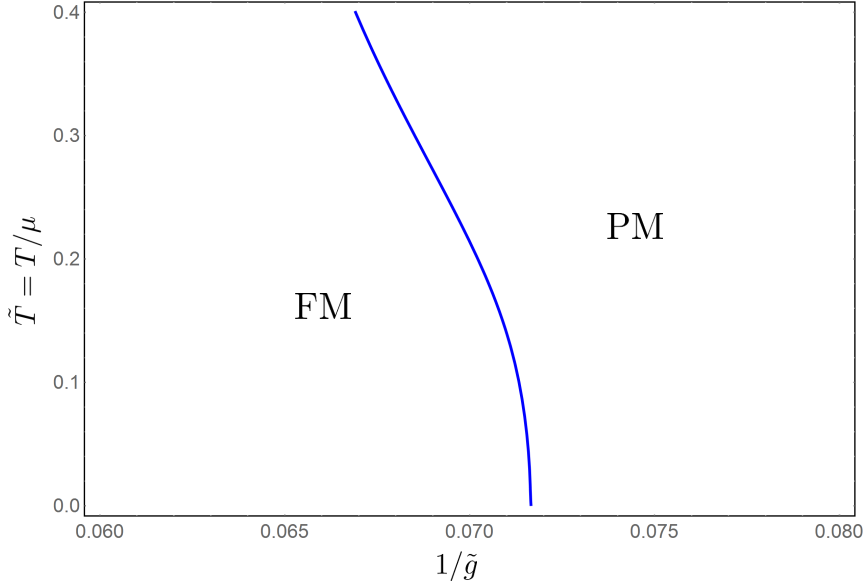
Figure 2.1 *The second-order ferromagnet-to-paramagnet transition as determined by Stoner mean-field theory, as a function of dimensionless temperature $\tilde{T} = T/\mu$ and interaction strength $\tilde{g} = g/\mu$.*

surface. This is the conventional form of the Stoner criterion, linking the onset of the ferromagnetic phase to the relation between the on-site repulsion and the density of states.

2.3 Hertz-Millis Theory

Having examined the mean-field model, now we can turn to a field theoretic approach. The conventional theory of ferromagnetic quantum criticality is the Hertz-Millis (also known as Hertz-Millis-Moriya) theory [80–82]. The philosophy behind this approach is to begin with a fermionic Hamiltonian and integrate out the fermionic degrees of freedom, replacing them instead with a bosonic field that serves as the order parameter. This then leads to an action which is analytic in powers of the bosonic field.

Starting from the fermionic Hubbard Hamiltonian in Eq. 2.1, we can express the partition function as a path integral over the fermionic Grassmann fields:

$$\mathcal{Z} = \int \mathcal{D}[\bar{\psi}, \psi] \mathrm{e}^{-S[\bar{\psi}, \psi]}, \tag{2.8}$$

$$S[\bar{\psi}, \psi] = \int_0^\beta \mathrm{d}\tau \int \mathrm{d}^3\mathbf{r} \left(\bar{\psi} \partial_\tau \psi + \mathcal{H}(\bar{\psi}, \psi) \right). \tag{2.9}$$

Following the original work of Hertz [80], we perform a Hubbard-Stratonovich [88, 89] decoupling in the spin channel (which we denote ϕ) and rewrite the

partition function as:

$$\mathcal{Z} = \int \mathcal{D}[\bar{\psi}, \psi]\mathcal{D}[\boldsymbol{\phi}]\mathrm{e}^{-S[\bar{\psi},\psi,\boldsymbol{\phi}]}, \qquad (2.10)$$

$$S[\bar{\psi}, \psi, \phi^*, \phi] = \int_{\mathbf{k},\omega} \bar{\psi}(G^{-1} - g\boldsymbol{\phi}\cdot\boldsymbol{\sigma})\psi + g\int_{\mathbf{r},\tau} \phi(\mathbf{r},\tau)^2, \qquad (2.11)$$

where $G = (k^2 - i\omega + \mu)^{-1}$ is the free electron Green's function, and $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is the vector of Pauli matrices. We can then integrate out the fermionic Grassmann fields to leave an effective field theory written solely in terms of fluctuations of the bosonic order parameter $\boldsymbol{\phi}$. This is the Hertz-Millis effective action:

$$\mathcal{Z} = \mathcal{D}[\boldsymbol{\phi}]\mathrm{e}^{-S[\boldsymbol{\phi}]}, \qquad (2.12)$$

$$S[\boldsymbol{\phi}] = -\mathrm{Tr}\ln\left[G^{-1} - g\boldsymbol{\phi}\cdot\boldsymbol{\sigma}\right] + g\int_{\mathbf{r},\tau} \phi(\mathbf{r},\tau)^2. \qquad (2.13)$$

This expression can be evaluated by expanding the Tr ln term to quartic order in $\boldsymbol{\phi}$. This expansion around $\boldsymbol{\phi} = 0$ restricts the resulting action to describing fluctuations of the bosonic order parameter in the paramagnetic phase only and it not valid in the magnetically ordered phase. The resulting action after this expansion has the form:

$$S = \sum_{\mathbf{q},\omega_n} \left(r + q^2 + \frac{|\omega_n|}{\gamma(q)}\right) \boldsymbol{\phi}_{\mathbf{q},\omega_n}\boldsymbol{\phi}_{-\mathbf{q},-\omega_n} + u\int_{\mathbf{r},\tau} \left[\phi(\mathbf{r},\tau)^2\right]^2 + ... \qquad (2.14)$$

where r and u can be defined in terms of the single particle Green's function and the density of states respectively [90].

The dynamic contribution $|\omega_n|/\gamma(q)$ is known as a Landau damping term and arises from the coupling of statics and dynamics at a quantum phase transition. It takes into account the excitation of particle-hole pairs across the Fermi level which lead to a damping of the spin fluctuations $\boldsymbol{\phi}_{\mathbf{q},\omega_n}$. The precise form of the damping term shown here relies on the assumption that the electronic quasiparticles obey Fermi liquid behaviour [90], taking the form $\gamma(q) = v_F q$ in a ferromagnet (in the limit $|\omega| \ll v_F q$ where v_F is the Fermi velocity) and $\gamma_q \sim \gamma_0$ in an antiferromagnet, independent of q. More sophisticated calculations which go beyond the Hertz-Millis and Fermi liquid paradigm predict quantitatively different forms of Landau damping [91].

Following, for example, Ref. [82], the quantum critical behaviour of the itinerant ferromagnet can be studied using perturbative renormalisation group techniques. As we have already discussed, however, the resulting analysis does not fully describe many materials. Having sketched Hertz-Millis theory, now we are in a position to explain its shortcomings.

2.4 Breakdown of Hertz-Millis Theory

Though successful at explaining many thermodynamic properties, Hertz-Millis theory does not capture the behaviour of certain materials close to quantum critical points. In particular, many materials exhibit first-order transitions at low temperatures [15–19], which Hertz-Millis theory cannot explain. Nor can it account for the formation of new phases around quantum critical points seen in experiment and predicted by various other theoretical methods. The avoidance of quantum critical points has been suggested to be a general principle [92].

The many successes of Hertz-Millis theory [90] suggest that it isn't irreparably incorrect, simply incomplete. The hope is that by rethinking the approximations and assumptions used, it may be possible to repair the theory and correctly describe the properties of ferromagnetic quantum criticality.

The main reason for the breakdown of Hertz-Millis theory is now known to be because it does not take into account the so-called 'soft modes' [93], low energy particle-hole fluctuations that couple to the order parameter. The resulting non-analytic correction to the free energy of the ferromagnet was first derived in [94] and takes the form $\mathcal{F}_{fl} = M^4 \ln(M^2 + T^2)$ in three spatial dimensions, where M is the magnetisation and T the temperature. This term contributes a $\ln T$ divergence to the quartic coefficient of the Landau free energy $\mathcal{F}_{fl}^{(4)} \approx M^4 \ln T$, guaranteeing that at a sufficiently low temperature the quartic term will turn negative and the transition will become first-order. This suggests that fluctuation-induced first-order behaviour is a generic feature of ferromagnetic quantum criticality, with the microscopics of the material in question governing what constitutes a 'sufficiently low' temperature[1].

Furthermore, inclusion of the non-analyticity opens up the possibility for incommensurate ordering as well as fluctuation-induced first-order behaviour. This has been suggested previously by numerous authors using a series of different frameworks [93, 95–102]. For our purposes in this thesis, the quantum order-by-disorder method provides a clear and intuitive picture as to why Fermi surface reconstruction is permitted near ferromagnetic quantum critical points (FM QCPs) and in turn why this leads to incommensurate ordering, so it is to this method that we turn to explain why incommensurate phases exist near FM QCPs.

---

[1]It has been pointed out [91] that should non-Fermi liquid properties of the normal state become significant before this temperature is reached, the first-order transition may be avoided.

2.5   The Order-by-Disorder Mechanism

First proposed in the context of condensed matter physics in 2009 [83], the fermionic quantum order-by-disorder mechanism has its roots in the Coleman-Weinberg theory of high-energy physics [103]. The central idea of the condensed matter form of this technique is that certain Fermi surface deformations can open up more phase space for low-energy particle-hole fluctuations which can self-consistently act to lower the energy of the system and stabilise new ground states. Order-by-disorder is essentially equivalent to self-consistent second-order perturbation theory [104], but this framework is a particularly accessible and transparent exposition of the physics at hand.

The concept of allowing the system some additional degree of freedom which it can use to lower its energy is not a new one. It is similar to the idea of entropic lowering of free energy at non-zero temperatures, where $\mathcal{F} = E - TS$ can be minimised by picking a state of high entropy.

In a quantum condensed matter context, an illustrative example of the order-by-disorder mechanism can be found in the insulating quantum antiferromagnet. If all the electron spins were oriented ferromagnetically, virtual hopping between sites would be prohibited by the Pauli exclusion principle. Confinement to a single site comes with an energy cost, so this configuration of spins is not ideal.

If instead the electron spins are oriented antiferromagnetically, Pauli exclusion no longer applies. Virtual hopping is allowed and the kinetic energy benefit gained by relaxing the confinement of the electrons self-consistently lowers the free energy of the antiferromagnet below that of the ferromagnet. In other words, quantum fluctuations associated with an additional degree of freedom can stabilise antiferromagnetic order.

In the case of fermionic quantum order-by-disorder, the fluctuation corrections correspond to deformations of the Fermi surfaces for spin-up and spin-down, as shown in Fig. 2.2. The excitations happen at or near the Fermi level, i.e. on the edges of the Fermi surfaces shown in the figure. By deforming the Fermi surfaces away from equal-sized circles (shaded), the system increases the surface-area-to-volume ratio and the increased total circumference allows for a larger number of fluctuations. If the fluctuations lower the free energy, the system will pick a deformation that maximises the allowed number of fluctuations. These deformations correspond to different phases, e.g. if the spin-up Fermi surface becomes smaller than the spin-down Fermi surface, the system will become magnetic, or if an anisotropic deformation is favoured the system could exhibit spiral magnetism or electron nematic order, for example.

The key idea behind the fermionic order-by-disorder mechanism is that one self-consistently calculates the free energy in the presence of a given type of order (ferromagnetic, spiral etc.) and checks whether the fluctuations can stabilise that phase over other competing phases considered. Order-by-disorder cannot *a*
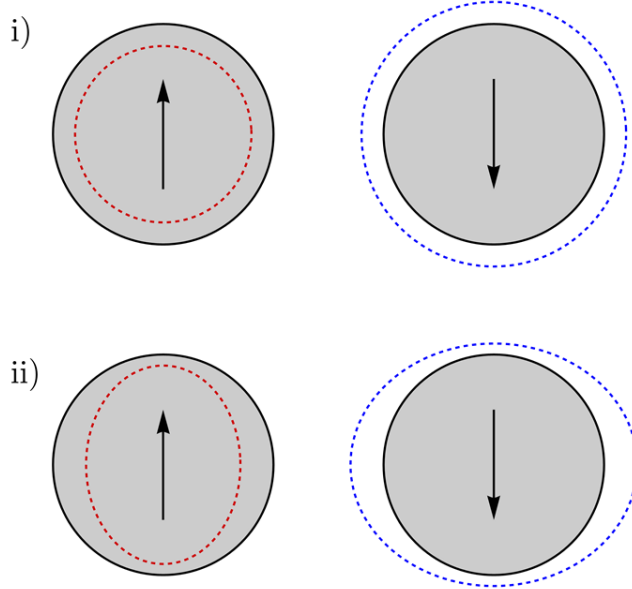
Figure 2.2   *A two dimensional cartoon of Fermi surface deformations corresponding to i) uniform ferromagnetism and ii) spiral magnetism. The filled circles correspond to a paramagnetic phase, i.e. equal numbers of spin-up and spin-down particles which lead to no net magnetic moment. In both cases, the Fermi surface deformations lead to an increase the surface area available for low-energy particle-hole fluctuations. This increase is relatively modest in the case of ferromagnetism, but larger in the case of the spiral magnet. Other deformations such as spin nematic phases are possible but are not considered in this thesis.*

priori determine the ground state - one must first guess a likely type of order, then calculate the fluctuation corrections and check whether it is stabilised or not. This limitation means that constant feedback is required from experiments in order to guide the search for new phases in the vicinity of quantum critical points.

2.6   Fluctuation Corrections

Full details on the calculation of the fluctuation corrections using the order-by-disorder framework can be found in Refs. [83, 104–107]. I first sketch the main steps here as they have been presented in previous work before deriving the clean-case fluctuation corrections in full along with the disorder-induced fluctuations in Section 2.8.

Derivation of the fluctuation corrections

Starting from the fermionic partition function defined in Eq. 2.8 and the associated action, we can perform a Hubbard-Stratonovich decoupling of the interaction term in both spin and charge channels (denoted $\boldsymbol{\phi}$ and ρ respectively) to obtain the action:

$$\mathcal{Z} = \int \mathcal{D}[\bar{\psi}, \psi]\mathcal{D}[\boldsymbol{\phi}]\mathcal{D}[\rho]\mathrm{e}^{-S[\bar{\psi},\psi,\boldsymbol{\phi},\rho]}, \tag{2.15}$$

$$S[\bar{\psi}, \psi, \boldsymbol{\phi}, \rho] = \int_{\mathbf{k},\omega} \bar{\psi} \left[G^{-1} + g(\rho - \boldsymbol{\phi} \cdot \boldsymbol{\sigma}) \right] \psi + g \int_{\mathbf{r},\tau} (\phi^2 - \rho^2). \tag{2.16}$$

where as before $G = (-i\omega - k^2 + \mu)^{-1}$ is the free electron Green's function, and $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is the vector of Pauli matrices. Decoupling only in the spin channel as in Hertz-Millis theory results in an incomplete description of the fluctuations present in the system, as both ρ and $\boldsymbol{\phi}$ appear here on the same footing. The fermionic degrees of freedom can be integrated out exactly as in the Hertz-Millis case, leading to the action:

$$\mathcal{Z} = \int \mathcal{D}[\boldsymbol{\phi}]\mathcal{D}[\rho]\mathrm{e}^{-S[\boldsymbol{\phi},\rho]}, \tag{2.17}$$

$$S[\phi] = -Tr\ln\left[G^{-1} + g(\rho - \boldsymbol{\phi} \cdot \boldsymbol{\sigma}) \right] + g \int_{\mathbf{r},\tau} (\phi^2 - \rho^2). \tag{2.18}$$

In the case of Hertz-Millis, the expansion of the free energy was performed around the paramagnetic state. In contrast, here we wish to perform an expansion in terms of the fluctuations $\tilde{\boldsymbol{\phi}}$ around the static magnetic order parameter $\mathbf{M}$. With this aim in mind, we make the first major deviation from Hertz-Millis and decompose the bosonic fields into static (zero frequency) and fluctuation (finite frequency) components:

$$\rho(\mathbf{r}, \tau) = \rho_0(\mathbf{r}) + \tilde{\rho}(\mathbf{r}, \tau), \tag{2.19}$$

$$\boldsymbol{\phi}(\mathbf{r}, \tau) = \mathbf{M}(\mathbf{r}) + \tilde{\boldsymbol{\phi}}(\mathbf{r}, \tau), \tag{2.20}$$

leading to an action

$$S[\phi] = -Tr\ln\left[G_{\mathbf{M}}^{-1} + g(\tilde{\rho} - \tilde{\boldsymbol{\phi}} \cdot \boldsymbol{\sigma}) \right] + g \int_{\mathbf{r},\tau} (\mathbf{M}^2 + \tilde{\phi}^2 - \rho_0^2 - \tilde{\rho}^2), \tag{2.21}$$

where $G_{\mathbf{M}}^{-1} = G^{-1} + g(\rho_0 - \mathbf{M} \cdot \boldsymbol{\sigma})$ is the inverse Green's function in the presence of magnetic order. In the following, we do not consider the possibility of charge ordering, so we set $\rho_0 = const.$ and absorb it into the chemical potential.

Expanding the action to quadratic order in finite-frequency fluctuations reproduces the mean-field contribution from Eq. 2.3 plus an additional fluctuation-correction term. After a lengthy calculation (shown in Section 2.8), the

fluctuation correction can be shown to be:

$$\mathcal{F}^{\mathrm{fl}} = -2g^2 \int_{\mathbf{k}_1 \ldots \mathbf{k}_4}^{(*)} \frac{f_+(\mathbf{k}_1) f_-(\mathbf{k}_2)(f_+(\mathbf{k}_3) + f_-(\mathbf{k}_4))}{\varepsilon_+(\mathbf{k}_1) + \varepsilon_-(\mathbf{k}_2) - \varepsilon_+(\mathbf{k}_3) - \varepsilon_-(\mathbf{k}_4)},$$

(2.22)

where f is the Fermi-Dirac distribution function with $f_\sigma(\mathbf{k}) = f(\varepsilon_\sigma(\mathbf{k}))$ for brevity and the asterisk on the integral indicates momentum conservation, $\mathbf{k}_1 - \mathbf{k}_2 = \mathbf{k}_3 - \mathbf{k}_4$. It is possible to see from this term the particle-hole excitation structure discussed earlier - the Fermi functions that enter here in the numerator create particle-hole pairs with opposite spin and equal-and-opposite momenta (enforced by the momentum conservation).

Corrections to the Ginzburg-Landau coefficients

In the case of the uniform ferromagnet, the fluctuation corrections to the quadratic and quartic coefficients can be calculated from Eq. 2.22 in the low-temperature limit and found in Refs. [104, 106, 107]. They are given by:

$$a_{2,fl} \approx -\frac{16\sqrt{2}}{3(2\pi)^6}(1 + 2\ln 2)g^4,$$

(2.23)

$$a_{4,fl} \approx \frac{16\sqrt{2}}{3(2\pi)^6}(1 + \ln \tilde{T})g^6,$$

(2.24)

where $\tilde{T} = T/\mu$ is the dimensionless temperature. The $\ln \tilde{T}$ contribution to the quartic coefficient guarantees the presence of a first-order transition at a low enough temperature.

Taking these fluctuation corrections into account leads to the phase diagram shown in Fig. 2.3 in the $(1/\tilde{g}, \tilde{T})$ plane, where $\tilde{g} = g\rho_F$ and ρ_F is the density of states at the Fermi surface. We find that the transition turns first order below a tricritical point located at approximately $\tilde{T}_{tcp} \approx 0.33$. However, the slope of the phase boundary below the tricritical point is unphysical. While it is not unreasonable to expect the fluctuations to stabilise ferromagnetic order in a region which was previously paramagnetic, the phase boundary should still terminate at a finite value of $1/\tilde{g}$. To obtain physical behaviour, we need to include terms beyond the leading fluctuation corrections.

Resummation of higher-order terms

The previous approach has relied upon expanding the free energy in terms of the order parameter M, however this is only valid when M is small, i.e. in the vicinity of the tricritical point. At lower temperatures, the subleading terms diverge more strongly and the validity of our expansion breaks down, leading to
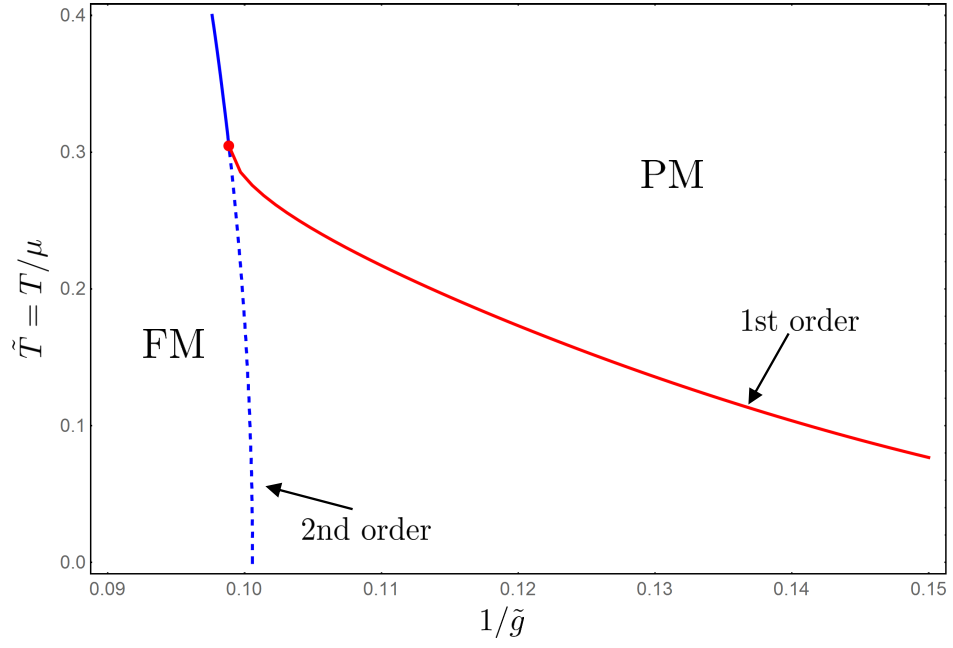
Figure 2.3 *After incorporating the leading fluctuation corrections, the transition turns first-order below a tricritical point but the phase boundary displays unphysical behaviour at low temperature.*

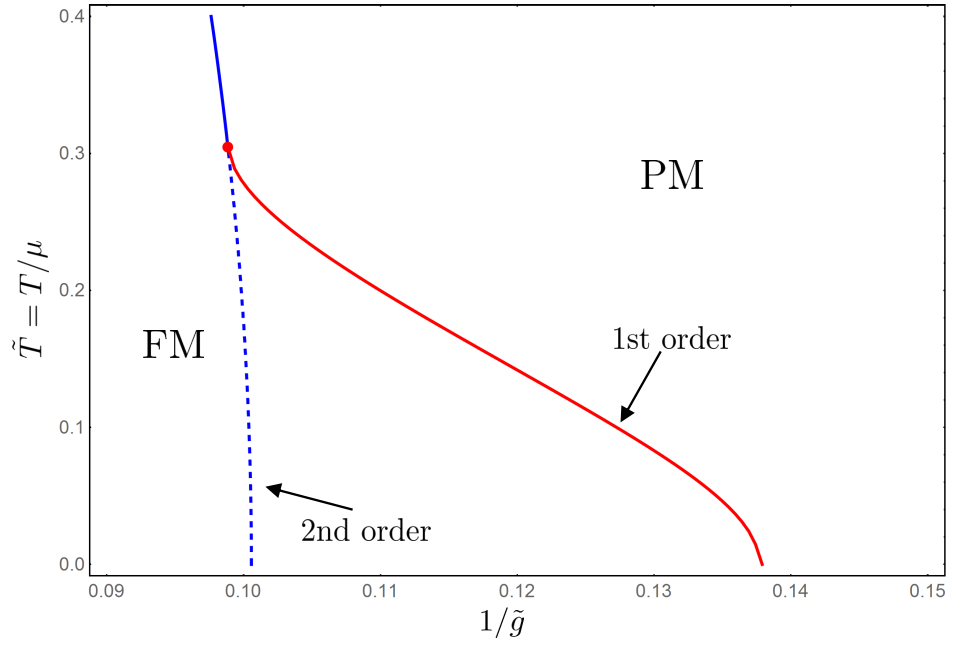
Figure 2.4 *The first-order phase boundary after including the resummation of higher-order divergences. The unphysical asymptotic behaviour is repaired and the transition now terminates at $\tilde{T} = 0$ at a finite value of $1/\tilde{g}$.*

the unphysical asymptotic behaviour in the $\tilde{T} \to 0$ limit shown in Fig. 2.3. Higher order terms in the expansion must be retained in order to accurately compute the low temperature properties and ultimately to get the true $\tilde{T} \to 0$ behaviour we must retain all orders of M.

The result from Ref. [107] can be used instead to resum the leading divergences, generating a resummed fluctuation correction of the form:

$$\mathcal{F}_{fl} = \frac{1}{2}\lambda \tilde{g}^2 m^4 \ln\left(\kappa m^2 + \tilde{T}^2\right), \tag{2.25}$$

where $\lambda = 16\sqrt{2}/3(2\pi)^6$, $m = Mg/(\rho_F \mu)$ and κ is some constant which controls where the first-order transition terminates on the $1/\tilde{g}$ axis at $\tilde{T} = 0$. This term arises only because it is not possible to resum all of the divergences exactly except at $T = 0$. This form of resummed fluctuation correction was suggested in Ref. [107] as a suitable function that interpolates between the leading corrections near the tricritical point but gives the same $\tilde{T} = 0$ intercept as the full resummation of all divergences. Following the value used in Ref. [79] I will set $\kappa = 0.001$ in the remainder of this chapter.

The results of including this resummation are shown in Fig. 2.4 where we see a more physically realistic phase boundary with the first-order transition terminating on the $\tilde{T} = 0$ axis at finite interaction strength $\tilde{g}$.

Thus far we have only allowed for the possibility of uniform ferromagnetism and we have seen that the fluctuations both stabilise ferromagnetism over a larger region of the phase diagram and change the nature of the transition at low temperatures. If fluctuations can change the order of the phase transition by lowering the ground state energy, it stands to reason that fluctuations could go further and stabilise new phases that are not favoured at mean-field level.

2.7   Fluctuation-Induced Spiral Magnet

Motivated by the suggestion that incommensurately ordered spiral ferromagnetic phases can be stabilised by fluctuations in the vicinity of ferromagnetic quantum critical points [83, 97, 104, 106], we can repeat the analysis of the previous section by allowing for a more general magnetisation

$$\mathbf{m}(\mathbf{r}) = m[\hat{x}\cos(\mathbf{Q} \cdot \mathbf{r}) + \hat{y}\sin(\mathbf{Q} \cdot \mathbf{r})], \tag{2.26}$$

where without loss of generality we assume $\mathbf{Q} = Q\hat{z}$, where $\hat{\alpha}$ is the unit vector in the $\alpha = \{x, y, z\}$ direction, and we continue to express the magnetisation in dimensionless units $m = Mg/(\rho_F \mu)$. The free energy can now be expanded around the tricritical point in both m and Q:

$$\mathcal{F} = (a_2 + b_2 Q^2 + c_2 Q^4)m^2 + (a_4 + b_4 Q^2)m^4 + \dots \tag{2.27}$$

Similarly to the case of the homogeneous ferromagnet, analysis of the coefficients a_{2n}, b_{2n} and c_{2n} then allows us determine when (and if) the system favours a modulated magnetic state, characterised by $m \neq 0$ and $Q \neq 0$. In the simplest instance, in order for a non-zero Q to be favoured we require $b_2 < 0$.

Free energy

At mean-field level we again find that $\mathcal{F} = gm^2 - T \sum_\sigma \int_k \ln(\mathrm{e}^{-\beta(\varepsilon_\sigma(\mathbf{k})-\mu)} + 1)$ but with the key difference that $\varepsilon_\sigma(\mathbf{k})$ is now the free fermion dispersion in the presence of spiral order, given by:

$$\varepsilon_\sigma(\mathbf{k}) = k^2 + \sigma\sqrt{(k_z q)^2 + m^2}, \tag{2.28}$$

where $q = Q/k_F$. We can derive expressions for the coefficients in Eq. 2.27 by Taylor expanding the free energy in two variables. The coefficient of the first mixed term involving both m and q becomes:

$$b_2 = \frac{g^4}{3!} \int_\mathbf{k} \left(\frac{\mathbf{k}\cdot\mathbf{q}}{k_F q}\right)^2 f^{(3)}(\varepsilon(k)). \tag{2.29}$$

This coefficient contains the third derivative of the Fermi function, similarly to a_4, the the coefficient of m^4. This suggests already that perhaps there is some relation between the coefficient of $m^2 q^2$ and the bare m^4 coefficient from the homogeneous case. To see this further we need to evaluate the $\mathbf{k}$- and $\mathbf{q}$-dependent prefactor. This can be dealt with in the low temperature approximation, where for temperatures $\tilde{T} \ll \mu$, the derivatives of the Fermi function are dominated by contributions close to the Fermi energy, e.g. the first derivative is a narrow peak centred at the Fermi energy. We may then approximate $|\mathbf{k}| = k_F$, leading to:

$$\frac{\mathbf{k}\cdot\mathbf{q}}{k_F q} \approx \cos\theta, \tag{2.30}$$

where θ is just the relative angle between $\mathbf{k}$ and $\mathbf{q}$ which can be straightforwardly integrated over. By writing the coefficients in this manner and expressing the free energy as an integral over the solid angle Ω in three spatial dimensions, e.g. $F(m, q) = \langle F(\sqrt{k_z^2 q^2 + m^2})\rangle_\Omega$, we find that the coefficients of m^2 and q^2 are indeed related to one another. In general, all n^{th} order coefficients are proportional.

In particular, this means that the $m^2 q^2$ coefficient is proportional to the m^4 coefficient. Since it is the m^4 coefficient which has the logarithmic divergence at low temperature, the $m^2 q^2$ term also inherits this same $\ln \tilde{T}$ divergence. This implies that rather than the ferromagnetic transition simply turning first-order at the tricritical point, instead a spiral phase becomes favoured over the homogeneous ferromagnet. Specifically, Refs. [104, 106] show that the proportionalities between the coefficients are $b_2 \approx \frac{2}{3}a_4$, $c_2 \approx \frac{3}{5}a_6$ and $b_4 \approx a_6$.

43

Fluctuation corrections

Fluctuations can be calculated self-consistently in the presence of spiral order by following the same procedure as for the homogeneous ferromagnet but now allowing for the order parameter shown in Eq. 2.26. The process again leads to a correction to the free energy of the same form as Eq. 2.22 with the important difference that the dispersion relation appearing in all terms under the momentum integral $\varepsilon_\sigma(\mathbf{k}) = k^2 + \sigma\sqrt{(k_z q)^2 + m^2}$ is now the dispersion in the presence of spiral order.

As in the case of the homogeneous ferromagnet, the fluctuation corrections must be resummed to obtain physical behaviour in the $\tilde{T} \to 0$ limit. To do this, we use the same functional form as of the resummed free energy correction in the homogeneous case, where $\mathbf{m}$ is now the spiral order parameter, and expand the full free energy in powers of q:

$$\mathcal{F}(m, q) = f_0(m) + f_2(m)q^2 + \frac{1}{2}f_4(m)q^4 + ... \tag{2.31}$$

where, after writing the free energy as an angular average over k_z using $\langle k_z^{2n} \rangle = 1/(2n + 1)$, we can write the coefficients as:

$$f_0(m) = [\mathcal{F}_{mf} + \mathcal{F}_{fl}]_{q=0}, \tag{2.32}$$

$$f_2(m) = \frac{2}{3}a_4 m^2 + a_6 m^4 + \frac{1}{6}\lambda\tilde{g}^2 m^2 \left[\frac{\kappa m^2}{\kappa m^2 + \tilde{T}^2} + 2\ln(\kappa m^2 + \tilde{T}^2) \right], \tag{2.33}$$

$$f_4(m) = \frac{6}{5}a_6 m^2 + \frac{1}{5}\lambda\tilde{g}^2 \left[-\frac{1}{2}\frac{\kappa^2 m^4}{(\kappa m^2 + \tilde{T}^2)^2} + 2\frac{\kappa m^2}{\kappa m^2 + \tilde{T}^2} + \ln\left(\frac{\kappa m^2 + \tilde{T}^2}{\tilde{T}^2}\right) \right], \tag{2.34}$$

where $f_0(m)$ is the resummed, fluctuation-corrected free energy expression for the $q = 0$ homogeneous ferromagnet [79]. As before, $\kappa = 0.001$ is a constant which controls the location of the $\tilde{T} \to 0$ termination of the first-order phase transition.

Minimising $\mathcal{F}$ with respect to q leads us to find that the optimum value of the spiral ordering wavevector is given by $q^2 = -f_2/f_4$. Plugging this back into Eq. 2.31 gives an expression which need be minimised only in terms of m:

$$\mathcal{F}(m) = f_0(m) - \frac{1}{2}f_2(m)^2/f_4(m^2) + ... \tag{2.35}$$

We can compare this finite-q free energy with the free energy of the homogeneous ferromagnet to see which state the system favours. From this, we can construct the phase diagrams shown in Fig. 2.5.

We find that the first-order transition from the paramagnet to the homogeneous ferromagnet is pre-empted by a transition into the spiral ferromagnetic phase. Quantum fluctuations have stabilised the incommensurate spiral over a region of the phase diagram where it did not exist at mean-field level. The
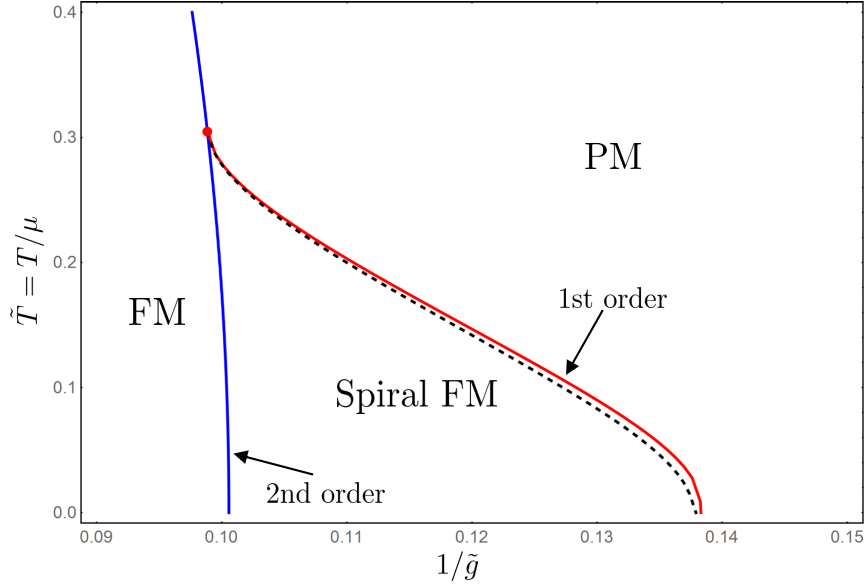
44

Figure 2.5 *The resulting phase diagram as a function of dimensionless temperature $\tilde{T} = T/\mu$ and inverse interaction strength $1/\tilde{g} = 1/(g\rho_F)$ showing the fluctuation-induced spiral magnetic order emerging near a ferromagnetic quantum critical point. The black dashed line is the fluctuation-induced first-order ferromagnetic transition, which is pre-empted here by the transition into the modulated spiral state.*

transition from the spiral to the homogeneous ferromagnet is second-order and occurs at the $a_2 = 0$ line, however the anisotropies present in any real system will render this transition weakly first-order. In addition to anisotropies, any real material will also have some form of randomness or disorder present, and that is where my work began.

2.8   Disorder-induced fluctuations

My original contribution to the fermionic order-by-disorder mechanism was to investigate how disorder couples to the quantum fluctuations in the vicinity of the quantum critical point. In the following, we specify to quenched charge disorder. This is a convenient type of disorder for such a study as at mean-field level it does not change the nature of either the paramagnetic or ferromagnetic phases, and as such the only effects it can have on the nature of the phases present is through disorder-induced fluctuations.

Replica Field Theory

We begin with the fermionic Hubbard Hamiltonian as before, however this time we write it entirely in real-space such that we can include the disorder like so:

$$\mathcal{H} = \int \mathrm{d}^3\mathbf{r} \left[\sum_{\nu=\pm} \psi_\nu^\dagger \left[\nabla^2 - \mu + v_c(\mathbf{r}) \right] \psi_\nu + g\psi_+^\dagger \psi_+ \psi_-^\dagger \psi_- \right], \qquad (2.36)$$

where the charge disorder $v_c(\mathbf{r})$ is local, uncorrelated between sites and is Gaussian distributed with zero mean and with a variance σ_c^2.

$$\overline{v_c(\mathbf{r})} = 0, \qquad (2.37)$$

$$\overline{v_c(\mathbf{r})v_c(\mathbf{r}')} = \sigma_c^2\,\delta(\mathbf{r} - \mathbf{r}'). \qquad (2.38)$$

We can write an action as a path integral over fermionic coherent states in the usual manner, leading to:

$$S_{\bar{\psi},\psi} = \int_0^\beta \mathrm{d}\tau \int \mathrm{d}^3\mathbf{r} \left\{ \sum_{\nu=\pm} \bar{\psi}_\nu \left(\partial_\tau - \nabla^2 - \mu + v_c(\mathbf{r}) \right) \psi_\nu + g\bar{\psi}_+\psi_+\bar{\psi}_-\psi_- \right\}, \qquad (2.39)$$

We again employ a Hubbard-Stratonovich transformation to decouple the interaction and introduce two bosonic auxiliary fields $\boldsymbol{\phi}$ and ρ representing spin and charge respectively:

$$S_{\bar{\psi},\psi,\rho,\boldsymbol{\phi}} = \int_0^\beta \mathrm{d}\tau \int \mathrm{d}^3\mathbf{r} \; \big[g(\boldsymbol{\phi}^2 - \rho^2)$$
$$+ \bar{\psi} \left\{ \partial_\tau - \nabla^2 - \mu + v_c(\mathbf{r}) + g(\rho - \boldsymbol{\phi} \cdot \boldsymbol{\sigma}) \right\} \psi \big], \qquad (2.40)$$

where $\bar{\psi} = (\psi_+, \psi_-)^T$ and $\psi = (\psi_+, \psi_-)$. To average over the random disorder, we use the replica trick (see Appendix B):

$$\overline{\ln \mathcal{Z}} = \lim_{n \to 0} \frac{\overline{\mathcal{Z}^n} - 1}{n}. \qquad (2.41)$$

After application of the replica trick, we arrive at the replica action with disorder-averaged coefficients:

$$S_n = \sum_{\alpha=1}^n \int \mathrm{d}\tau \int \mathrm{d}^3\mathbf{r} \left\{ g\left[(\boldsymbol{\phi}^\alpha)^2 - (\rho^\alpha)^2 \right] \right.$$
$$+ \bar{\psi}^\alpha(\mathbf{r},\tau) \left[\partial_\tau - \nabla^2 - \mu + g(\rho^\alpha - \boldsymbol{\phi}^\alpha \cdot \boldsymbol{\sigma}) \right] \psi^\alpha(\mathbf{r},\tau) \Big\}$$
$$- \frac{\sigma_c^2}{2} \sum_{\alpha,\beta} \sum_{\nu,\nu'=\pm1} \int_{\tau,\tau',\mathbf{r}} \bar{\psi}_\nu^\alpha(\mathbf{r},\tau)\psi_\nu^\alpha(\mathbf{r},\tau)\bar{\psi}_{\nu'}^\beta(\mathbf{r},\tau')\psi_{\nu'}^\beta(\mathbf{r},\tau'), \qquad (2.42)$$

where α and β label the replicas. The first part of this equation is just the clean action replicated n times, but the disorder-averaging procedure has generated an additional disorder vertex with a coefficient σ_c^2 that is just the variance of the disorder distribution. This vertex is non-diagonal in both replica index and imaginary time. We may then proceed to decompose the fields into static and fluctuation terms exactly as before, with the same spiral magnetic order parameter:

$$\mathbf{m}(\mathbf{r}) = m[\mathbf{n}_x \cos(\mathbf{q} \cdot \mathbf{r}) + \mathbf{n}_y \sin(\mathbf{q} \cdot \mathbf{r})], \qquad (2.43)$$

where $\mathbf{q} = q\mathbf{n}_z$. Again, we include the dispersion of the spiral magnet in the free-fermion action in order to self-consistently calculate the fluctuation corrections in the presence of spiral order.

Integration over fermionic degrees of freedom

The free-fermion action may be diagonalised by a unitary transformation:

$$\begin{pmatrix} \varphi_+^\alpha(\mathbf{k}, \omega_n) \\ \varphi_-^\alpha(\mathbf{k}, \omega_n) \end{pmatrix} = e^{-i\frac{\theta_{\mathbf{k}}}{2}\sigma_y} \begin{pmatrix} \psi_+^\alpha(\mathbf{k}+\mathbf{q}/2, \omega_n) \\ \psi_-^\alpha(\mathbf{k}-\mathbf{q}/2, \omega_n) \end{pmatrix}, \qquad (2.44)$$

with $\tan\theta_{\mathbf{k}} = gM/(k_z Q)$. This corresponds to transforming the fermionic fields into the rotating spiral frame.

The fluctuation-independent part of the resulting action, from the transformation of the $\bar{\psi}^\alpha(\mathbf{r}, \tau)\left[\partial_\tau - \nabla^2 - \mu + g(\rho^\alpha - \boldsymbol{\phi}^\alpha \cdot \boldsymbol{\sigma})\right]\psi^\alpha(\mathbf{r}, \tau)$ term in Eq. 2.42, can be written as:

$$S_0[\bar{\varphi}, \varphi, \mathbf{m}] = \sum_{\omega_n} \int_{\mathbf{k}} \sum_{\alpha,\nu} G_\nu^{-1}(\mathbf{k}, \omega_n)\bar{\varphi}_\nu^\alpha(\mathbf{k}, \omega_n)\varphi_\nu^\alpha(\mathbf{k}, \omega_n), \qquad (2.45)$$

where $\omega_n = (2n+1)\pi T$ denote fermionic Matsubara frequencies and $G_\nu(\mathbf{k}, \omega_n)$ is the fermionic Green's function in the presence of spiral order, given by:

$$G_\nu(\mathbf{k}, \omega_n) = \frac{1}{-i\omega_n + \varepsilon_\nu(\mathbf{k}) - \mu}, \qquad (2.46\text{a})$$

$$\varepsilon_\nu(\mathbf{k}) = k^2 + \nu\sqrt{k_z^2 q^2 + m^2}, \qquad (2.46\text{b})$$

where we have absorbed a q^2 term into the chemical potential. In the $\mathbf{q} \to 0$ limit, we recover the theory for the homogeneous ferromagnet, as we would expect.

The fluctuation terms resulting from transforming this same part of the action become:

$$\tilde{S} = g \sum_{\omega_n, \omega_n'} \int_{\mathbf{k},\mathbf{k}'} \sum_\alpha \bar{\varphi}^\alpha(\mathbf{k}, \omega_n) \left[\tilde{\rho}^\alpha(\Delta\mathbf{k}, \delta\omega_n) - \tilde{\boldsymbol{\phi}}^\alpha(\Delta\mathbf{k}, \delta\omega_n) \cdot \boldsymbol{\sigma}\right] \varphi^\alpha(\mathbf{k}', \omega_n'),$$
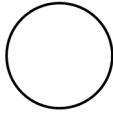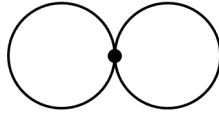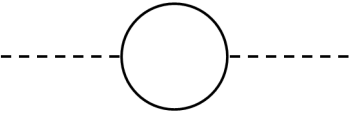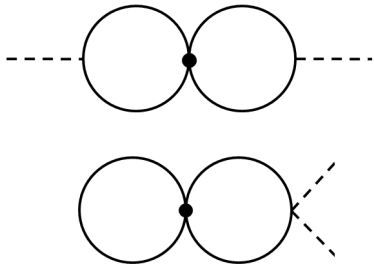
$$(2.47)$$

	Clean	Disordered
Mean-Field		
Fluctuations		

Figure 2.6 *The various diagrams representing integration over the fermionic fields (solid lines) to obtain an effective field theory in terms of the bosonic fields (dashed lines).*

where $\Delta\mathbf{k} = \mathbf{k} - \mathbf{k}'$ and $\delta\omega_n = \omega_n - \omega_n'$.

The disorder vertex $S_{dis} = -\frac{\sigma_c^2}{2} \sum_{\alpha,\beta} \sum_{\nu,\nu'=\pm1} \int_{\tau,\tau',\mathbf{r}} |\psi_\nu^\alpha(\mathbf{r},\tau)|^2 |\psi_{\nu'}^\beta(\mathbf{r},\tau')|^2$ remains most compactly expressed in terms of the initial variables, rather than transforming into the rotating frame. It is easier to leave it in the initial frame and later define a set of mixed correlation functions to allow us to compute the required expectation values.

We can now integrate out the fermionic fields to obtain an effective bosonic action which may be written as the sum of the mean-field and fluctuation contributions:

$$\tilde{S}_n = S_{\mathrm{mf}}[\mathbf{M}] + \tilde{S}_{\mathrm{fl}}[\mathbf{M}, \tilde\rho, \tilde\phi], \tag{2.48}$$

$$S_{\mathrm{mf}} = \beta n g M^2 - n \sum_{\omega_n} \int_{\mathbf{k}} \sum_\nu \ln G_\nu^{-1}(\mathbf{k}, \omega_n) + \left\langle S_{\mathrm{dis}}[\bar\varphi, \varphi] \right\rangle_{S_0[\bar\varphi,\varphi,\mathbf{M}]}, \tag{2.49}$$

$$\tilde{S}_{\mathrm{fl}}[\mathbf{M}, \tilde\rho, \tilde\phi] = \tilde{S}_0[\tilde\rho, \tilde\phi] - \frac{1}{2}\left\langle \tilde{S}^2[\bar\varphi, \varphi, \tilde\rho, \tilde\phi] \right\rangle_0 + \frac{1}{2}\left\langle \tilde{S}^2[\bar\varphi, \varphi, \tilde\rho, \tilde\phi] S_{\mathrm{dis}}[\bar\varphi, \varphi] \right\rangle_0^{(c)}, \tag{2.50}$$

where the superscript (c) signifies that only the connected diagrams contribute. Fig. 2.6 shows a schematic representation of the diagrams we are required to evaluate in integrating out the fluctuation fields.

48

Longitudinal and Transverse Fluctuations

To obtain the fluctuation corrections, we must integrate out the fluctuation fields $\phi(\mathbf{q}, \eta_\mathbf{n})$ and $\rho(\mathbf{q}, \eta_\mathbf{n})$ from the fluctuation-dependent part of the action, where η_n denotes a bosonic Matsubara frequency. We can split the correction terms into longitudinal and transverse fluctuations like so:

$$
\begin{aligned}
\tilde{S}_{\text{fl}} &= \tilde{S}_0[\tilde{\rho}, \tilde{\phi}] + \frac{g^2}{2} \sum_{\alpha \eta_n} \int_\mathbf{q} L^\perp(\mathbf{q}, \eta_n) \left(|\tilde{\phi}_x^\alpha(\mathbf{q}, \eta_n)|^2 + |\tilde{\phi}_y^\alpha(\mathbf{q}, \eta_n)|^2 \right) \\
&\quad + \frac{g^2}{2} \sum_{\alpha \eta_n} \int_\mathbf{q} \left(\tilde{\rho}^\alpha(\mathbf{q}, \eta_n), \tilde{\phi}_z^\alpha(\mathbf{q}, \eta_n) \right)^* \mathbf{L}^\parallel(\mathbf{q}, \eta_n) \left(\begin{array}{c} \tilde{\rho}^\alpha(\mathbf{q}, \eta_n) \\ \tilde{\phi}_z^\alpha(\mathbf{q}, \eta_n) \end{array} \right), \quad (2.51)
\end{aligned}
$$

where the longitudinal matrix elements, which account for charge fluctuations and longitudinal magnetic fluctuations, are

$$
L_{\rho\rho}^\parallel = \sum_\nu \left[\Pi_{\nu,\nu} + \Omega_{\nu,\nu,\nu,\nu} + \Omega_{\nu,\nu,-\nu,-\nu} + 2 \left(\Gamma_{\nu,\nu,\nu} + \Gamma_{\nu,-\nu,-\nu} \right) \right], \qquad (2.52)
$$

$$
L_{\rho\phi}^\parallel = -\sum_\nu \nu \left[\Pi_{\nu,\nu} + \Omega_{\nu,\nu,\nu,\nu} - \Omega_{\nu,\nu,-\nu,-\nu} + 2 \left(\Gamma_{\nu,\nu,\nu} - \Gamma_{\nu,-\nu,-\nu} \right) \right], \quad (2.53)
$$

$$
L_{\phi\rho}^\parallel = -\sum_\nu \nu \left[\Pi_{\nu,\nu} + \Omega_{\nu,\nu,\nu,\nu} + \Omega_{\nu,\nu,-\nu,-\nu} + 2 \left(\Gamma_{\nu,\nu,\nu} - \Gamma_{\nu,-\nu,-\nu} \right) \right], \quad (2.54)
$$

$$
L_{\phi\phi}^\parallel = \sum_\nu \left[\Pi_{\nu,\nu} + \Omega_{\nu,\nu,\nu,\nu} - \Omega_{\nu,\nu,-\nu,-\nu} + 2 \left(\Gamma_{\nu,\nu,\nu} + \Gamma_{\nu,-\nu,-\nu} \right) \right]. \qquad (2.55)
$$

and the transverse terms, which only describe transverse magnetic fluctuations, are

$$
L^\perp = \sum_\nu \left[\Pi_{\nu,-\nu} + \Omega_{\nu,-\nu,-\nu,\nu} + 2 \left(\Gamma_{\nu,\nu,-\nu} + \Gamma_{\nu,-\nu,\nu} \right) \right]. \qquad (2.56)
$$

The Π_{ν_1,ν_2} terms are Lindhard functions containing two Green functions, while the $\Gamma_{\nu_1,\nu_2,\nu_3}$ and $\Omega_{\nu_1,\nu_2,\nu_3,\nu_4}$ are analogous terms containing three and four Green functions respectively:

$$
\Pi_{\nu_1,\nu_2} = T \sum_{\omega_n} \int_\mathbf{k} G_{\nu_1}(\mathbf{k}, \omega_n) G_{\nu_2}(\mathbf{k} + \mathbf{q}, \omega_n + \eta_n), \qquad (2.57)
$$

$$
\Gamma_{\nu_1,\nu_2,\nu_3} = T \sum_{\omega_n} \int_{\mathbf{k},\mathbf{k}'} g_{\nu_1,\nu_2,\nu_3}(\mathbf{k}, \mathbf{k}') \qquad\qquad (2.58)
$$
$$
\times G_{\nu_1}(\mathbf{k}', \omega_n) G_{\nu_2}^2(\mathbf{k}, \omega_n) G_{\nu_3}(\mathbf{k} + \mathbf{q}, \omega_n + \eta_n),
$$
$$
\Omega_{\nu_1,\nu_2,\nu_3,\nu_4} = T \sum_{\omega_n} \int_{\mathbf{k},\mathbf{k}'} f_{\nu_1,\nu_2,\nu_3,\nu_4}(\mathbf{k}, \mathbf{k}', \mathbf{q})
$$
$$
\times G_{\nu_1}(\mathbf{k}, \omega_n) G_{\nu_2}(\mathbf{k} + \mathbf{q}, \omega_n + \eta_n)
$$
$$
\times G_{\nu_3}(\mathbf{k}' + \mathbf{q}, \omega_n + \eta_n) G_{\nu_4}(\mathbf{k}', \omega_n). \qquad (2.59)
$$

The **q**-dependent prefactors contained within them which describe the effect of disorder are:

$$f_{\nu,\nu,\nu,\nu} = \sigma_c^2 \left[1 - \frac{1}{4}(k_z - k_z')^2 \right] \frac{q^2}{m^2}, \tag{2.60}$$

$$f_{\nu,\nu,-\nu,-\nu} = \frac{\sigma_c^2}{4}(k_z - k_z')^2 \frac{q^2}{m^2}, \tag{2.61}$$

$$f_{\nu,-\nu,-\nu,\nu} = \sigma_c^2 \left[1 - \frac{1}{4}(k_z - k_z')^2 \right] \frac{q^2}{m^2}, \tag{2.62}$$

$$h_{\nu,-\nu}(\mathbf{k},\mathbf{k}') = 1 - h_{\nu,\nu}(\mathbf{k},\mathbf{k}') \approx \frac{1}{4}(k_z - k_z')^2 \frac{q^2}{m^2}, \tag{2.63}$$

$$g_{\nu,\nu,\nu'} = h_{\nu,\nu}, \tag{2.64}$$

$$g_{\nu,-\nu,\nu'} = h_{\nu,-\nu}. \tag{2.65}$$

The Π terms arise from the clean case fluctuations (the lower-left panel of Fig. 2.6) while the Ω and Γ terms arise from the disorder fluctuation integrals (the upper and lower loops respectively in the lower-left panel of Fig. 2.6). The coefficients $h_{\nu,-\nu}$ actually break the proportionalities between the coefficients of the Landau expansion discussed earlier in this chapter, however their contribution turns out to be negligible compared to that of $h_{\nu,\nu}$ and so the previous results linking the proportionality of various terms are assumed to still hold.

After the integration over the fluctuation fields, we reach the following expressions:

$$\mathcal{F}_{\text{clean}}^{\text{fl}} = \frac{T}{2} \sum_{\eta_n} \int_{\mathbf{q}} \left\{ g \left(\Pi_{+-} + \Pi_{-+} \right) + g^2 \left(\Pi_{+-}\Pi_{-+} - \Pi_{++}\Pi_{--} \right) \right\}, \tag{2.66}$$

$$\mathcal{F}_{\text{dis}}^{\text{fl}} = gT \sum_{\nu,\eta_n} \int_{\mathbf{q}} \left\{ \Omega_{\nu,-\nu,-\nu,\nu} - \Omega_{\nu,\nu,-\nu,-\nu} + 2 \left(\Gamma_{\nu,\nu,-\nu} + \Gamma_{\nu,-\nu,\nu} \right) \right\}, \tag{2.67}$$

which we can then evaluate by performing the summation of Matsubara frequencies to obtain the final fluctuation-corrected free energy.

2.9 Evaluating the Free Energy

Mean-field

The mean-field free energy requires no integration over fluctuation fields and after appropriate Matsubara frequency summation we recover the mean-field clean case result plus a mean-field disorder term:

$$\mathcal{F}_{\text{clean}}^{\text{mf}} = gM^2 - T \sum_{\nu} \int_{\mathbf{k}} \ln\left(e^{\beta(\varepsilon_\nu(\mathbf{k})-\mu)} + 1\right), \tag{2.68}$$

$$\mathcal{F}_{\text{dis}}^{\text{mf}} = \frac{\sigma_c^2}{2} \sum_{\nu,\nu'} \int_{\mathbf{k},\mathbf{k}'} h_{\nu,\nu'}(\mathbf{k},\mathbf{k}') \chi_{\nu,\nu'}(\mathbf{k},\mathbf{k}'), \tag{2.69}$$

where

$$\chi_{\nu,\nu'}(\mathbf{k},\mathbf{k}') = T \sum_{\omega_n} G_\nu(\mathbf{k},\omega_n) G_{\nu'}(\mathbf{k}',\omega_n) = \frac{f(\varepsilon_\nu(\mathbf{k})) - f(\varepsilon_{\nu'}(\mathbf{k}'))}{\varepsilon_\nu(\mathbf{k}) - \varepsilon_{\nu'}(\mathbf{k}')}. \tag{2.70}$$

The mean-field disorder term is always positive, therefore destabilising the ferromagnet in favour of the paramagnet. This term leads only to a small quantitative change in the phase boundary and does not change the nature of either phase.

Clean case fluctuations

The Matsubara sums in Eq. 2.66 can be evaluated to generate the clean-case fluctuation term:

$$\mathcal{F}^{\text{fl}} = -2g^2 \int_{\mathbf{k}_1 \dots \mathbf{k}_4}^{(*)} \frac{f_+(\mathbf{k}_1) f_-(\mathbf{k}_2)(1 - f_+(\mathbf{k}_3))(1 - f_-(\mathbf{k}_4))}{\varepsilon_+(\mathbf{k}_1) + \varepsilon_-(\mathbf{k}_2) - \varepsilon_+(\mathbf{k}_3) - \varepsilon_-(\mathbf{k}_4)}, \tag{2.71}$$

where, as before, f is the Fermi-Dirac distribution function with $f_\nu(\mathbf{k}) = f(\varepsilon_\nu(\mathbf{k}))$ for brevity and the asterisk on the integral indicates momentum conservation, $\mathbf{k}_1 - \mathbf{k}_2 = \mathbf{k}_3 - \mathbf{k}_4$. This form of $\mathcal{F}_{fl}$ contains an ultraviolet divergence due to the contact (i.e. zero-range) interaction that must be regularised by setting $g \to g_{\mathbf{k}_1,\mathbf{k}_2} = g - 2g^2 \sum_{\mathbf{k}_3,\mathbf{k}_4} (\varepsilon_{\mathbf{k}_1}^+ + \varepsilon_{\mathbf{k}_2}^- - \varepsilon_{\mathbf{k}_3}^+ - \varepsilon_{\mathbf{k}_4}^-)^{-1}$ following the procedure of Refs. [83, 104], which leads to the final non-divergent form of the fluctuation correction:

$$\mathcal{F}^{\text{fl}} = -2g^2 \int_{\mathbf{k}_1 \dots \mathbf{k}_4}^{(*)} \frac{f_+(\mathbf{k}_1) f_-(\mathbf{k}_2)(f_+(\mathbf{k}_3) + f_-(\mathbf{k}_4))}{\varepsilon_+(\mathbf{k}_1) + \varepsilon_-(\mathbf{k}_2) - \varepsilon_+(\mathbf{k}_3) - \varepsilon_-(\mathbf{k}_4)}, \tag{2.72}$$

as used earlier in the chapter and in Refs. [83, 104, 106, 107].

Disorder-induced fluctuations

The leading fluctuation corrections due to disorder are given by:

$$\mathcal{F}_{\mathrm{dis}}^{\mathrm{fl}} = g\sigma_{\mathrm{c}}^2 \sum_{\nu,\nu'} \left[\int_{\mathbf{k}_1\ldots\mathbf{k}_4}^{(*)} (-1)^{\nu+\nu'} h_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) \chi_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) \chi_{-\nu,-\nu'}(\mathbf{k}_3,\mathbf{k}_4) \right.$$

$$\left. + \int_{\mathbf{k}_1\mathbf{k}_2\mathbf{k}_3} 2h_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) K_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) n_{-\nu'}(\mathbf{k}_3) \right], \qquad (2.73)$$

where

$$K_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) = T \sum_{\omega_n} G_\nu(\mathbf{k}_1,\omega_n) G_{\nu'}^2(\mathbf{k}_2,\omega_n). \qquad (2.74)$$

In practice, evaluation of this term can be made simpler by writing $K_{\nu,\nu'}$ as combinations of derivatives of $\chi_{\nu,\nu'}$ with respect to the chemical potential and magnetisation.

$$K_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2) = -\frac{1}{2}\left[\partial_\mu + \frac{1}{2}(\nu - \nu')\partial_M \right] \chi_{\nu,\nu'}(\mathbf{k}_1,\mathbf{k}_2). \qquad (2.75)$$

We may now carry out the expansion of the free energy in powers of q, as before, while retaining the full functional form of the m-dependence to take into account the resummation of higher-order divergences. This must be done numerically. All coefficients remain smooth and well-controlled throughout the region of interest. The resulting phase diagram is shown in Fig. 2.7.

The presence of quenched charge disorder has not destroyed the incommensurate phase, but has in fact slightly enhanced the area over which it is stabilised by the fluctuations. Numerically, we find that the fluctuations contribute a positive correction to the m^2 coefficient, destabilising the homogeneous ferromagnet, but the dominant contribution is a negative correction to the m^2q^2 term which acts to enhance the stability of the incommensurately ordered phase.

The quenched charge disorder considered here cannot have changed the nature of the ferromagnetic or paramagnetic phases, but it has resulted in additional disorder-induced fluctuation corrections. In the clean case, the fluctuation corrections generate a long-range ordered spiral magnet. While the disorder-induced fluctuations cannot change the nature of the ferromagnetic or paramagnetic phases, they do change the nature of the fluctuation-induced incommensurate phase.
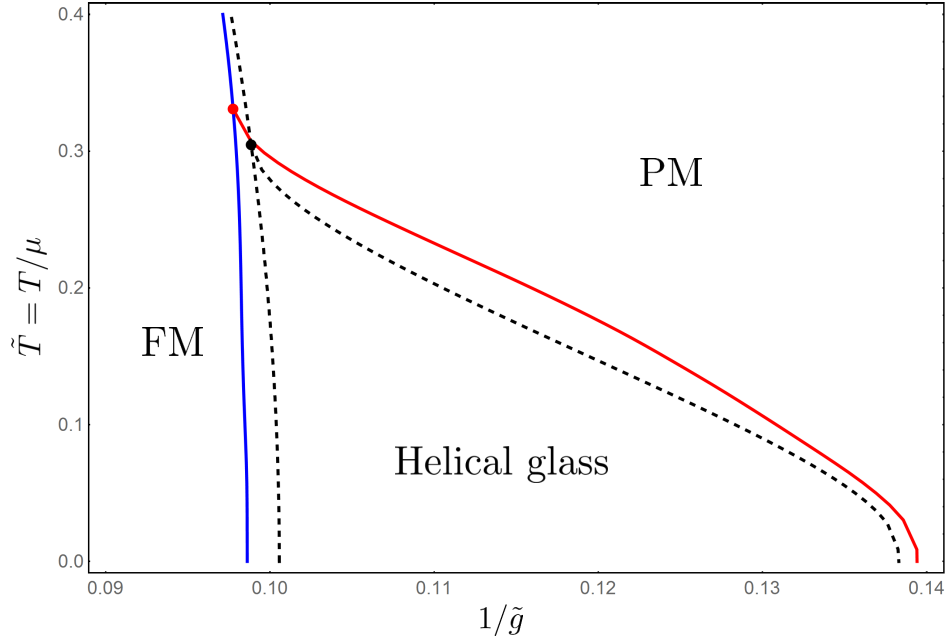
Figure 2.7 *The phase diagram of an itinerant ferromagnet with quenched charge disorder in the vicinity of the (avoided) quantum critical point, plotted with disorder strength $\sigma_c^2 = 0.8$. The black dashed lines represent the spiral magnetic region in the clean case. Upon the addition of disorder, the incommensurate phase is stabilised over a slightly larger area and its nature is changed from a long-ranged ordered spiral phase to a novel helical glass with local deformations in the pitch of the spiral due to the effects of disorder.*

2.10   The Helical Glass

Goldstone Modes

To see how the disorder affects the fluctuation-stabilised spiral phase, we need to examine its effects on the Goldstone modes of the system. The approach thus far has treated the Goldstone modes in a slightly naive way, and has assumed the Hamiltonian to be invariant under rotation of the spiral ordering wavevector.

The resulting Goldstone modes associated with such rotations are artefacts of this theoretical approach. In any real system, there will be anisotropies generated by the crystal field which, among other effects, will act to pin the spiral ordering wavevector along a particular lattice direction. Here, we add a cubic anisotropy term of the following form:

$$\mathcal{F}_{\mathrm{an}} = \lambda \int \mathrm{d}^3\mathbf{r} \sum_{\alpha=x,y,z} (\partial_\alpha m_\alpha)^2. \tag{2.76}$$

This can be derived, for example, by expanding around the tight-binding

53

dispersion at low filling fractions. For $\lambda > 0$ this is minimised by having the spiral wavevector point along one of the equivalent crystal axes. This pinning along a crystal axis does not change the form of the saddle-point equation used earlier to minimise $\mathbf{q}$. Here, as before, we choose $\mathbf{q} = q\hat{z}$ without loss of generality.

Now that we have restricted the direction of the spiral ordering wavevector, we have eliminated the Goldstone modes associated with continuous rotations of the direction of $\mathbf{q}$. The only remaining Goldstone modes are rotations in the phase of the ordering wavevector. The spiral order parameter becomes:

$$\mathbf{m}(\mathbf{r}) = m[\hat{x}\cos\left(qz + \phi(\mathbf{r})\right) + \hat{y}\sin\left(qz + \phi(\mathbf{r})\right)]. \tag{2.77}$$

The Goldstone modes of the spiral phase correspond to continuous deformations $\phi(\mathbf{r})$ of the phase of the spiral order parameter.

By expanding the action around the saddle point value $q^2 = -f_2(m^2)/f_4(m^2)$ in terms of this phase slip $\phi(\mathbf{r})$, we find that these deformations can be captured by a classical, anisotropic $3d$-XY action of the following form:

$$S_\phi \approx \int d^3\mathbf{r} \left[2f_4(m^2)q^2(\partial_z\phi)^2 + \frac{1}{2}\lambda m^2 \sum_{i=x,y} (\partial_i\phi)^2 \right], \tag{2.78}$$

which can be derived by re-inserting Eq. 2.77 into the free energy expression in Eq. 2.31, after re-expressing the q^2 terms in the free energy as gradient terms and allowing for a spatially varying phase $\phi(\mathbf{r})$.

The charge disorder that we consider here generates a random mass disorder for the ferromagnet, i.e. it appears in the coefficient of the m^2 term in the Landau expansion. This disorder leads to some quantitative changes but otherwise has no effect on the long-range magnetic correlations in the ferromagnetic phase. However, from the saddle point equation for $\mathbf{q}$, we see that what is a relatively weak random mass disorder in the ferromagnetic phase has a much stronger effect in the spiral, inducing a spatial variation in the spiral ordering wavevector $\mathbf{q}$.

In conjunction with the cubic anisotropy, this disorder in the pitch of the spiral leads to a much stronger random anisotropy disorder in the Goldstone modes $\phi(\mathbf{r})$. We can see this by plugging Eq. 2.77 into the anisotropy term Eq. 2.76, allowing for variations in the pitch of the spiral $q \to q + \delta q(\mathbf{r})$ and expanding in powers of δq. The net result of this procedure is that the disorder generates a random anisotropy term for the phase of the spiral:

$$S_\phi^{\text{dis}} = \frac{1}{2}\lambda m^2 \int d^3\mathbf{r}\, g(\mathbf{r})\cos\left[2\phi + \alpha(\mathbf{r})\right], \tag{2.79}$$

where $\alpha(\mathbf{r}) \sim \delta q(\mathbf{r})z$ is the random phase and $g(\mathbf{r}) = (1/4)[(\partial_y\alpha)^2 - (\partial_x\alpha)^2]$. This random anisotropy is only induced in the fluctuation-stabilised phase, not in the homogeneous ferromagnet. This demonstrates that the disorder has a distinctly different - and much stronger - effect on the spiral phase than on the

ferromagnet. It has previously been shown that random anisotropy disorder of this type is sufficient to destroy long range order in dimensions $d < 4$ [108–110]. The disordered spiral phase is therefore not a long-range ordered phase.

It is possible that the system may display quasi-long-range-order, though this is unlikely [111]. Weak random anisotropy disorder [112] in $O(N)$ magnets with $N < N_c = 9.44$ has been shown to lead to algebraic quasi-long-range-order for $d_{lc} < d < 4$ with a lower critical dimension $d_{lc} \approx 4 - 0.00158(N - N_c)^2$. In the case of the XY system considered here, $d_{lc}(N = 2) \approx 3.91$. Similar lower critical dimensions have been found using non-perturbative functional renormalisation group for the case of random field disorder [113], adding up to the conclusion that it is very likely the disordered spiral displays short-ranged magnetic correlations.

Correlation Length

To obtain an expression for the correlation length, we use the result of Ref. [114] which shows that the correlation length of a $3d$-XY model with random anisotropy disorder may be written as:

$$\xi = \frac{n^2 \rho_s^2}{8\pi D^2 K_n^2}, \tag{2.80}$$

where ρ_s is the spin stiffness, the random anisotropy disorder takes the form $D\cos(n\phi + \alpha(\mathbf{r}))$ and $K_n = \langle \cos(n\phi) \rangle_0$ is computed in the absence of disorder. For the random anisotropy we consider here, $n = 2$. We may extend this result to our anisotropic $3d$-XY model by rescaling the z co-ordinate, allowing us to write down an expression for the correlation length in the plane perpendicular to $\mathbf{q}$ and along the direction of the spiral:

$$\xi_{x,y} = \frac{1}{2\pi\sigma_g^2} \left(\frac{\lambda m^2}{4f_4(m^2)q^2} \right)^2 \exp\left(\frac{\tilde{T}}{\pi\sqrt{4\lambda f_4(m^2)m^2 q^2}} \right), \tag{2.81}$$

$$\xi_z = \frac{1}{2\pi\sigma_g^2} \left(\frac{\lambda m^2}{4f_4(m^2)q^2} \right)^{5/2} \exp\left(\frac{\tilde{T}}{\pi\sqrt{4\lambda f_4(m^2)m^2 q^2}} \right). \tag{2.82}$$

Note that in the limit of zero disorder, $\sigma_g^2 \to 0$, the correlation length diverges and long range order is recovered. This correlation length is highly anisotropic, with a ratio given by:

$$\frac{\xi_z}{\xi_{xy}} = \sqrt{\frac{4f_4(m^2)q^2}{\lambda m^2}}. \tag{2.83}$$

Figure 2.8 shows the pronounced anisotropy of the correlation length. The anisotropy is not entirely a surprise, since the 'background' ordering of the phase is a spiral ferromagnet, but the strong dependence on both temperature and interaction strength is unusual. The same is true for the ordering wavevector.

This potentially offers a new avenue for the detection of helical ordering in addition to, or instead of, neutron scattering. Establishing the presence of such a correlation length (in the disordered system only) and/or ordering wavevector that scales so strongly with temperature and some relevant experimental tuning parameter would be a strong indication for the presence of this type of phase.

2.11   Discussion

This result shows that disorder couples to fluctuations around ferromagnetic quantum critical points in a highly unusual way, destroying long range order and leading to the formation of an unusual glassy phase. This phase differs from typical glasses in that the short-range nature of its ordering stems entirely from the coupling of disorder to the soft modes of the system, rather than any frustration or frozen-in magnetic randomness.

Due to this unique formation method, the helical glass discussed here is different to other disordered spirals previously seen in the literature [115, 116]. This form of spiral ordering is not due to a mean-field Dzyaloshinskii-Moriya type interaction, nor to any form of nesting instability. The resultant glass produced by the addition of quenched charge disorder is therefore unique. A possible candidate for this phase has been found in CeFePO [84] which exhibits unusual glassy dynamics (seemingly short-ranged correlations and glassy relaxation times) in the vicinity of a putative quantum critical point, but curiously does not show typically glassy effects under field cooling. The experimental results show what appears to be a long-range ordered AFM state, however this could be the helical glass in a region of the phase where the correlation length is larger than the sample size, rendering the order effectively long-ranged on the scale of the experimental samples.

The material MnSi also displays an unusual partially ordered phase around its avoided quantum critical point [28, 29], but this phase is distinct from the helical glass which is described in this chapter. Disorder in this phase leads to the generation of skyrmions, or vortices in the orientation of $\mathbf{Q}$. The effects of the disorder on the helical glass are on the pitch of the spiral, not the direction of it, corresponding to a vortex line defect. It has been argued that fluctuations of any type of line defects should lead to a $T^{3/2}$ contribution to the electrical resistivity [117], and though the specifics of our system are different, the same logic should hold equally well in the case of the helical glass.

In the future it would be extremely interesting to see what effects disorder may have on fluctuations around spin density wave quantum critical points such as those thought to exist in certain superconducting materials, however application of the order-by-disorder framework to anything other than ferromagnetic systems has so far proven problematic. The parent materials of the high-temperature superconductors usually require doping in order to trigger superconductivity, and
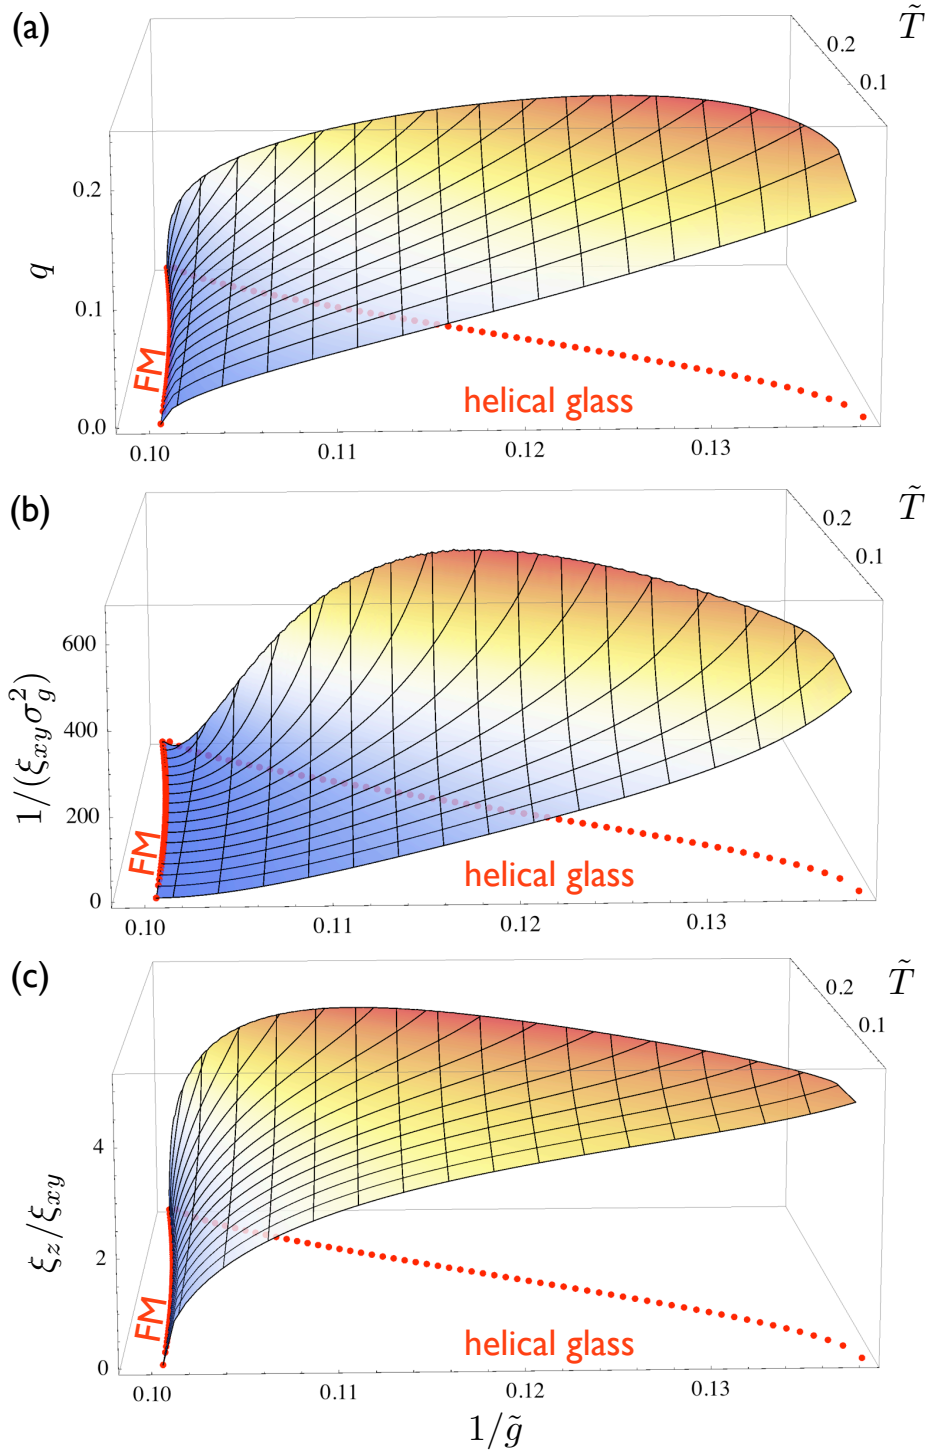
56

Figure 2.8 *Three dimensional plots showing the variation of a) the ordering wavevector q, b) the inverse xy correlation length and c) the correlation length anisotropy across the helical glass region. The dramatic variation of these quantities across the helical glass region is unusual and very distinctive. Figure reproduced from Ref. [79].*

though the effects of doping are typically considered to be limited to structural changes (i.e. applying stresses and strains to the chemical structure of the lattice) or simple increases and decreases in carrier concentrations, it seems likely that the disorder added in the doping process leads to non-trivial quantum fluctuations that may have a significant effect on the physics in the vicinity of these putative quantum critical points.

2.12   Conclusions

The main conclusions of the work presented in this chapter are as follows:

- The addition of quenched charge disorder into an itinerant ferromagnet leads to the destruction of the fluctuation-stabilised spiral phase previously seen in the vicinity of the quantum critical point, and instead leads to the establishment of an unusual short-range-ordered helical glass phase.

- The helical glass displays an unusual variation of ordering wavevector q, correlation length and correlation length anisotropy across the region where it is stabilised. This variation is distinctive and within reach of current experiments.

- In the vicinity of the helical glass to ferromagnet transition, the correlation length diverges. For experimental samples smaller than the correlation length, the system will appear to be long-range ordered but with unusual pseudo-glassy properties such as those seen in Ref. [84].

In the context of this thesis, the main importance of this work has been to establish that disorder can couple to quantum fluctuations stemming from a zero-temperature quantum critical point, yet still lead to significant changes in the behaviour of the material at finite temperatures. The helical glass phase discussed here is unique in that the effects of disorder are limited to the fluctuation-induced phase and do not affect the phases of the clean system. This calculation serves as a demonstration of the effects of disorder-induced quantum fluctuations at finite temperatures.

Chapter 3

Replica Symmetry Breaking in the Bose Glass

The Bose-Hubbard model is cosmetically similar to the fermionic Hubbard model studied in Chapter 2, with the difference being that rather than describing spin-1/2 fermions, its degrees of freedom are instead spinless bosons. Originally examined prior to the fermionic Hubbard Hamiltonian [118], the model of spinless bosons on a lattice was later re-introduced in its modern form to study superconducting systems where Cooper pairs of fermions tunnel between superconducting islands, for example in granular superconductors or Josephson junction arrays. The model rose to prominence after a seminal study [119] showed that it captured the essential physics of the superfluid-insulator transition and was more mathematically tractable than other metal-insulator models.

At zero temperature, the clean Bose-Hubbard model contains two possible phases: a gapped, incompressible Mott insulator (MI) and a gapless, compressible superfluid (SF). When disorder is added into the system, an intermediate phase known as the Bose glass (BG) intervenes. The Bose glass is a Griffiths phase characterised by rare superfluid regions existing within an otherwise Mott-insulating background. These rare ordered regions render the system globally gapless and compressible. The MI-to-SF quantum critical point is completely destroyed and instead replaced with a pair of MI-to-BG and BG-to-SF quantum critical points.

The nature of the Bose glass is still poorly understood, despite decades of theoretical work. My contribution to this problem was to investigate just how 'glassy' the Bose glass really is. We found that the Bose glass is characterised by non-zero Edwards-Anderson order parameters and the MI-to-BG transition is controlled by a finite-disorder fixed point which exhibits exact one-step replica symmetry breaking. The original work presented in this chapter was published in *"Replica symmetry breaking in the Bose glass"*, S. J. Thomson and F. Krüger, *Europhysics Letters* **108**, 30002 (2014) [120].

3.1 Background

The problem of disorder in strongly interacting bosonic systems has had a long and often controversial history. The question of disorder in the Bose-Hubbard model became relevant after it was investigated in the context of the superfluid-insulator transition [119, 121, 122] and contrasted with the much more difficult to deal with fermionic metal-insulator transition. The fermionic problem lacked the simple superfluid order parameter of the Bose-Hubbard model and incorporated numerous other complications that obscured the essential physics of the superfluid-insulator transition.

Despite its origins, the relevance of this problem is not limited to toy systems of interacting bosons as a poor substitute for fermionic metal-insulator transitions. It turns out that the physics underpinning this model is in fact very generic and is able to describe a wide class of phenomena. Bose-Hubbard-like models have also been used in coupled light-matter systems (in the form of the Jaynes-Cummings-Hubbard model [123] and Rabi-Hubbard model [124]) and as we'll see in the following chapter, the hard-core Bose-Hubbard model also describes dimerised quantum antiferromagnets. With the advent of modern optical lattice technology, the obvious experimental system nowadays in which to realise the conventional Bose-Hubbard model is that of ultracold atoms in an optical lattice, however there are now a wide variety of experimental systems capable of realising Bose-Hubbard-like physics.

The superfluid-insulator transition in a bosonic system can be expressed simply as the competition between kinetic energy (which favours delocalising the particles) and an on-site contact repulsion (which favours localising them). The transition between gapless superfluid and gapped Mott insulator is well described by a mean-field theory across most of the parameter space, with the exception of certain multicritical points which we shall largely avoid in this thesis. The addition of disorder into the Bose-Hubbard model was an interesting and logical next step, as disorder also acts to localise the particles. In this particular system, however, the disorder does more than simply destabilise the superfluid in favour of the localised insulating phase: it instead leads to the emergence of an entirely different phase known as the Bose glass which displays some unexpected properties and has been the subject of intense study ever since.

The Bose glass is a Griffiths phase, as introduced in Chapter 1. It is characterised by the presence of rare disconnected superfluid regions within a Mott-insulating background and is a gapless, compressible insulating phase. The Bose glass can be thought of as a 'smeared' MI-SF transition where regions of both phases coexist and the system is essentially a Mott insulator with disconnected superfluid pockets. For increased values of hopping, the 'puddles' eventually join up through a percolation transition and the system becomes superfluid.

For strong enough disorder ($\Delta \approx U$) we move from the interaction-dominated regime to the disorder-dominated regime where Anderson localisation starts to
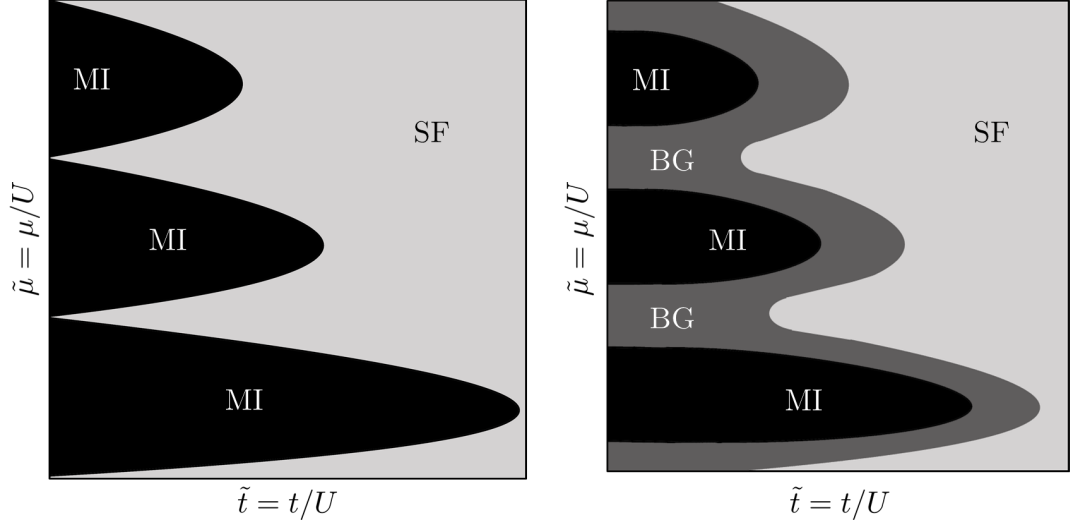
Figure 3.1 *A sketch of the phase diagram of the Bose-Hubbard model. i) The clean case phase diagram, showing Mott insulating (MI) lobes with integer filling surrounded by a superfluid (SF) region. ii) The disordered phase diagram, showing how the Mott lobes shrink and the Bose glass (BG) phase intervenes everywhere between the MI and SF.*

play a role, resulting in a phase known as an Anderson or Lifshitz glass [125]. In this thesis, we do not consider such strong disorder. A table of properties of the phases relevant to this thesis is shown in Table 3.1 - foreshadowing the findings of this chapter and in agreement with Ref. [126], the table displays the Edwards-Anderson order parameter and shows it takes a non-zero value in the glassy phases.

Although the Bose glass has been the subject of extensive theoretical work using analytic techniques spanning stochastic mean-field theory [127], strong coupling expansions [128–130] and renormalisation group [119, 120, 122, 131–137], as well as numerical techniques including quantum Monte Carlo [138–143], density matrix renormalisation group [144] and other methods [126, 145, 146], there are still many outstanding questions regarding the nature of the Bose glass.

For a long time, it remained unclear whether the Bose glass always intervenes between the MI and SF in the presence of disorder, or whether there could be a direct MI-SF transition, most likely at the tip of the Mott lobe where the universality class changes [119]. This was the subject of some controversy, but eventually a rigorous 'theorem of inclusions' was developed [147, 148] which forbade a direct MI-SF transition in the presence of even infinitesimal disorder. Due to the nature of a phase dominated by arbitrarily large rare regions, finite-size effects can present a serious problem to numerical simulations and it is likely that the direct MI-SF transition seen by many was an artefact of the strange suppression of compressibility in the vicinity of the multicritical points, seen even

Phases of the disordered Bose-Hubbard model			
	κ	ρ_{SF}	q_{EA}
Mott insulator	0	0	0
Bose glass	✓	0	✓
Mott glass	0	0	?
Superfluid	✓	✓	≈ 0

Table 3.1   *A table of properties of the phases being considered in this chapter showing how they differ from one another. The Bose glass has a finite compressibility κ and Edwards-Anderson order parameter q_{EA} but zero superfluid stiffness ρ_{SF}, while the homogeneous superfluid has a finite compressibility and superfluid stiffness but the Edwards-Anderson order parameter is zero. (In a disordered, inhomogeneous superfluid, q may be non-zero but small [126].) All three quantities are zero in the Mott insulator. The Mott glass has vanishing compressibility and superfluid fraction, and likely has a non-zero q. The Mott glass will be dealt with in more detail in Chapter 4.*

in the clean case [127].

Later work [149–154] has suggested the presence of an incompressible Mott glass at these multicritical points, suggesting that many of the previous works which used compressibility as the means of distinguishing the glass transition at the tip of the lobes would necessarily have missed the MG and may have interpreted their results as a direct MI-SF transition. We shall look at this in more detail in Chapter 4 in the context of dimerised antiferromagnets described by a hard-core Bose-Hubbard model, but for the purposes of this chapter we are first concerned with establishing the nature of the Bose glass phase itself and extracting an order parameter.

3.2   Clean Bose-Hubbard Model

The Hamiltonian in its modern second-quantised form is as follows:

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle} (\hat{b}_i^\dagger \hat{b}_j + \hat{b}_j^\dagger \hat{b}_i) - \mu \sum_i \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i (\hat{n}_i - 1), \qquad (3.1)$$

where $\hat{b}_i^\dagger$ and $\hat{b}_i$ are respectively the usual bosonic creation and annihilation operators and $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$ is the number operator on site i. The hopping amplitude is given by the coefficient t, the chemical potential by μ and the on-site Coulomb repulsion by U.

This Hamiltonian has a global $U(1)$ symmetry, such that it is invariant under the transformation $\hat{b}_j \rightarrow e^{i\theta} \hat{b}_j$. In the clean system, the model contains a gapped

Mott insulator (MI) phase which respects this symmetry, and a gapless superfluid (SF) in which the $U(1)$ symmetry is spontaneously broken.

We can see this by considering the limiting cases of interaction-dominated ($t = 0$) and hopping-dominated ($U = 0$) regimes. For $t = 0$, the Hamiltonian reduces to a purely local Hamiltonian, the eigenstates of which are n-particle Fock states $|n\rangle$, where the number n depends on the ratio μ/U.

$$\hat{\mathcal{H}}_0 = \sum_i \left[-\mu \hat{n}_i + \frac{U}{2} \hat{n}_i(\hat{n}_i - 1) \right] \rightarrow \text{eigenstate of } \prod_i |n_i\rangle. \qquad (3.2)$$

This gives a gapped spectrum with no long range order. This is the Mott insulating phase, characterised by an integer number of particles per site and an energy gap to excitations. The particles are localised on their respective lattice sites and cannot move throughout the lattice.

In the opposite regime where $U = 0$, the Hamiltonian is effectively that of a non-interacting gas of bosons which can be diagonalised in momentum space via Fourier transform, leading to a gapless spectrum and eigenstates that may be written as a product over sites of coherent states.

$$\hat{\mathcal{H}}_t = -t \sum_{\langle ij \rangle} (\hat{b}_i^\dagger \hat{b}_j + \hat{b}_j^\dagger \hat{b}_i) - \mu \sum_i \hat{n}_i \qquad (3.3)$$

$$= \sum_q (\varepsilon_q - \mu) \hat{n}_q, \qquad (3.4)$$

where $\varepsilon_q = -2t \sum_\lambda \cos(q e_\lambda)$ with e_λ the unit vector along the λ axis of the lattice. This leads to a gapless dispersion for $q = 0$. In the truly non-interacting case, though all of the bosons will condense into the lowest mode there will be no long-range order and no superfluid stiffness. The presence of non-zero but arbitrarily small interactions, however, will lead to the establishment of long-range order. This is the superfluid phase, characterised by the uninhibited motion of particles throughout the lattice and the spontaneous breaking of the $U(1)$ symmetry (i.e. the eigenstates of this Hamiltonian are coherent states which pick out a phase).

Given that we have two very different states here depending on whether U or t dominates, we can expect to see a quantum phase transition between these states as a function of t/U. To capture this, we first turn to a mean-field theory.

3.2.1 Mean-Field Theory

In the clean system, the MI-SF transition is well-described by a mean-field (MF) theory [10]. We can make a mean-field decoupling in terms of the superfluid

order parameter ψ_i and write the Hamiltonian as a sum over sites:

$$\mathcal{H}_{\mathrm{MF}} = \sum_i \left[-\mu \hat{n}_i + \frac{U}{2} \hat{n}_i(\hat{n}_i - 1) - zt(\psi^* \hat{b}_i + \psi \hat{b}_i^\dagger) + zt\psi^*\psi \right], \qquad (3.5)$$

where $z = 2d$ is the lattice co-ordination number, using $\hat{b}_i \to \psi + \delta\hat{b}$ with $\psi = \langle \hat{b}_i \rangle$, neglecting terms quadratic in $\delta\hat{b}_i$ and relabelling $\delta\hat{b}_i$ back to $\hat{b}_i$. This MF decoupling gives rise to terms which break the $U(1)$ symmetry for $\psi \neq 0$ and do not conserve the number of particles.

As shown in Ref. [10], any interacting system of this form which breaks $U(1)$ symmetry will will develop a non-zero stiffness to any further rotations of the order parameter. The situation where $\psi \neq 0$ therefore corresponds to a superfluid state with a non-zero superfluid stiffness, therefore the mean-field decoupling here is precisely the right one to describe the Mott insulator to superfluid transition. The precise form of the mean-field decoupling used here assumes ψ to be the same on every site, precluding phases such as supersolids which are assumed not to exist in this model, though they can exist in related models with finite-range interactions and different lattice geometries.

We can obtain the ground state energy by computing the expectation value of the Hamiltonian in the Fock basis. We first rewrite the Hamiltonian as $\mathcal{H}_{\mathrm{MF}} = \sum_i \left[h_i^{(0)} - zt(\psi^* \hat{b}_i + \psi \hat{b}_i^\dagger) \right]$, where $h_i^{(0)} = -\mu \hat{n}_i + \frac{U}{2} \hat{n}_i(\hat{n}_i - 1) + zt\psi^*\psi$ and we treat the other piece as a perturbation due to the presumed smallness of ψ at the phase transition. The zeroth-order piece of the Hamiltonian $h^{(0)}$ is diagonal in the Fock basis, giving a contribution to the energy:

$$E_m^{(0)} = -\mu m + \frac{U}{2} m(m - 1) + zt|\psi|^2, \qquad (3.6)$$

where m is the integer that minimises $E_m^{(0)}$ for a given value of μ/U and I drop the site index as all sites are identical. The correction due to the additional terms may be calculated by second-order perturbation theory:

$$E_n^{(2)} = zt|\psi|^2 \sum_{n \neq m} \frac{|\langle m|(\hat{b}_i^\dagger + \hat{b}_i)|n\rangle|^2}{E_n^{(0)} - E_m^{(0)}}$$

$$= -(zt)^2 |\psi|^2 \left(\frac{m}{U(m-1) - \mu} + \frac{m+1}{\mu - Um} \right). \qquad (3.7)$$

We can use the usual Landau formalism to write the free energy in even powers of the order parameter ψ to get $E = const. + R|\psi|^2 + W|\psi|^4 + ...$ The second-order MI-SF transition is given by the point at which the coefficient $r = R/zt$ changes sign, i.e. the $r = 0$ line. Combining the zeroth-order and perturbative correction to the energy calculated above, we find a dimensionless mass coefficient:

$$r = \frac{R}{zt} = 1 + zt \left(\frac{m}{U(m-1) - \mu} + \frac{m+1}{\mu - Um} \right). \qquad (3.8)$$
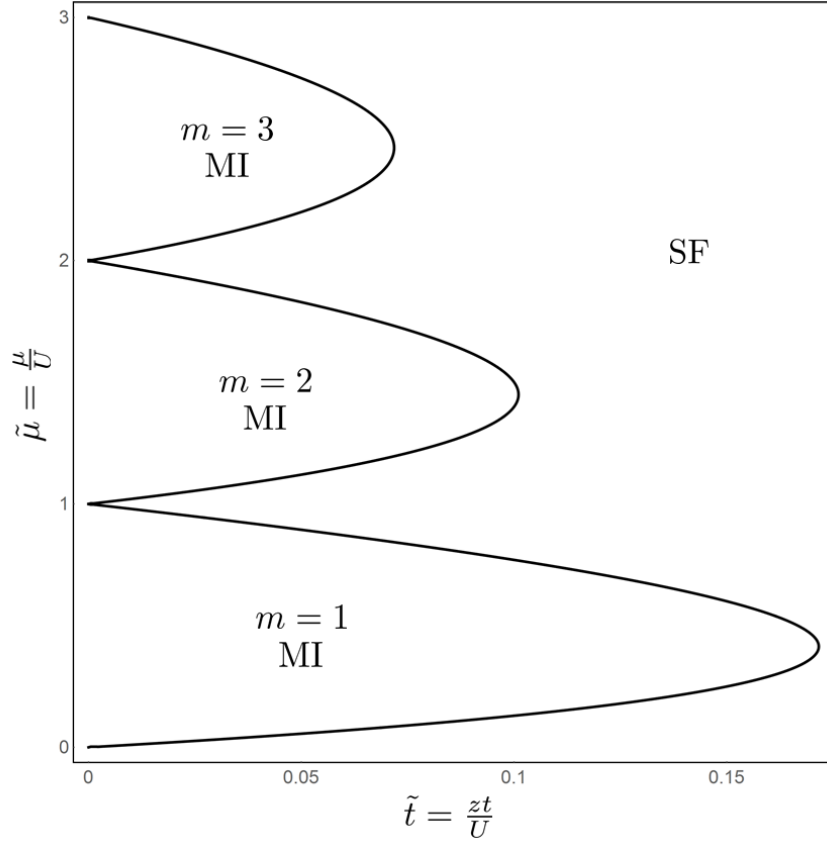
Figure 3.2　*The mean-field phase diagram of the clean Bose-Hubbard model given by Eq. 3.8, showing the first three Mott insulating (MI) lobes of integer filling m and the surrounding superfluid (SF) region.*

We can plot the $r = 0$ line for a variety of different integer fillings m and we obtain a phase diagram with the familiar Mott-lobe structure, shown in Fig. 3.2.

3.2.2　Dual Field Theory

Strong-Coupling Expansion

Since we ultimately want to examine the disordered Bose-Hubbard model, and as the glassy phases we are interested in cannot be captured by this mean-field approach, we need to utilise a more sophisticated method.

Following Refs. [120, 135, 136], I outline the construction of a strong-coupling dual field theory [10, 128, 130], with an eye to eventually employing a momentum-shell renormalisation group analysis to extract the behaviour of the system. The mass term we shall derive in the strong coupling theory to signify the mean-field MI-SF transition is identical to the mass term derived in the previous section.

The idea behind the development of this strong-coupling expansion is to partially resum the interaction vertex and express the action in terms of some new auxiliary fields which take into account the effects of interactions at quadratic order. It can be viewed as a type of dual field theory, mapping a strongly interacting problem (where the quartic terms are vital) onto a Gaussian action where the quartic terms are irrelevant (in the renormalisation group sense).

To construct our strong coupling field theory, we first express the action as a path integral over bosonic coherent states:

$$\mathcal{Z} = \int \mathcal{D}[\phi^*, \phi] \mathrm{e}^{-S[\phi^*, \phi]}, \tag{3.9}$$

$$S[\phi^*, \phi] = \int_\beta^0 \mathrm{d}\tau \left[\sum_i (\phi_i^* \partial_\tau \phi_i - \mu |\phi_i|^2) + U \sum_i |\phi_i|^4 - t \sum_{ij} (\phi_i^* \phi_j + \phi_j^* \phi_i) \right]. \tag{3.10}$$

where $\hat{b}_i |\phi_i\rangle = \phi_i |\phi_i\rangle$. We then employ a Hubbard-Stratonovich transform to decouple the hopping term (see Appendix A and Refs. [9, 10] for more details).

Though this is essentially the same technique as used in Chapter 2 to decouple the interaction term of the fermionic Hubbard Hamiltonian, here it is used differently. Decoupling the hopping in this manner is what allows us to express the action in terms of the dual fields with partially resummed interaction vertex, rather than being a spiritually mean-field manoeuvre as in the aforementioned decoupling of the fermionic Hubbard Hamiltonian interaction term. Once we have performed the decoupling, we can write the partition function as $\mathcal{Z} = \int \mathcal{D}[\phi^*, \phi, \psi^*, \psi] \exp[-(S_0 + S_{\phi\psi})]$ where:

$$S_0 = \int_\beta^0 \mathrm{d}\tau \left[\sum_i (\phi_i^* \partial_\tau \phi_i - \mu |\phi_i|^2) + U \sum_{ij} |\phi_i|^4 \right], \tag{3.11}$$

$$S_{\phi\psi} = \int_0^\beta \left[\sum_{ij} \psi_i^* T_{ij}^{-1} \psi_j + \psi_i^* \phi_i + \phi_i^* \psi_i \right]. \tag{3.12}$$

We can now trace out the original fields to obtain a partition function of the form $\mathcal{Z} = \mathcal{Z}_0 \int \mathcal{D}[\psi^*, \psi] \exp[-(S_\psi^{(0)} + S_\psi')]$ where

$$S_\psi^{(0)} = \int_0^\beta \mathrm{d}\tau \sum_{\langle ij \rangle} \psi_i^* T_{ij}^{-1} \psi_j, \tag{3.13}$$

$$S_\psi' = -\ln \left\langle T_\tau \exp \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^* \phi_i + \phi_i^* \psi_i) \right] \right\rangle_0. \tag{3.14}$$

where T_τ is the time-ordering operator and the average is taken with respect to the local action S_0.

Calculating the local terms

The average in Eq. 3.14 can be re-expressed as an operator average taken with respect to the local Hamiltonian:

$$\hat{\mathcal{H}}_0 = \sum_i \left[-\mu \hat{n}_i + \frac{U}{2} \hat{n}_i(\hat{n}_i - 1) \right]. \tag{3.15}$$

This Hamiltonian is not quadratic, meaning Wick's theorem cannot be employed to take the average in Eq. 3.14 and instead it must be calculated in full. This is the step which allows us to incorporate the interactions into the propagator of the dual field theory. As an operator average:

$$S'_\psi = -\ln \left\langle T_\tau \exp \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau)\hat{b}_i(\tau) + \hat{b}_i^\dagger(\tau)\phi_i(\tau)) \right] \right\rangle_0. \tag{3.16}$$

Expanding both the logarithm and the exponential to quartic order leads to:

$$\begin{aligned} S'_\psi = &-\frac{1}{2} \left\langle T_\tau \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau)\hat{b}_i(\tau) + \hat{b}_i^\dagger(\tau)\phi_i(\tau)) \right]^2 \right\rangle_0 \\ &-\frac{1}{4!} \left\langle T_\tau \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau)\hat{b}_i(\tau) + \hat{b}_i^\dagger(\tau)\phi_i(\tau)) \right]^4 \right\rangle_0 \\ &+\frac{1}{8} \left\langle T_\tau \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau)\hat{b}_i(\tau) + \hat{b}_i^\dagger(\tau)\phi_i(\tau)) \right]^2 \right\rangle_0^2. \end{aligned} \tag{3.17}$$

In principle, we could expand Eq. 3.17 to any order and construct an action with arbitrarily high order terms, however quartic order will prove to be sufficient for our purposes. Taking the quadratic terms alone for the moment:

$$S_\psi^{(2)} = -\frac{1}{2} \int_0^\beta \mathrm{d}\tau_1 \int_0^\beta \mathrm{d}\tau_2 \sum_{i,j} \psi_i^*(\tau_1)\psi_i(\tau_2) \left\langle T_\tau \hat{b}_i(\tau_1)\hat{b}_i^\dagger(\tau_2) \right\rangle_0. \tag{3.18}$$

With a change of variables $\tau_1 = \tau$, $\tau' = \tau_1 - \tau_2$ and an expansion of $\psi(\tau - \tau') = \psi(\tau) - \partial_\tau \psi_i(\tau) \cdot \tau' + (1/2)\partial_\tau^2 \psi_i(\tau) \cdot (\tau')^2$, we arrive at a contribution to the effective quadratic action of the form:

$$S_\psi^{(2)} = -\frac{1}{2} \int_0^\beta \sum_i \left(K_1 \psi_i^* \partial_\tau \psi_i + K_2 |\partial_\tau \psi_i|^2 + G|\psi_i|^2 \right), \tag{3.19}$$

where the coefficients are given by:

$$G = -\int_{-\beta}^{\beta} \mathrm{d}\tau \left\langle T_\tau \hat{b}_i(\tau)\hat{b}_i^\dagger(0) \right\rangle_0, \tag{3.20}$$

$$K_1 = \int_{-\beta}^{\beta} \mathrm{d}\tau\, \tau \left\langle T_\tau \hat{b}_i(\tau)\hat{b}_i^\dagger(0) \right\rangle_0, \tag{3.21}$$

$$K_2 = -\frac{1}{2}\int_{-\beta}^{\beta} \mathrm{d}\tau\, \tau^2 \left\langle T_\tau \hat{b}_i(\tau)\hat{b}_i^\dagger(0) \right\rangle_0. \tag{3.22}$$

At $T = 0$ we know the possible states of the system and so we can directly calculate $G = G_m$ in the vicinity of the m^{th} Mott lobe. Computing the time-ordered expectation values gives:

$$\tau > 0: \qquad \left\langle \hat{b}_i(\tau)\hat{b}_i^\dagger(0) \right\rangle_0 = \langle m|e^{\hat{\mathcal{H}}_0\tau}\hat{b}_i e^{-\hat{\mathcal{H}}_0\tau}\hat{b}_i^\dagger|m\rangle \tag{3.23}$$

$$= (m+1)\exp\left[-\tau(Um-\mu)\right], \tag{3.24}$$

$$\tau < 0: \qquad \left\langle \hat{b}_i^\dagger(0)\hat{b}_i(\tau) \right\rangle_0 = \langle m|\hat{b}_i^\dagger e^{\hat{\mathcal{H}}_0\tau}\hat{b}_i e^{-\hat{\mathcal{H}}_0\tau}|m\rangle \tag{3.25}$$

$$= m\exp\left[\tau(U(1-m)+\mu)\right]. \tag{3.26}$$

Performing the integral over τ in Eq. 3.20 then leads to the final result:

$$G_m = -\left(\frac{m+1}{mU-\mu} + \frac{m}{U(1-m)+\mu}\right). \tag{3.27}$$

K_1 and K_2 may be calculated explicitly in the same manner, or one may use the fact that they are given by derivatives of G, e.g. $K_1 = -\partial G/\partial\mu$ and $K_2 = -(1/2)\partial^2 G/\partial\mu^2$. This is a consequence of demanding that Eq. 3.19 remain invariant under the time-dependent $U(1)$ gauge transformation $\hat{b}_i \rightarrow \hat{b}_i e^{i\phi(\tau)}$, $\psi_i \rightarrow \psi_i e^{i\phi(\tau)}$ and $\mu \rightarrow \mu + i\partial\phi/\partial\tau$. As discussed in Ref. [10], although the time-dependent μ generated by this transformation means we are no longer in a physical parameter regime, it is useful for enabling us to determine K_1 without need for explicit evaluation of Eq. 3.21.

The relationship between the frequency coefficients K_1 and K_2 and the zero-frequency Green's function G_m leads to an important consequence. When $K_1 \neq 0$, the higher-order term K_2 may be neglected as it is irrelevant with respect to K_1. The derivative in K_1 evaluates to zero when the MI-SF phase boundary shown in Fig. 3.2 has a vertical tangent, i.e. at the tip of the Mott lobes. Since the dynamical critical exponent is determined by the scaling of lowest-order non-zero frequency term, the disappearance of K_1 tells us that the universality class changes at the tip of the lobes. Specifically, the $K_1 = 0$ points are described by a relativistic quantum field theory which maps onto the $N = 2$ quantum rotor model, while the $K_1 \neq 0$ points obey the non-relativistic theory where K_2 may be neglected. For the remainder of this chapter, the frequency dependence will be unimportant, however we shall see in Chapter 4 that correct treatment of these frequency terms is critical for accurate calculations of certain thermodynamic properties.

Evaluating the Hopping Terms

The hopping term $S_\psi^{(0)}$ must be treated using a continuum expansion in the long-wavelength limit, since the hopping matrix is singular and cannot be directly inverted. The singularity of the matrix is pathological to the case of nearest-neighbour hopping and the continuum limit approach is a way of avoiding the singularity without complicating the model with the addition of longer-range hopping terms:

$$\int_\tau d\tau \sum_{ij} T_{ij}\phi_i^*\phi_j = t\int_\tau d\tau \sum_{\langle ij\rangle}(\phi_i^*\phi_j + \phi_j^*\phi_i) \tag{3.28}$$

$$= \int_\tau \int_k \left(2t\sum_\alpha^d \cos(k_\alpha)\right)\phi^*(k)\phi(k) \tag{3.29}$$

$$\approx \int_\tau \int_k t(z - k^2)\phi^*(k)\phi(k), \tag{3.30}$$

where $z = 2d$ is the lattice co-ordination number. The corresponding term in the dual action is:

$$S_\psi^{(0)} = \int_\tau \int_k \frac{1}{t(z-k^2)}\psi^*(k)\psi(k) \tag{3.31}$$

$$\approx \int_\tau \int_k \left(\frac{1}{zt} + \frac{1}{z^2 t}k^2\right)\psi^*(k)\psi(k). \tag{3.32}$$

This yields the contribution:

$$S_\psi^{(0)} = \int_{\mathbf{r},\tau} \left(\frac{1}{zt}|\psi(\mathbf{r},\tau)|^2 + \frac{1}{z^2 t}|\nabla\psi(\mathbf{r},\tau)|^2\right). \tag{3.33}$$

Combining with Eq. 3.19, this leads to an effective real-space quadratic action of the form:

$$S_\psi^{(2)} = -\frac{1}{2}\int_{\mathbf{r},\tau} \left(K_1\psi^*(\mathbf{r},\tau)\partial_\tau\psi(\mathbf{r},\tau) + K_2|\partial_\tau\psi(\mathbf{r},\tau)|^2 + R|\psi(\mathbf{r},\tau)|^2 + K_3|\nabla\psi(\mathbf{r},\tau)|^2\right). \tag{3.34}$$

Calculating the quartic vertex

We can similarly calculate the coefficient of the quartic vertex W by taking the relevant terms from Eq. 3.17, effectively making a direct computation of the local connected two-particle Green's function.

$$S^{(4)} = \frac{1}{8} \left\langle T_\tau \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau) b_i(\tau) + b_i^\dagger(\tau) \phi_i(\tau)) \right]^2 \right\rangle_0^2$$
$$- \frac{1}{4!} \left\langle T_\tau \left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau) b_i(\tau) + b_i^\dagger(\tau) \phi_i(\tau)) \right]^4 \right\rangle_0, \qquad (3.35)$$

which, after taking the continuum limit, gives an expression of the form $W \int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4$. The full calculation of W is lengthy but is essentially only a more elaborate version of the calculation of the single-particle Green's function and can be evaluated in the same manner. I do not reproduce the calculation here but instead quote the final result from Refs. [130, 135]:

$$W = \left(\frac{m}{(1-m)U + \mu} + \frac{m+1}{mU - \mu} \right) \left(\frac{m}{((1-m)U + \mu)^2} + \frac{m+1}{(mU - \mu)^2} \right)$$
$$- \frac{m(m-1)}{((1-m)U + \mu)^2((3-2m)U + 2\mu)} - \frac{(m+1)(m+2)}{(mU - \mu)^2((1+2m)U - 2\mu)}. \qquad (3.36)$$

The coefficient W is always greater than zero, confirming that all transitions captured by this action are second-order.

Final Dual Field Theory

We then perform a rescaling, employing dimensionless units for imaginary time and length such that $\tau \to V\tau$ and $\mathbf{r} \to \Lambda\mathbf{r}$ where Λ is the momentum cutoff we shall eventually employ. We then rescale the other coefficients as $r = R/K_3$, $\gamma_1 = (K_1/K_3)V$, $\gamma_2 = (K_2/K_3)V^2$ and $w = W/K_3$ to obtain:

$$S = \sum_\alpha \int_{\mathbf{k},\omega} \Omega(\mathbf{k},\omega) \psi_\alpha^\dagger(\mathbf{k},\omega) \psi_\alpha(\mathbf{k},\omega) + w \int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4 + ... \qquad (3.37)$$

where

$$\Omega(\mathbf{k},\omega) = (k^2 - i\gamma_1\omega + \gamma_2\omega^2 + r), \qquad (3.38)$$

$$r = 1 + zt \left(\frac{n}{U(n-1) - \mu} + \frac{n+1}{\mu - Un} \right), \qquad (3.39)$$

which is the same as the mass term derived from the mean field theory in Eq. 3.8 and consequently leads to the same phase diagram as shown in Fig. 3.2.
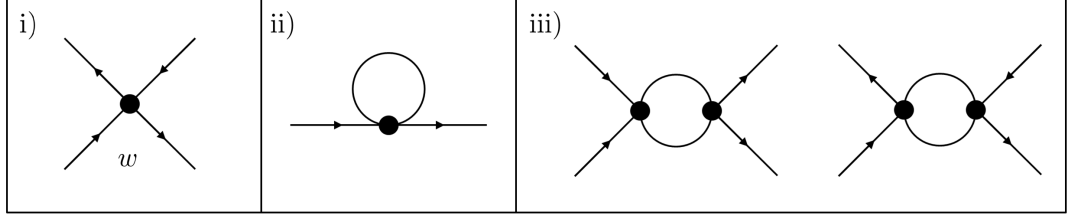
70

Figure 3.3 *i) Diagrammatic representation of the interaction vertex w in Eq. 3.37. ii) The contribution to the renormalisation of the mass coefficient r from the interaction vertex w. iii) The diagrams contributing to the renormalisation of the interaction vertex itself.*

3.3 Clean Case Renormalisation Group

We use a momentum-shell renormalisation group (RG) approach to extract the behaviour of the dual field theory. As described schematically in Chapter 1, we integrate out the highest energy modes which correspond to momenta in the infinitesimal shell $\mathrm{e}^{-dl} \leq |\mathbf{k}| \leq 1$. To regain resolution, we rescale the momenta $\mathbf{k} \to \mathbf{k}\mathrm{e}^{dl}$, frequencies $\omega \to \omega \mathrm{e}^{zdl}$ and the fields $\psi_\alpha \to \psi_\alpha^{-\lambda dl}$. All of the following is done to one-loop order.

We decompose the fields into 'fast' and 'slow' like so:

$$\psi(\mathbf{r}, \tau) = \psi(\mathbf{r}, \tau)^< + \psi(\mathbf{r}, \tau)^> \tag{3.40}$$

$$= \int_{|\mathbf{k}|<\mathrm{e}^{-dl}} \int_\omega \psi(\mathbf{k}, \omega)\mathrm{e}^{-i(\mathbf{k}\cdot\mathbf{r}-\omega\tau)}$$

$$+ \int_{\mathrm{e}^{-dl}<|\mathbf{k}|<1} \int_\omega \psi(\mathbf{k}, \omega)\mathrm{e}^{-i(\mathbf{k}\cdot\mathbf{r}-\omega\tau)}. \tag{3.41}$$

We then take the trace over the 'fast' fields and compute the correction of order dl to the 'slow' action. The diagrams we evaluate are shown in Fig. 3.3. As in previous work [120, 135, 136], we make the change of variables $I_0 = 1/(1 + r)$ where I_0 is the asymptotic inverse mass. The renormalisation group equations are:

$$\frac{\mathrm{d}r}{\mathrm{d}l} = 2r + 2I_1 w, \tag{3.42}$$

$$\frac{\mathrm{d}I_0}{\mathrm{d}l} = -2I_0 + 2I_0^2 - 2I_1 I_0^2 w, \tag{3.43}$$

$$\frac{\mathrm{d}\gamma_1}{\mathrm{d}l} = (2 - d_z)\gamma_1, \tag{3.44}$$

$$\frac{\mathrm{d}\gamma_2}{\mathrm{d}l} = (2 - 2d_z)\gamma_2, \tag{3.45}$$

$$\frac{\mathrm{d}w}{\mathrm{d}l} = (4 - d - d_z)w - (I_2 + 4I_3)w^2, \tag{3.46}$$

where d_z is the dynamical critical exponent, $I_1 = \int_\omega \Omega(\mathbf{k} = 1, \omega) = 1/\sqrt{\gamma_1^2 + 4\gamma_2(1 + r)}$, $I_2 = \int_\omega \Omega(\mathbf{k} = 1, \omega)^2 = 2\gamma_2 I_1^3$ and $I_3 = \int_\omega |\Omega(\mathbf{k} = 1, \omega)|^2 = (1/2)I_0 I_1$. Away from the tips of the Mott lobes, the dynamical critical exponent d_z is equal to 2 in order that γ_1 does not scale under the RG. At the tips of the Mott lobes where $K_1 = 0$ we instead require that γ_2 not scale under the RG and so there we find $d_z = 1$.

In all of the following we work in $d = 3$. Although w is an irrelevant variable away from the tip for $d > 2$, we include it for completeness to check for any dangerously irrelevant couplings. We find none: as expected, the RG including the interaction vertex leads to marginally more stable MI regions than predicted by mean-field theory but otherwise makes no qualitative difference. It has been previously established that mean-field theories overestimate superfluid coherence in this model so this finding is not surprising. In the following, we shall neglect this vertex and consider a simplified model including only the quadratic terms of the clean action.

To draw the phase diagram (Fig. 3.4), we interpret the RG flow such that when I_0 reaches some arbitrary cutoff value greater than 1, the flow is into the superfluid phase, whereas when I_0 flows to zero, the flow is into the Mott insulator. The position of the phase boundary drawn does not strongly change with the choice of cutoff, as the flow is extremely fast.

3.4   Replica Field Theory

Now we add disorder into the model in the form of a site-dependent chemical potential μ_i such that the Hamiltonian becomes:

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle} (\hat{b}_i^\dagger \hat{b}_j + \hat{b}_j^\dagger \hat{b}_i) - \sum_i \mu_i \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i (\hat{n}_i - 1). \qquad (3.47)$$

In the following, we consider only random site disorder such that $\mu_i = \overline{\mu} + \varepsilon_i$ where the random variation ε_i is drawn from a symmetric box distribution of width $2\delta = 2\Delta/U$. Physically, in an ultracold atomic lattice this form of disorder could be applied by superimposing a speckle pattern over the lattice potential [155–158], or by varying the lattice depth from site to site using a spatial light modulator (SLM) [159–161].

It is important for what follows that we use a bounded disorder distribution. Unbounded disorder will destroy the Mott lobes entirely and cause some of the quantities we calculate to become ill-defined. Note also that we do not consider hopping disorder in the following: the Hubbard-Stratonovich transform performed in Eq. 3.12 becomes much more complicated once disorder enters in
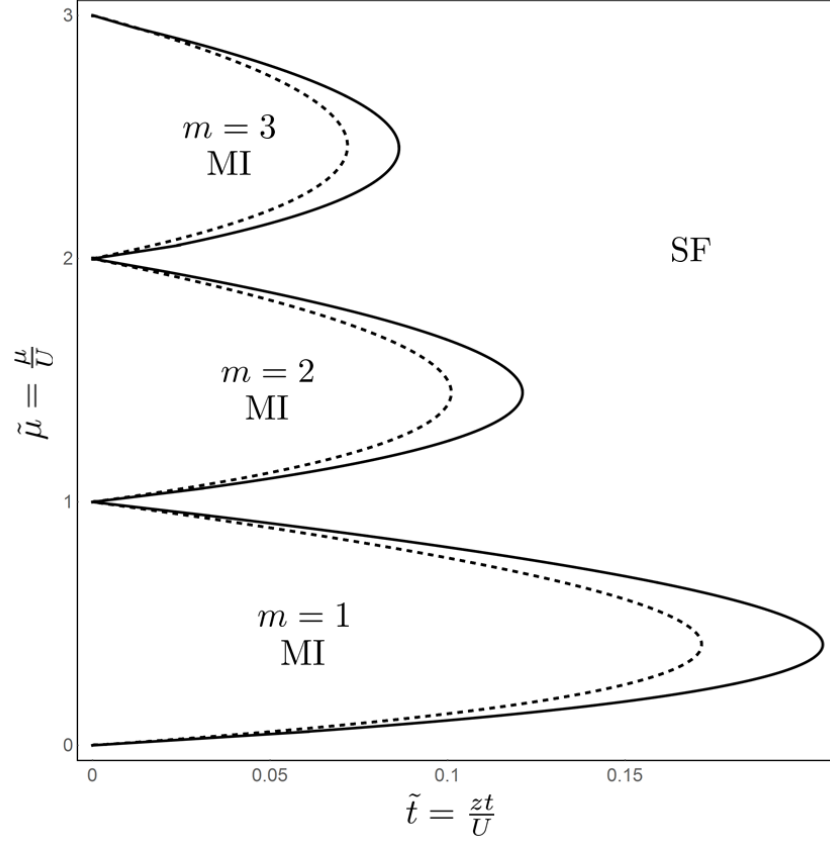
Figure 3.4 *The mean-field MI-SF transition line (dashed) shown alongside the phase boundary drawn from renormalisation group (solid). The mean-field theory overestimates the superfluid coherence and consequently underestimates the size of the Mott lobes. The RG predicts larger Mott lobes, in qualitative agreement with quantum Monte Carlo and other beyond-mean-field methods.*

the hopping matrix, so we restrict our analysis to on-site disorder only. A full renormalisation group analysis including hopping disorder was performed in Ref. [135].

Through the same procedure as in Section 3.1 we may derive an action for the system of the form:

$$S = \sum_\alpha \int_{\mathbf{k},\omega} (k^2 + -i\gamma_1\omega + \gamma_2\omega^2 + r)\psi_\alpha^\dagger(\mathbf{k},\omega)\psi_\alpha(\mathbf{k},\omega) + w \int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4 + ...$$

(3.48)

where all coefficients are now functions of the disorder. We take the disorder average using the replica trick, arriving at a disorder-averaged action:

$$S = \sum_\alpha \left[\int_{\mathbf{k},\omega} (k^2 + -i\overline{\gamma_1}\omega + \overline{\gamma_2}\omega^2 + \overline{r})\psi_\alpha^\dagger(\mathbf{k},\omega)\psi_\alpha(\mathbf{k},\omega) + \overline{w} \int_{\mathbf{r},\tau} |\psi_\alpha(\mathbf{r},\tau)|^4 \right]$$
$$- \frac{g}{2} \sum_{\alpha,\beta} \int_{\mathbf{k}_i,\omega,\omega'}^{'} \psi_\alpha^\dagger(\mathbf{k}_1,\omega)\psi_\alpha(\mathbf{k}_2,\omega)\psi_\beta^\dagger(\mathbf{k}_3,\omega')\psi_\beta(\mathbf{k}_4,\omega') + ... \quad (3.49)$$

where $g = (\overline{r^2} - \overline{r}^2)$ is the variance of the mass coefficient induced by the presence of disorder and the overline indicates the disorder average. The new disorder vertex, controlled by the coefficient g, is off-diagonal in both replica index and imaginary time (transformed to frequency in Eq. 3.49 above). This non-locality is a consequence of the perfect correlation of the disorder in imaginary time and it a generic feature of $T = 0$ systems with quenched disorder [119, 121].

In Ref. [119], this same effective action was dealt with using a double power series expansion in ε_τ and $\varepsilon = 4 - d - \varepsilon_\tau$ (i.e. like a regular epsilon-expansion but with an additional ε_τ variable due to the perfect correlation in imaginary time). They find that the renormalisation group equation for g takes the form:

$$\frac{\partial g}{\partial l} = (\varepsilon + \varepsilon_\tau)g + Bg^2 + ... \quad (3.50)$$

where $B > 0$ and therefore g increases without bound, signifying runaway flow to some infinite randomness fixed point not captured by their calculation. In other words, it is not possible to deal with this using a standard epsilon-expansion and we must find a more suitable variable to use that does not suffer from such a runaway flow.

3.4.1 Replica Symmetric System

First, I briefly recap the previous replica symmetric analysis [135, 136] where we take the action in Eq. 3.49 and put it through the renormalisation group machinery. The addition of a disorder vertex means the RG equations must be modified, with an additional RG equation for g and additional contributions to
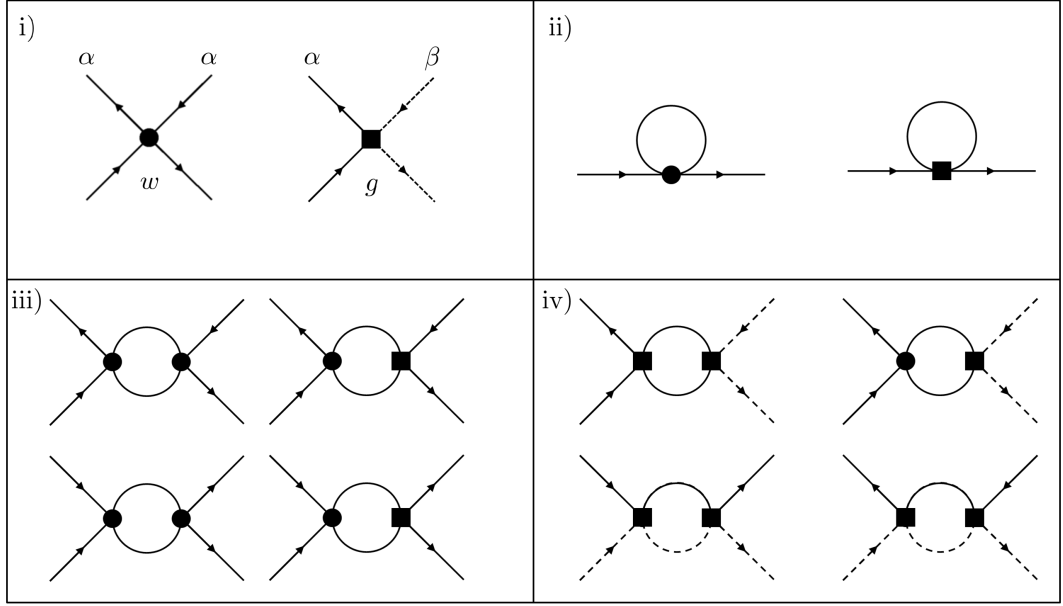
Figure 3.5 *A list of all the diagrams computed in calculating the renormalisation group equations, based on Ref. [135]. i) Diagrammatic representation of the interaction vertex (left) and disorder vertex (right). The interaction vertex is diagonal in replica index, whereas the disorder vertex is not (indicated by dashed lines). ii) The diagrams which contribute to the renormalisation of the mass coefficient r. iii) Diagrams which renormalise the interaction vertex w, which is ultimately found to be an irrelevant variable. iv) Diagrams which renormalise the disorder vertex g.*

the other equations arising from this new vertex. The resulting diagrams are shown in Fig. 3.5.

The replica-symmetric RG equations are as follows, again using the substitution $I_0 = 1/(1 + \overline{r})$:

$$\frac{\mathrm{d}r}{\mathrm{d}l} = 2\overline{r} + 2I_1\overline{w} - I_0 g, \tag{3.51}$$

$$\frac{\mathrm{d}I_0}{\mathrm{d}l} = (I_0^2 g - 2)I_0 + 2I_0^2, \tag{3.52}$$

$$\frac{\mathrm{d}g}{\mathrm{d}l} = (4 - d)g + 4g^2 I_0^2 - 4I_1 g\overline{w}, \tag{3.53}$$

$$\frac{\mathrm{d}\overline{\gamma}_1}{\mathrm{d}l} = (2 - d_z)\overline{\gamma}_1 + I_0^2 \overline{\gamma}_1 g, \tag{3.54}$$

$$\frac{\mathrm{d}\overline{\gamma}_2}{\mathrm{d}l} = (2 - 2d_z)\overline{\gamma}_2 + I_0^2(I_0 \overline{\gamma}_1^2 + \overline{\gamma}_2)g, \tag{3.55}$$

$$\frac{\mathrm{d}\overline{w}}{\mathrm{d}l} = (4 - d - d_z)\overline{w} - (I_2 + 4I_3)\overline{w}^2. \tag{3.56}$$

We see that $\overline{w}$ is still an irrelevant variable in the presence of disorder so we

neglect it from here on in. For our current purposes we can also neglect all of the frequency dependence since in the absence of w, the frequency terms decouple from the disorder and mass coefficients and so they have no effect on the replica structure, the main focus of this chapter. We'll see in Chapter 4, however, that inclusion of the frequency dependence is in fact crucial to an accurate description of certain thermodynamic properties of the phase and the frequency terms are retained in the numerical solutions of the RG equations used to generate the phase diagrams shown in Figs. 3.7 and 3.11.

The variance of the disorder distribution g still diverges everywhere under the RG, meaning this is not an appropriate variable to use to distinguish the different phases present. Instead, following Refs. [135, 136], we employ the relative variance $\lambda = I_0^2 g \approx (\overline{r^2} - \overline{r}^2)/\overline{r}^2$, the RG equation for which is given by:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}l} = (4I_0 - d)\lambda + 6\lambda^2. \tag{3.57}$$

This new variable no longer diverges uncontrollably. By inspection, the fixed point solution is $\lambda^* = (d - 4I_0)/6$. Since Eq. 3.52 tells us that I_0 is an irrelevant variable, we can set it to zero in the equation for λ^* and we find that there is a fixed point at $\lambda^* = d/6$ corresponding to a threshold below which the relative variance renormalises to zero and above which it diverges. The relative variance λ is therefore capable of distinguishing between the Mott insulator ($\lambda \to 0$ and $\overline{r} \to \infty$) and the Bose glass ($\lambda \to \infty$).

The flow diagram in the (I_0, λ) plane is shown in Fig. 3.6. It exhibits three fixed points: $P_{MI} = (0,0)$ is the attractive fixed point corresponding to the Mott insulator phase, $P_{MI/SF} = (1,0)$ is an unstable fixed point corresponding to the clean-case MI-SF transition and $P_{MI/BG} = (0, d/6)$ is the unstable fixed point corresponding to the MI-BG transition in the presence of disorder. Even for infinitesimal disorder, the flow is away from $P_{MI/SF}$ and towards $P_{MI/BG}$.

The strong-coupling approach does not capture the BG-SF transition: there is no corresponding fixed point and this transition is beyond the reach of our calculation. For sufficiently large values of hopping the system will always transition back into a superfluid with long-range order, but the method presented here can only access the transition out of the Mott insulator.

The relative variance λ diverges in the Bose glass phase. Using this criterion, we can generate the phase diagram: we know the dependence of I_0 and λ on the microscopics of the system, so we can plot the phase diagram in the $(t/U, \mu/U)$ plane and establish which regions are characterised by a diverging relative variance (identifying this with the BG) and in which regions it flows to zero (MI). The result is show in Fig. 3.7. We see that the Mott lobe shrinks commensurately with the disorder strength, illustrating why unbounded disorder (or any $\delta > 0.5$) is a problem for our calculation, as it would destroy the Mott lobes entirely.
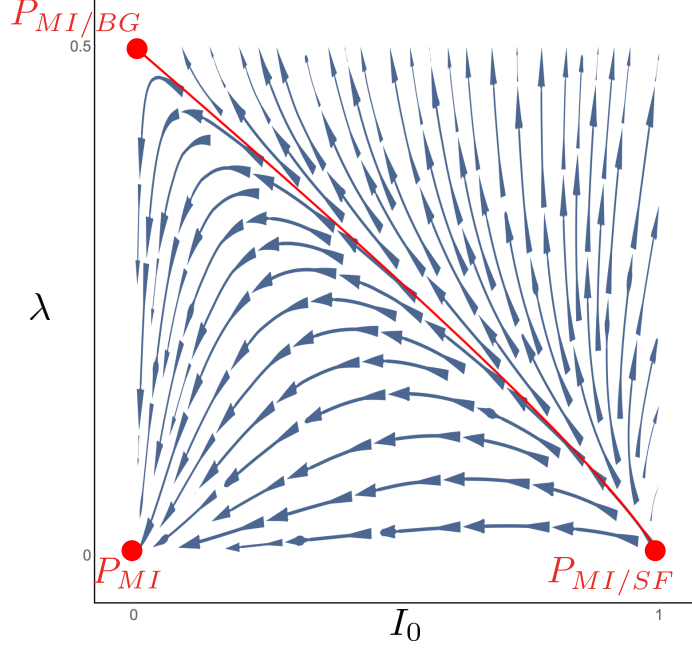
76

Figure 3.6 *The renormalisation group flow diagram in $d = 3$ of the replica-symmetric system in the (I_0, λ) plane, with fixed points P_{MI} at $(0,0)$, $P_{MI/SF}$ at $(1,0)$ and $P_{MI/BG}$ at $(0, d/6)$.*

3.4.2 Replica Symmetry Breaking

Characteristic of glassy systems, replica symmetry breaking (RSB) is the property whereby the coefficient of the replica-mixing vertex (here denoted g) is not a constant, but itself depends on the replica indices. This properly is indicative of the non-ergodic nature of glassy systems and is linked to the breakdown of self-averaging. RSB is essentially a measure of the free energy landscape of a system, and can be understood in terms of long-time correlation functions. If a system exhibits RSB, its free energy landscape is characterised by deep valleys separated by infinitely tall energy barriers into which the system can 'freeze' and even over an infinite timescale will never escape from.

Motivated by previous work on the analogous fermionic Anderson glasses [162] and on random-mass disordered ferromagnets [163, 164] and in the knowledge that renormalisation of a disordered system without allowing for replica symmetry breaking is a dangerous game [41], we now go on to test for the presence of replica symmetry breaking in the disordered Bose-Hubbard model.

Following the general Parisi hierarchical form [70, 74], we set $g \rightarrow g_{\alpha\beta}$ and the action becomes:
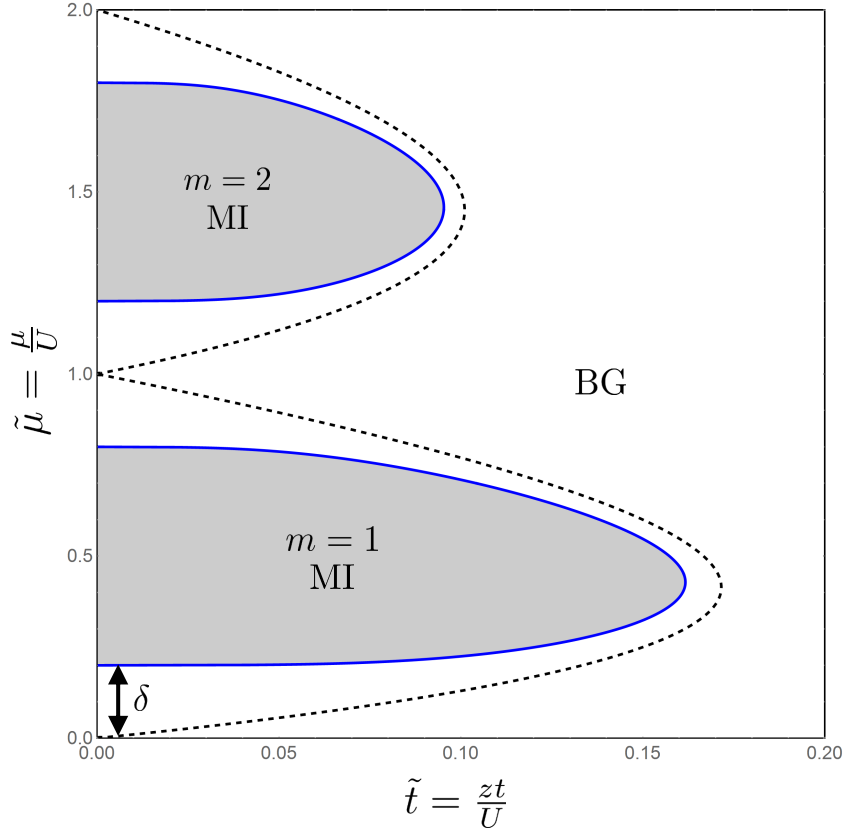
Figure 3.7 *A renormalisation group calculation of the MI-BG phase boundary in $d = 3$ with a disorder strength of $\delta = \Delta/U = 0.2$ using the replica-symmetric field theory. The black dashed line is the mean-field clean-case phase boundary for comparison.*

$$S = \sum_{\alpha} \int_{\mathbf{k},\omega} (k^2 - i\overline{\gamma_1}\omega + \overline{\gamma_2}\omega^2 + \overline{r})\psi_\alpha^\dagger(\mathbf{k},\omega)\psi_\alpha(\mathbf{k},\omega) \tag{3.58}$$

$$- \frac{1}{2}\sum_{\alpha,\beta} g_{\alpha\beta} \int_{\mathbf{k}_i,\omega,\omega'}^{'} \psi_\alpha^\dagger(\mathbf{k}_1,\omega)\psi_\alpha(\mathbf{k}_2,\omega)\psi_\beta^\dagger(\mathbf{k}_3,\omega')\psi_\beta(\mathbf{k}_4,\omega') + ...$$

Physically, $g_{\alpha\beta}$ can be loosely thought of as a new coefficient that measures the overlap between different replicas [74]. Replica symmetry breaking occurs when $g_{\alpha\beta}$ varies for different choices of α and β, i.e. some replicas have small (or even zero) overlap with each other, corresponding to the suppression of tunnelling between different free energy minima, resulting in glassy freezing.

In the Parisi scheme the parameter $g_{\alpha\beta}$ takes the form of a blockwise-diagonal matrix (as shown in Fig. 3.8) where the values along the diagonal are equal to the variance of the induced mass distribution $\tilde{g} = (\overline{r^2} - \overline{r}^2)$. In a replica symmetric
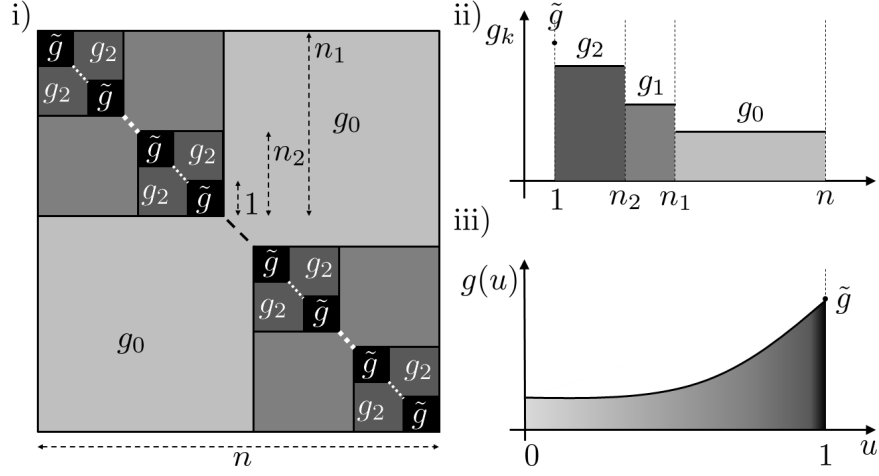
Figure 3.8 *i) Schematic representation of Parisi hierarchical replica symmetry breaking, showing how the matrix $g_{\alpha\beta}$ breaks into a blockwise-diagonal structure. ii) The values are largest along the diagonal and get progressively smaller towards the edges. iii) In the $n \to 0$ replica limit, the series of steps g_k becomes becomes reflected in the $u = 1$ line and contracted such that it becomes a continuous function defined only on the interval $u \in [0, 1]$. The hierarchical structure of the Parisi ansatz requires that $g(u)$ be a monotonically increasing function with $g(1) = \tilde{g}$. The function need not be continuous, but is shown as such for illustrative purposes.*

system, the off-diagonal entries are also equal to $\tilde{g}$, however in the presence of replica symmetry breaking the matrix $g_{\alpha\beta}$ adopts the hierarchical structure shown in Fig. 3.8 where the off-diagonal entries get progressively smaller the further away from the diagonal they are.

Replica Limit

In the replica limit $n \to 0$, the matrix $g_{\alpha\beta}$ becomes a monotonically increasing function $g(u)$ defined on the interval $u \in [0, 1]$ where $g(1) \equiv \tilde{g}$ is the coefficient of the replica diagonal terms. We can then perform functional RG on $g(u)$ and extract its form at the MI-BG transition.

In taking the replica limit, we make use of the following standard identities

for Parisi matrices, denoting $g_{\alpha=\beta}$ by $\tilde{g}$:

$$g_{\alpha\neq\beta} \to g(u), \tag{3.59}$$

$$\sum_{\alpha\beta} g_{\alpha\beta} = \tilde{g} + (n-1) \sum_{\alpha\neq\beta} g_{\alpha\neq\beta}$$

$$\to \tilde{g} + \int_0^1 g(u)\mathrm{d}u, \tag{3.60}$$

$$\sum_{\alpha\beta} g_{\alpha\beta}^2 = \tilde{g}^2 + (n-1) \sum_{\alpha\neq\beta} g_{\alpha\neq\beta}^2$$

$$\to \tilde{g}^2 + \int_0^1 g(u)^2\mathrm{d}u, \tag{3.61}$$

$$\sum_\gamma g_{\alpha\gamma}g_{\gamma\beta} \to (\tilde{c}; c(u)), \tag{3.62}$$

$$\tilde{c} = \tilde{g}^2 - \int_0^1 \mathrm{d}u\, g(u)^2, \tag{3.63}$$

$$c(u) = 2\left(\tilde{g} - \int_0^1 \mathrm{d}v\, g(v)\right) g(u) - \int_0^u \mathrm{d}v\, [g(u) - g(v)]^2. \tag{3.64}$$

Further details on Parisi matrix algebra are included in Appendix C.

Renormalisation Group Equations

We can repeat the same renormalisation group analysis as before, now allowing for the more complicated replica structure. We now have separate renormalisation group equations for the diagonal entry of the Parisi matrix $\tilde{g}$ and for the function $g(u)$ which parameterises the off-diagonal components of the Parisi matrix.

The RG equations for the bare variables become:

$$\frac{\mathrm{d}r}{\mathrm{d}l} = 2r - I_0\tilde{g} - \tilde{g}\int \Omega(\omega) + \int \Omega(\omega)\int_0^1 g(u)\mathrm{d}u, \tag{3.65}$$

$$\frac{\mathrm{d}\tilde{g}}{\mathrm{d}l} = (4-d)\tilde{g} + \rho\left(\tilde{g}^2 - \int_0^1 \mathrm{d}v g(v)^2\right) + 4\tilde{g}^2 I_0^2, \tag{3.66}$$

$$\partial_l g(u) = (4-d)g(u) + 2(g(u) + \tilde{g})g(u)I_0^2,$$

$$\frac{\mathrm{d}\overline{\gamma}_1}{\mathrm{d}l} = (2 - d_z)\overline{\gamma}_1 + I_0^2\overline{\gamma}_1\tilde{g}, \tag{3.67}$$

$$\frac{\mathrm{d}\overline{\gamma}_2}{\mathrm{d}l} = (2 - 2d_z)\overline{\gamma}_2 + I_0^2(I_0\overline{\gamma}_1^2 + \overline{\gamma}_2)\tilde{g}, \tag{3.68}$$

$$+ \rho\left(2\left(\tilde{g} - \int_0^1 \mathrm{d}v g(v)\right)g(u) - \int_0^u (g(u) - g(v))^2\right), \tag{3.69}$$

where $\Omega(\omega) = \Omega(\mathbf{k} = 1, \omega)$ and again we make use of the rescalings $I_0 = 1/(1+r)$,

$\tilde{\lambda} = I_0^2 \tilde{g}$ and $\lambda(u) = I_0^2 g(u)$. The inclusion of RSB does not affect the frequency structure to any order, as only $\tilde{g}$ enters into the RG equations for the frequency coefficients. The RG equations for the rescaled variables are:

$$\frac{\mathrm{d}I_0}{\mathrm{d}l} = (\tilde{\lambda} - 2)I_0 + 2I_0^2 + \rho(\tilde{\lambda} - \langle\lambda\rangle), \tag{3.70}$$

$$\frac{\mathrm{d}\tilde{\lambda}}{\mathrm{d}l} = (4I_0 - d)\tilde{\lambda} + 6\tilde{\lambda}^2 + 2\rho\left[\tilde{\lambda}(\tilde{\lambda} - \langle\lambda\rangle) + 2(\tilde{\lambda}^2 - \langle\lambda^2\rangle)\right], \tag{3.71}$$

$$\partial_l \lambda(u) = (4I_0 - d)\lambda(u) + 2(2\tilde{\lambda} + \lambda(u))\lambda(u)$$
$$+ 2\rho\left[5\lambda(u)(\tilde{\lambda} - \langle\lambda\rangle) - 2\int_0^u \mathrm{d}v(\lambda(u) - \lambda(v))^2)\right], \tag{3.72}$$

where $\langle\lambda\rangle = \int_0^1 \lambda(u)$.

RSB Fixed Point

To determine the replica structure at the MI-BG fixed point, we can follow the method outlined in Refs. [165, 166] and differentiate Eq. 3.72 with respect to u. We find that the only two possible solutions are $\lambda'(u) = 0$ or $(4I_0 - d) + 4(\tilde{\lambda} + \lambda(u)) + 10\rho(\tilde{\lambda} - \langle\lambda\rangle) - 8\rho(u\lambda(u) - \int_0^u \lambda(v)\mathrm{d}v) = 0$. Differentiating the second solution again with respect to u leads to the conditions:

$$4\lambda'(u)(1 - 2\rho u) = 0, \tag{3.73}$$

$$\rightarrow u_c = \frac{1}{2\rho} \text{ or } \lambda'(u \neq u_c) = 0. \tag{3.74}$$

The only solutions to the fixed point equation for $\lambda(u)$ are either a constant function, or a step-like function with the step at u_c, as shown in Fig. 3.9.

Refs. [163, 164] considered random-mass disorder in ferromagnets and employed a very similar analysis but did not fix the position of the step, which led to the non-universal behaviour they found in their critical exponents. Repeating this analysis and fixing the step position in the same manner as in Eq. 3.74 shows that both of their non-universal fixed points in fact reduce to a single RSB fixed point which gives universal exponents, however this fixed point is unphysical. For more details, see Appendix E.

If we allow the $\lambda(u)$ to be a step-like function with a single step at u_c, we
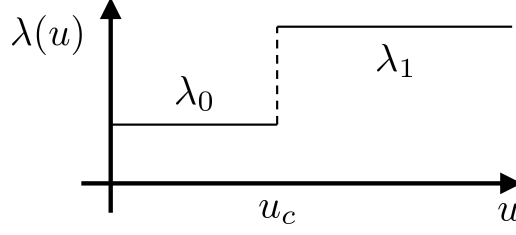
Figure 3.9 *An illustration of the form of the Parisi function $\lambda(u)$ after solving using Eq. 3.74. It becomes a step function with the first part of the step denoted λ_0, the second part denoted λ_1 and the step located at $u_c = 1/(2\rho_0)$.*

find the following RG equations:

$$\frac{d\lambda_0}{dl} = (4I_0 - d)\lambda_0 + 2(2\tilde{\lambda} + \lambda_0) + 10\rho\lambda_0(\tilde{\lambda} - u_c\lambda_0 - (1 - u_c)\lambda_1), \qquad (3.75)$$

$$\frac{d\lambda_1}{dl} = (4I_0 - d)\lambda_1 + 2(2\tilde{\lambda} + \lambda_1)$$
$$+ 10\rho\lambda_1(\tilde{\lambda} - u_c\lambda_0 - (1 - u_c)\lambda_1) - 4\rho u_c(\lambda_1 - \lambda_0)^2, \qquad (3.76)$$

$$\frac{d\tilde{\lambda}}{dl} = (4I_0 - d)\tilde{\lambda} + 6\tilde{\lambda}^2 + 2\rho\left[\tilde{\lambda}(\tilde{\lambda} - u_c\lambda_0 - (1 - u_c)\lambda_1)\right.$$
$$\left. + 2(\tilde{\lambda}^2 - u_c\lambda_0^2 - (1 - u_c)\lambda_1^2)\right], \qquad (3.77)$$

where λ_0 is the first piece of the step and λ_1 is the second, as shown in Fig. 3.9, with $\lambda_0 < \lambda_1$. Numerically, we find that the only solutions are those for which $\lambda_0 = 0, \lambda_1 \equiv \lambda = \tilde{\lambda}$. We're left with a single RG equation:

$$\frac{d\lambda}{dl} = (4I_0 - d)\lambda + 6\lambda^2(1 + \rho u_c) \qquad (3.78)$$

$$= (4I_0 - d)\lambda + 9\lambda^2 \qquad \text{using } u_c = \frac{1}{2\rho}. \qquad (3.79)$$

The solutions to this are:

$$\lambda^* = 0 \qquad \text{or} \qquad \lambda^* = \frac{d - 4I_0}{9}. \qquad (3.80)$$

So, for $I_0 = 0$, we find the new RSB fixed point solution is $\lambda^* = d/9$, in contrast to the replica symmetric $\lambda^* = d/6$ fixed point found in previous work.

Phase and Flow Diagrams

The new RSB fixed point lies at a lower disorder strength than the RS fixed point ($d/9$ rather than $d/6$), already suggesting that this fixed point should pre-
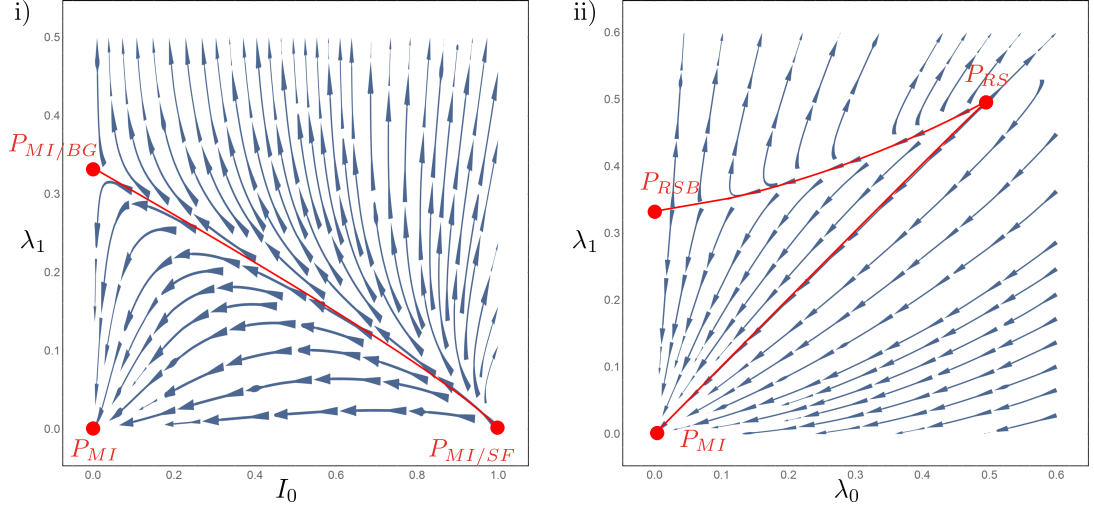
82

Figure 3.10 *Renormalisation group flow diagrams in the replica-symmetry broken case, with approximate separatrices drawn as guides to the eye. i) RG flow in the (I_0, λ) plane showing $P_{MI} = (0,0)$, $P_{MI/SF} = (1,0)$ and the new $P_{RSB} = (0, d/9)$ which occurs at a lower disorder strength than the previous replica symmetric fixed point shown in Fig. 3.6. ii) RG flow in the (λ_0, λ_1) plane showing the instability of the system towards 1-step RSB. The replica symmetric fixed point is shown on the diagonal at $(d/6, d/6)$ and the clean fixed point at $(0,0)$. Any perturbation away from $\lambda_0 = \lambda_1$ results in flow away from the replica symmetric fixed point. For $\lambda_0 > \lambda_1$, the flow is to the clean fixed point as this unphysical disorder distribution renormalises to zero. For $\lambda_1 > \lambda_0$, the flow is controlled by a new RSB fixed point at $(0, d/9)$.*

empt the replica symmetric transition and be the physical one that controls the phase transition, but to be sure of this we must conduct further analysis. Figure 3.10 shows renormalisation group flow diagrams in the (λ_0, λ_1) plane and in the (I_0, λ_1) plane.

Plotting the flow in the (λ_0, λ_1) plane, we see that for $\lambda_0 = \lambda_1$ (the replica symmetric solution along the diagonal), the situation exactly reproduces the replica symmetric model with an MI-BG fixed point at $(d/6, d/6)$. If we allow any infinitesimal perturbation such that $\lambda_0 \neq \lambda_1$ we find that the system flows to a new RSB fixed point at $(0, d/9)$. Note that for the unphysical situation of $\lambda_0 > \lambda_1$ (everything below the diagonal) the system flows to the Mott insulating fixed point, indicating that this unphysical type of disorder renormalises to zero and has no effect.

Plotting the phase diagram for the situation of 1-step RSB (Fig. 3.11), we find that it is almost identical to the replica symmetric case with only exceedingly minor quantitative differences. This means that experimentally testing for RSB cannot be achieved by mapping out the phase diagram alone. The phase diagram
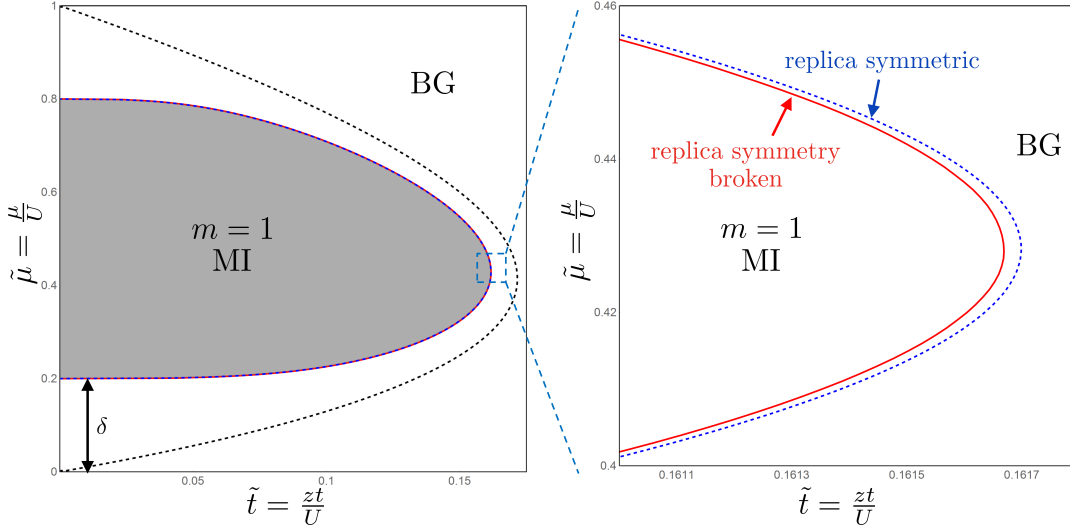
Figure 3.11 *The phase diagram for the disordered Bose-Hubbard model in $d = 3$, showing the negligible difference between the RS and RSB phase boundaries. Since this difference is quantitative and extremely small, experimental determination of the phase diagram alone is insufficient to distinguish whether the physical system exhibits replica symmetry breaking.*

in Fig. 3.11 differs from that shown in Ref. [120]. The latter is incorrect. An arithmetic error led to a rescaling of the hopping axis such that it exaggerated the difference between the replica symmetric and replica symmetry broken phase boundaries. This error has no effect on the conclusions stated in Ref. [120], as the phase boundaries were already considered to be sufficiently close together that there was no practical method for distinguishing the two. There is also a further quantitative shift in boundaries due to including the frequency terms in the renormalisation group equation used to plot Fig. 3.11 which were not retained in Ref. [120].

Additionally, the correlation length exponent ν turns out to be the same for both RS and RSB situations. We can linearise the RG equations around the RSB fixed point $P_{\mathrm{MI/BG}}^{rsb}$, yielding $\delta\lambda_0'(\ell) = d\cdot\delta\lambda_0$ for a deviation $\delta\lambda_0 = \lambda_0 - d/9$ from the critical point, which gives rise to the same correlation length exponent $\nu = 1/d$ as in the RS situation [136]. Other experimentally measurable parameters are therefore required.

Stability of RSB

While our results show that the disordered Bose-Hubbard model displays an inherent and seemingly strong instability towards 1-step RSB, our calculation cannot prove that the requisite perturbation exists to physically break replica symmetry. Indeed, a rigorous proof of this seems impossible, though this is

no different to every other case of spontaneous symmetry breaking in physics. It has, however, been argued [164, 167] that such replica symmetry breaking perturbations exist in any disordered system, and our results from Chapter 4 certainly suggest this to be the case.

Though there is some evidence that 1-step RSB is sufficient to capture qualitative effects of glassy physics [168], no spin glass systems exhibit this type of exact 1-step RSB. This suggests that to 2-loop or higher order, the nature of the replica symmetry breaking would change and perhaps become something more complex. It is unlikely that the inclusion higher-order vertices would undo the finding of RSB, although without explicitly calculating their contributions it is impossible to know for sure.

The accuracy of the replica symmetric approach (in terms of the phase diagram and scaling of the correlation length exponent) is not unprecedented in the literature. For example, the Cardy-Ostlund random-field XY model in $d = 2$ also displays an instability towards replica symmetry breaking [41] but certain properties can be accurately computed within a replica symmetric theory [169].

3.5 Order Parameters and Experimental Tests

In Chapter 4 we look more closely at the effects of replica symmetry breaking on bulk thermodynamic probes, but here let us look at the more conventional glassy correlation functions known as the Edwards-Anderson order parameters. The Edwards-Anderson order parameters are a well-established piece of theoretical lexicon from the spin glass community [70, 75].

By analogy with spin glass systems, the Edwards-Anderson order parameters take the following form in the Bose glass:

$$q_{\text{EA}} = \lim_{\tau \to \infty} \left(\overline{\langle \rho_i(\tau) \rho_i(0) \rangle} - \overline{\langle \rho_i(\tau) \rangle} \; \overline{\langle \rho_i(0) \rangle} \right), \qquad (3.81\text{a})$$

$$q = \overline{\langle \rho_i(\tau) \rangle^2} - \overline{\langle \rho_i(\tau) \rangle}^2. \qquad (3.81\text{b})$$

where we define the density $\rho_i(\tau) = I_0^{-1} |\psi_i(\tau)|^2$ such that $\langle \rho_i(\tau) \rangle = \rho_0 \sim m$ in the MI state. In the Mott insulator, both order parameters are zero, but both become finite in the Bose glass [126].

The first order parameter is the long-time density-density correlation function, while the second is an equal-time density-density correlation function. They can be recast in the language of replicas, as discussed in more detail in Chapter 1, where we saw that q_{EA} measures long-time correlations within a single replica and q measures equal-time correlations between different replica systems. In a replica symmetric system, $q = q_{\text{EA}}$, but the two order parameters will become different once replica symmetry is broken. This corresponds to the free energy

landscape becoming truly 'glassy', i.e. ergodicity is non-trivially broken and thermal averages are no longer equal to long-time averages. The system can become stuck in local minima for thermodynamically long timescales and so whether or not the system can tunnel between replicas determines whether the order parameters are equal to each other or not.

We may define the quantity $\Delta = q_{\mathrm{EA}} - q$ to quantify the degree of ergodicity breaking [70]. Using the one-loop one-step RSB solution, we obtain $q_{\mathrm{EA}} = \rho_0^2 \tilde{\lambda}$ and $q = \rho_0^2 \langle \lambda \rangle$. For the present one-step RSB solution we obtain $\Delta = \rho_0(\lambda_1 - \lambda_0)/2 = \rho_0 \lambda_1 / 2$.

Experimentally measuring these order parameters is challenging. Though a proposal exists for measuring replica correlation functions on physical replica systems using ultracold atomic gases [126], this experiment has never been performed. At the time when the work in this chapter was originally performed, there seemed to be no easily accessible experimental method to test for the presence of replica symmetry breaking in ultracold atomic gases. The long-time correlations in particular are impossible to perform in an ultracold atomic gas due in part to the short lifetime of current traps, but more so to the fact that all suitable measurement techniques are destructive.

Further consideration of potential ways to test for replica symmetry breaking led to the development of the work presented in Chapter 5 where we use mean-field numerical simulations to reproduce the experimental results achieved by single-atom resolution fluorescence imaging in quantum gas microscopes and show that these setups are able to measure the Edwards-Anderson order parameter in the form suggested by Ref. [126]. It may be possible to experimentally test for replica symmetry breaking using this method, however our mean-field simulations are not sufficiently complex to be able to take this possibility into account. We shall show in Chapter 4 that replica symmetry breaking is vital for our model to correctly describe the thermodynamics of a disordered Bose-Hubbard-like system, but before we get to that we can perform a further check on the feasibility of replica symmetry breaking in the Bose glass.
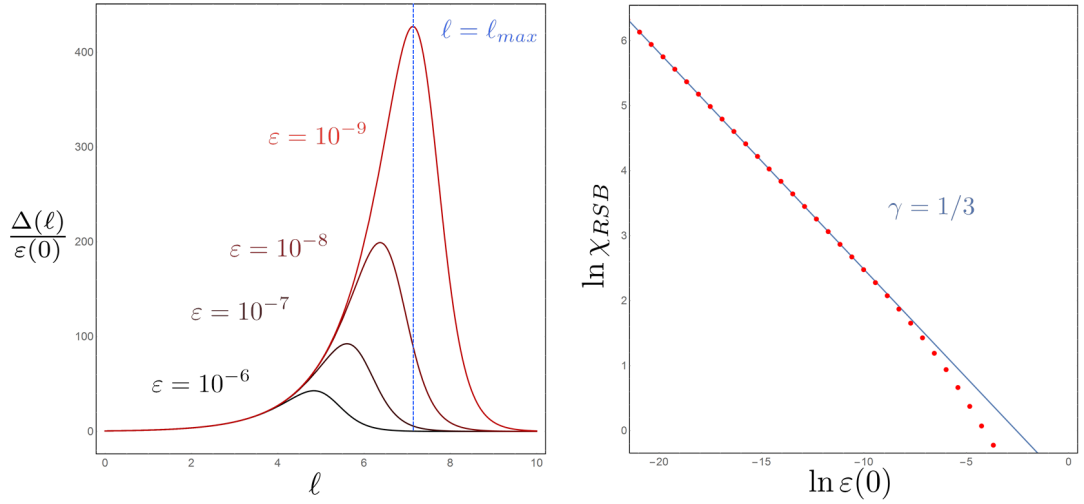
Figure 3.12 *i) RG flow of the degree of ergodicity breaking $\Delta(\ell)$ for four different values of initial RSB perturbation $\varepsilon(0)$. The limit of $\varepsilon(0) \to 0$ corresponds to approaching the RS fixed point from the MI side. The divergence of $\Delta(\ell)$ as $\varepsilon(0) \to 0$ is indicative of a diverging susceptibility χ_{RSB}. ii) Susceptibility χ_{RSB} plotted against $\varepsilon(0)$, confirming that it diverges as $\varepsilon(0)^{-1/3}$ at the transition.*

3.6 Susceptibility to RSB

So far, the renormalisation group analysis shows a clear instability to RSB in the Bose glass phase, however the RG flow of the disorder strength in this regime diverges uncontrollably and a stable Bose glass fixed point exhibiting RSB is not accessible with this strong-coupling calculation. Following the analysis of Ref. [41], we calculate a susceptibility towards an infinitesimal RSB perturbation in the Mott insulating regime where an exact weak-coupling fixed point exists and the vanishing disorder vertex under RG ensures that the effects of disorder remain small across all length scales.

We define the susceptibility to a replica symmetry breaking perturbation as:

$$\chi_{RSB} = \lim_{\substack{\ell \to \infty \\ \varepsilon \to 0}} \frac{\Delta(\ell)}{\varepsilon(0)}, \tag{3.82}$$

where $\Delta(\ell)$ is the degree of ergodicity breaking defined previously and $\varepsilon(0)$ is some infinitesimal perturbation away from replica symmetry such that the limit $\varepsilon(0) \to 0$ corresponds to approaching the replica symmetric fixed point from within the Mott insulating regime.

In computing χ_{RSB} directly, both limits $\varepsilon(0) \to 0$ and $\ell \to \infty$ must be taken simultaneously, which is a somewhat ambiguous process. A more well-defined procedure is to instead evaluate the susceptibility at the point ℓ_{max} where $\Delta(\ell)/\varepsilon(0)$ is maximum. As shown in Fig. 3.12, both the height and position of

this maximum diverge as $\varepsilon(0) \to 0$ and we find a scaling form of the susceptibility $\chi_{RSB} \sim \varepsilon(0)^{-\gamma}$ with $\gamma = 1/3$ in $d = 3$. This generalises to a scaling form of $\gamma = 1/d$ for all $d \geq 3$.

The divergent susceptibility to replica symmetry breaking even in the weak-coupling regime where the renormalisation procedure is well-controlled confirms that replica symmetry breaking is a physical feature of this model, that any infinitesimal perturbation away from perfect replica symmetry will inevitably lead to replica symmetry breaking and that the replica symmetry broken fixed point is the one which controls the physical phase transition.

3.7 Discussion

The presence of replica symmetry breaking in the Bose glass has been controversial, and there are a number of commonly raised objections which bear further examination.

- **Where does the ultrametricity associated with RSB come from in this model?** - In conventional spin glass literature, the existence of a fragmented free energy manifold with many local minima stems from the inability of the system to uniquely satisfy all bonds, e.g. the geometric frustration shown in Fig. 1.6. In the Bose-Hubbard model, the glassiness enters in terms of boson number density but there's no clear picture one can adopt to understand where the fragmented free energy landscape truly comes from. Instead, we must appeal to the well-established mapping between bosons and spin systems (Ref. [10] and which we shall encounter in detail in Chapter 4) where we see geometrical frustration emerge naturally from the type of disorder we consider here.

- **Do RSB perturbations really exist?** - This analysis can never truly say whether RSB perturbations exist or not, only that if any such infinitesimal perturbation away from replica symmetry does occur, the system will make use of it. Though there are arguments for why RSB perturbations should occur in disordered systems [164, 167], there is no rigorous proof. However, there is no rigorous proof of any form of spontaneous symmetry breaking - just as one can supercool water below its freezing point if one is sufficiently careful, it may well be possible to encounter a disordered system that (despite divergent RSB susceptibility) somehow does not break replica symmetry. Our analysis, however, suggests that such a replica-symmetric fixed point would be highly unstable.

- **Could RSB arise as an artefact of an incorrect approximation?** - This is a particularly interesting possibility. Applying renormalisation group techniques to a disordered system is a matter of some subtlety due to the implicit dependence of short length scales upon the long [41]. Could

it be that allowing for RSB is still not enough to correct the problem of renormalising a glassy system, or even that this calculation somehow enforces RSB upon the system? This is possible, but in light of other works which have used similar methods and found stable replica-symmetric solutions (e.g. Ref. [170] and Appendix D), it seems unlikely. Additionally, the divergent RSB susceptibility is a strong indication that the calculation is realistic. The most consistent story currently available is that RSB is a genuine physical feature of the system.

- **Does RSB really matter, given that the phase boundary is virtually unchanged and it's not clear what it does?** - Though the lack of some immediately obvious effect of RSB on the phase diagram or correlation length exponents was disappointing, we shall see in Chapter 4 that allowing for RSB can have a dramatic effect on certain thermodynamic properties which were not considered in this chapter.

In this chapter, we have shown the existence of replica symmetry breaking in the Bose glass and demonstrated that the replica symmetric system is highly unstable to any infinitesimal RSB perturbation. The breaking of replica symmetry is linked to the breaking of ergodicity in the system, and we have demonstrated the existence of glassy Edwards-Anderson order parameters which can in principle be measured to test for RSB. We shall see in Chapter 5 that the thermally averaged order parameter q is accessible by quantum gas microscopes, however the long-time correlation function q_{EA} is not. Motivated by the need to find an alternative experimental environment which demonstrates the same physics and allows measurement of temporal correlation functions, we now turn to dimerised quantum antiferromagnets in Chapter 4.

Chapter 4

Bose and Mott Glasses in Dimerised Antiferromagnets

As part of the previous work on the breaking of replica symmetry in the Bose glass which was presented in Chapter 3, it became apparent that the measurements required to test our predictions were possible in principle but often highly impractical to conduct on ultracold atomic gases. Most imaging measurements are destructive, meaning quantities such as two-time correlation functions are impossible to measure on a single system.

Dimerised quantum antiferromagnets (AFMs) provided an alternative environment in which to explore the same physics. Motivated by experiments on bromine-doped dichloro-tetrakis-thiourea-nickel (DTN) [171], the work presented in this chapter is an investigation into the properties of disordered dimerised quantum AFMs seen through the lens of the hard-core Bose gas formalism.

It turned out that this would lead to a very clear and dramatic demonstration of the importance of replica symmetry breaking in the Bose glass. Without including RSB, the field theory formalism presented in the previous chapter does not correctly recover a finite compressibility in the Bose glass. With RSB, we find the compressibility becomes finite at the MI-BG transition, but also find that the compressibility is strongly suppressed at the multicritical point at the tip of the Mott lobe. We identify this anomalous suppression of compressibility with a rare Mott glass phase never before analytically predicted to exist in this model.

The work in this chapter was published in "Bose and Mott glasses in dimerized quantum antiferromagnets," S. J. Thomson and F. Krüger, *Physical Review B: Rapid Communications* **92**, 180201(R) (2015) [154].

4.1   Spin Dimer Systems

Bose-Einstein condensation of the spin degrees of freedom in magnetic insulators is a well-established phenomenon, having been observed in a wide variety of clean materials such as $TlCuCl_3$ [172–174], Cs_2CuCl_4 [175, 176], $BaCuSi_2O_6$ [177, 178], $SrCu_2(BO_3)_2$ [179, 180], $NiCl_2$-$4SC(NH_2)_2$ [181] and $Ba_3Mn_2O_8$ [182–184].

The central idea is that a dimer (i.e. a pair of spins) can adopt either a spin-0 singlet configuration or any of three spin-1 triplet configurations. Loosely speaking, one can treat the singlet as a 'vacuum' and the triplets as three distinct species of boson. Similar bosonisation methods have been explored in various contexts, such as the Matsuda-Matsubara transform [185], Schwinger bosons (see, e.g. Ref. [9]) or the Holstein-Primakoff formalism [186]. The latter has been used to describe effects seen in dimerised quantum antiferromagnets [171], however its reliance on expanding around a broken symmetry state seems inherently unsatisfactory. Here in this chapter we use a bond operator technique derived independently of but exactly equivalent to that presented in Refs. [187–191], which does not rely on an expansion around a broken-symmetry state.

The bosonic excitations out of the spin-0 vacuum are known as triplons, however they are sufficiently similar to the magnetic excitations seen in ordered antiferromagnets (known as magnons) that in the majority of the literature the two terms are used interchangeably [192]. This distinction becomes further muddled in a glassy state with locally ordered regions so for the purposes of this thesis I will also use the terms 'magnon' and 'triplon' interchangeably.

Just as regular bosons can form a Bose-Einstein condensate, so too can these triplons condense in a magnetic system (for reviews, see Refs. [192, 193]). Typically the spins of the triplons are taken as being aligned along the z direction and condensation of these triplons corresponds to the establishment of long-range order in the x-y plane. The correspondence between the spin system and the bosonic description is shown in Table 4.1, based on Ref. [192]. The Mott insulating phase of the regular Bose-Hubbard model is instead a singlet state of the dimerised model, while the superfluid phase is a long-range ordered antiferromagnet. In an external field, this is a canted antiferromagnet but in zero applied field this condensed state becomes a perfect Néel antiferromagnet.

As the excitations of the magnetic system can be described by this effective bosonic model, it should come as no surprise that the addition of disorder leads to very similar effects to those presented in Chapter 3. Specifically, a magnetic Bose glass phase has been predicted to occur, characterised by rare patches of magnetic order in an otherwise spin-singlet background. Such phases have been experimentally observed in several materials [194–197], but the case of bromine-doped dichloro-tetrakis-thiourea-nickel (DTN) [171] is perhaps the most interesting. There, the authors find not only evidence for a magnon Bose glass, but also find a glassy phase with vanishing magnetic susceptibility

Correspondence between Bose gas and spin system	
Bose gas	Spin system
Bosons	Triplons (spin excitations)
$\hat{n}_{i\pm}$	$\pm\hat{S}_i^z$
$U(1)$ symmetry	$O(2)$ symmetry
$\langle\psi(\mathbf{r})\rangle$	$\langle\hat{S}_i^x + i\hat{S}_i^y\rangle$
Chemical potential $\tilde{\mu}$	Applied field h
SF density ρ_s	Transverse spin stiffness
MI state	Magnetisation plateau
SF state	Long-range xy-plane order
Compressibility κ	Magnetic susceptibility χ

Table 4.1 *The properties of the hard-core Bose gas with nearest-neighbour interactions can be mapped onto analogous properties of a system of spins. This correspondence will be formally derived in the following sections. Here, $\hat{S}_i^\alpha$ refers to the overall spin of the dimer, not the spin of the individual spin-1/2 particles within the dimer. This table is based on that shown in Ref. [192] and extended to include more parameters.*

(compressibility, in the bosonic language) which they identify with a Mott glass (which we met previously in Table 3.1 in Chapter 3).

4.2   The Mott Glass

The Mott glass is an unusual incompressible glassy phase first analytically predicted in one dimensional fermionic systems [162, 198] using functional renormalisation group techniques. It was then seen in real-space renormalisation group calculations of the $O(2)$ quantum rotor model in one dimension [149, 151, 199, 200] and evidence for this phase has also been seen in quantum Monte Carlo calculations performed on spin models and related Bose-Hubbard-like systems in both one and two dimensions [150, 152, 153, 201]. A related - and very possibly identical - incompressible glassy phase in the $O(2)$ quantum rotor model known as a random-rod glass has been shown to exist only in the presence of particle-hole symmetry [131, 202, 203]. The only experimental observation thus far of a Mott glass is that of Ref. [171], which includes a full quantum Monte Carlo study that backs up their conclusion.

The Mott glass remains a curious phase that has arisen in a variety of different theoretical scenarios in different systems and with different types of disorder. Some work suggests that the Mott glass requires finite-range interactions and could not be stabilised in the regular disordered Bose-Hubbard model; yet other work suggests that finite-range interactions are not necessary but that the essential ingredient is instead off-diagonal disorder [150]. The effects of

dimensionality are not known, other than that evidence for it now exists in one [149, 150, 162, 198–200], two [153, 201] and three [154, 171] dimensions.

There has never been a comprehensive study of the Mott glass itself across all its possible realisations, and in this chapter I do not attempt one. Whatever the root cause of the phase, it seems certain that something not yet fully understood is happening in these systems. In fact, there is reason to be suspicious as to whether all the phases given the name 'Mott glass' are really phases with the same underlying physics or simply phases with similar cosmetic properties but different fundamental physics.

In the Bose-Hubbard model, there is a large amount of evidence that the Mott glass exists at and around the special multicritical points where particle-hole symmetry is restored. In the hard-core boson model which describes the experiments on DTN [171], the zero-field limit where the Mott glass is seen corresponds to one of these multicritical points. Given the proximity of the Mott glass to the these high-symmetry points, it seems very likely that its physics has something to do with the change in the nature of excitations precisely at the tip of the Mott lobes [119].

Motivated by the experimental observation of a Mott glass in a system that can be well-described by a hard-core Bose-Hubbard model with finite range interactions and local disorder [171], in this chapter we derive a strong-coupling replica field theory for this system following the procedure laid out in Chapter 3 and calculate the magnetic susceptibility (or compressibility, in boson language) to see if we see any evidence for a Mott glass in this model.

4.3   Mapping to Hard-Core Bosons

We start from the following Hamiltonian, which is reminiscent of a typical Heisenberg Hamiltonian for dimers in the presence of a magnetic field but with an additional anisotropy term which couples to the square of the z-projection of the total spin of the dimer:

$$\mathcal{H} \;=\; \sum_{ijmn} J_{ijmn}\hat{\mathbf{S}}_{i,n} \cdot \hat{\mathbf{S}}_{j,m} - D\sum_{i}(\hat{S}^z_{i,1} + \hat{S}^z_{i,2})^2 - h\sum_{i}(\hat{S}^z_{i,1} + \hat{S}^z_{i,2}), \quad (4.1)$$

where i, j label the dimers and m, n label the component spins of the dimers. The coefficient J_{ijmn} is always greater than zero and becomes J_0 for the intra-dimer coupling, J_1 for the nearest-neighbour inter-dimer coupling and J_2 for the next-nearest-neighbour inter-dimer coupling, as shown in Fig. 4.1. The anisotropy D
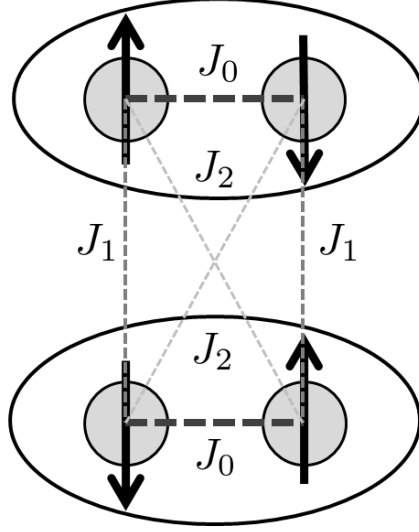
Figure 4.1 *An illustration of the inter-dimer and intra-dimer couplings defined in Eq. 4.1. All couplings are positive (antiferromagnetic).*

is the dimer equivalent of a single-ion anisotropy and is of the standard form used in dimerised magnetic systems.

Using the method first developed in Ref. [187], we map the fermionic spins to bosonic creation and annihilation operators corresponding to singlet and triplet dimer configurations as follows:

$$\hat{s}^\dagger |0\rangle = |0,0\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle), \tag{4.2}$$

$$\hat{t}_0^\dagger |0\rangle = |1,0\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle), \tag{4.3}$$

$$\hat{t}_x^\dagger |0\rangle = -\frac{1}{\sqrt{2}}(|\uparrow\uparrow\rangle - |\downarrow\downarrow\rangle), \tag{4.4}$$

$$\hat{t}_y^\dagger |0\rangle = \frac{i}{\sqrt{2}}(|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle). \tag{4.5}$$

where $|0\rangle$ represents the vacuum. This allows us to rewrite the spin operators as:

$$\hat{S}_1^\alpha = \frac{1}{2}\left[\hat{s}^\dagger \hat{t}_\alpha + \hat{t}_\alpha{}^\dagger - i\epsilon_{\alpha\beta\gamma}\hat{t}_\beta^\dagger \hat{t}_\gamma\right], \tag{4.6}$$

$$\hat{S}_2^\alpha = \frac{1}{2}\left[-\hat{s}^\dagger \hat{t}_\alpha - \hat{t}_\alpha{}^\dagger - i\epsilon_{\alpha\beta\gamma}\hat{t}_\beta^\dagger \hat{t}_\gamma\right], \tag{4.7}$$

where the subscripts 1 and 2 number the spins within a single dimer and $\epsilon_{\alpha\beta\gamma}$ is the Levi-Civita symbol. The Hamiltonian is no longer diagonal in this basis, so

we make a rotation to the following basis:

$$\hat{t}_0^\dagger \ket{0} = \hat{t}_z^\dagger \ket{0} = \ket{1, 0}, \tag{4.8}$$

$$\hat{t}_1^\dagger \ket{0} = \ket{1, 1} = \ket{\uparrow\uparrow}, \tag{4.9}$$

$$\hat{t}_{-1}^\dagger \ket{0} = \ket{1, -1} = \ket{\downarrow\downarrow}. \tag{4.10}$$

The presence of the anisotropy term D allows us to project out the $\ket{1, 0}$ state. As long as $h < D$, this state is energetically unfavourable and can be projected out.

Following the harmonic approximation used in Ref. [188–191], we can use the hard core constraint $\hat{s}^\dagger \hat{s} + \hat{t}_\alpha^\dagger \hat{t}_\alpha = 1$ (where the summation over α is implied) to rewrite the singlet operators as $\hat{s}^\dagger = \hat{s} = \sqrt{1 - \hat{t}_\alpha^\dagger \hat{t}_\alpha}$, allowing us to write a Hamiltonian entirely in terms of the triplet creation and annihilation operators.

$$
\begin{aligned}
\mathcal{H} = & -\sum_{i,\sigma=\pm} \mu_\sigma \hat{n}_{i\sigma} + V \sum_{\langle i,j \rangle} \sum_{\sigma_1,\sigma_2} \sigma_1 \sigma_2 \, \hat{n}_{i\sigma_1} \hat{n}_{j\sigma_2} \\
& + t \sum_{\langle i,j \rangle} \left[(\hat{t}_{i-}^\dagger - \hat{t}_{i+})(\hat{t}_{j-} - \hat{t}_{j+}^\dagger) + \text{H.c.} \right],
\end{aligned}
\tag{4.11}
$$

where $t = (J_1 - J_2)/2$, $V = (J_1 + J_2)/2$ and $\mu_\sigma = (J_0 - D) + \sigma h$. We have also used the hard core constraint to rewrite products of number operators using $\hat{n}_b^2 = \hat{n}_b$ and $\hat{n}_b \hat{n}_c = 0$ and we have shifted the energies by a constant such that the singlet state has zero energy.

We can also do the same mapping without the anisotropy term, leading to an additional contribution coming from the coupling to the $\hat{t}_z$ triplet:

$$
\begin{aligned}
\mathcal{H}_z = & \sum_i J_0 \hat{t}_{i,z}^\dagger \hat{t}_{i,z} + t \sum_{\langle ij \rangle} \left(\hat{t}_{i,z}^\dagger \hat{t}_{j,z} + \hat{t}_{j,z}^\dagger \hat{t}_{i,z} + \hat{t}_{j,z}^\dagger \hat{t}_{i,z}^\dagger + \hat{t}_{i,z}^\dagger \hat{t}_{j,z}^\dagger + \hat{t}_{i,z} \hat{t}_{j,z} + \hat{t}_{j,z} \hat{t}_{i,z} \right) \\
& + V \sum_{\langle ij \rangle} \left(\hat{t}_{i,z}^\dagger \hat{t}_{j,z} \hat{t}_{j-}^\dagger \hat{t}_{i-} + \hat{t}_{i,z}^\dagger \hat{t}_{j,z}^\dagger \hat{t}_{i-} \hat{t}_{j+} + \hat{t}_{i-}^\dagger \hat{t}_{j+}^\dagger \hat{t}_{i,z} \hat{t}_{j,z} + \hat{t}_{i+}^\dagger \hat{t}_{j+} \hat{t}_{j,z}^\dagger \hat{t}_{i,z} \right).
\end{aligned}
\tag{4.12}
$$

This additional contribution contains the same physics as Eq. 4.11, but coupled to the triplet with zero z-projection rather than to the bosonic vacuum. Since this term adds no significant new behaviour but does significantly increase the algebraic complexity, we include the anisotropy term in Eq. 4.1 which allows us to project out these terms.

The form of Eq. 4.11 is that of an extended Bose-Hubbard model with hard-core repulsions. The extended Bose-Hubbard model is well studied and the effect of the finite-range interactions is to stabilise fractionally-filled Mott insulating phases that depend upon the underlying lattice geometry. For example, in a square lattice this repulsion can stabilise half-filled checkerboard and stripe phases, whereas in a triangular lattice it will stabilise fillings of 1/3, 2/3 and so on.

The stabilisation of fractional filling factors adjacent to superfluid regions leads naturally to the question of whether the model exhibits supersolid behaviour at the transition from fractionally-filled Mott lobes to the superfluid states. Many studies indicate this [204–211], though other authors claim that in reality this model does not contain a supersolid phase but instead a first-order transition from Mott insulator to homogeneous superfluid [212].

In its current form, our renormalisation group calculation is not sensitive to the nature of the condensed phase, nor can it distinguish whether the MI-SF transition should be first-order. Should there be a first-order transition in the vicinity of the fractionally-filled lobes, then that may introduce additional factors that supersede our work here. Consequently, in the later parts of this chapter where the behaviour at the tip of the lobe becomes important we focus on the $m = 0$ central Mott lobe as we know our calculation remains well-controlled and captures all of the essential physics in this region.

4.4   Clean Mean Field Theory

Using a Gutzwiller variational wavefunction, we can calculate a mean-field energy expression and determine the ground-state phase diagram by minimising it. The wavefunction we use is given by:

$$|\Psi\rangle = \prod_i |\psi_i\rangle = \prod_i \left(\sqrt{1 - |\alpha_i|^2 - \beta_i|^2} + \sigma_i \alpha_i a_i^\dagger + \sigma_i \beta_i b_i^\dagger \right) |0\rangle , \qquad (4.13)$$

where $\sigma_i = \pm 1$ describes the antiferromagnetic ordering. We can extract a mean-field energy from this wavefunction by calculating $E_{MF} = \langle \Psi | \mathcal{H} | \Psi \rangle$.

$$
\begin{aligned}
E_{MF} = &- \sum_i (\mu_{a,i} |\alpha_i|^2 + \mu_{b,i} |\beta_i|^2) \\
&+ V \sum_{\langle i,j \rangle} \left(|\alpha_i|^2 |\alpha_j|^2 |\beta_i|^2 |\beta_j|^2 - |\alpha_i|^2 |\beta_j|^2 - |\beta_i|^2 |\alpha_j|^2 \right) \\
&+ t \sum_{\langle i,j \rangle} \left(\alpha_i^* \alpha_j + \beta_i^* \beta_j - \alpha_i^* \beta_j - \beta_j^* \alpha_i + H.c. \right) \sqrt{1 - |\alpha_i|^2 - |\beta_i|^2} \sqrt{1 - |\alpha_j|^2 - |\beta_j|^2},
\end{aligned}
$$

$$(4.14)$$

using

$$\langle \Psi | \hat{n}_{i,a} | \Psi \rangle = |\alpha_i|^2, \tag{4.15}$$

$$\langle \Psi | \hat{n}_{i,b} | \Psi \rangle = |\beta_i|^2, \tag{4.16}$$

$$\langle \Psi | \hat{n}_{i,a} \hat{n}_{j,a} | \Psi \rangle = |\alpha_i|^2 |\alpha_j|^2, \tag{4.17}$$

$$\langle \Psi | \hat{n}_{i,b} \hat{n}_{j,b} | \Psi \rangle = |\beta_i|^2 |\beta_j|^2, \tag{4.18}$$

$$\langle \Psi | \hat{n}_{i,a} \hat{n}_{j,b} | \Psi \rangle = |\alpha_i|^2 |\beta_j|^2, \tag{4.19}$$

$$\langle \Psi | \hat{b}_i^\dagger \hat{b}_j | \Psi \rangle = \beta_i^* \beta_j \sqrt{1 - |\alpha_i|^2 - |\beta_i|^2} \sqrt{1 - |\alpha_j|^2 - |\beta_j|^2}, \tag{4.20}$$

$$\langle \Psi | \hat{a}_i^\dagger \hat{a}_j | \Psi \rangle = \alpha_i^* \alpha_j \sqrt{1 - |\alpha_i|^2 - |\beta_i|^2} \sqrt{1 - |\alpha_j|^2 - |\beta_j|^2}, \tag{4.21}$$

$$\langle \Psi | \hat{b}_i^\dagger \hat{b}_j | \Psi \rangle = \beta_i^* \beta_j \sqrt{1 - |\alpha_i|^2 - |\beta_i|^2} \sqrt{1 - |\alpha_j|^2 - |\beta_j|^2}, \tag{4.22}$$

$$\langle \Psi | \hat{a}_j^\dagger \hat{a}_i | \Psi \rangle = \alpha_j^* \alpha_i \sqrt{1 - |\alpha_i|^2 - |\beta_i|^2} \sqrt{1 - |\alpha_j|^2 - |\beta_j|^2}, \tag{4.23}$$

which make use of the fact that for the product-state Gutzwiller ansatz wavefunction we employ here, $\langle \hat{a}_i \hat{a}_j \rangle = \langle \hat{a}_i \rangle \langle \hat{a}_j \rangle$.

The phase diagram resulting from Eq. 4.14 is shown in Fig. 4.2. We see a large central $m = 0$ Mott insulating lobe flanked on either side by half-filled checkerboard $m = \pm 1/2$ states of $\hat{t}_+$ bosons and $\hat{t}_-$ bosons for $h > 0$ and $h < 0$ respectively, followed by the fully polarised $m = \pm 1$ Mott insulating states when $|h|$ is sufficiently large. Numerically solving Eq. 4.14 in a finite-sized system using simulated annealing methods reveals the presence of supersolid regions in the vicinity of the tips of the half-filled Mott lobes. The renormalisation group methods we are going to develop in the next sections are not sensitive to this type of ordering and as it may be that the supersolid regions are pre-empted by first-order transitions [212], we do not consider them from here on in.

4.5   Strong Coupling Field Theory

The field theory is constructed in the same way as in the previous chapter, differing only in the microscopics (manifesting as differences in the Green's functions) and that there are two bosonic fields rather than one, though after the Hubbard-Stratonovich transform this difference is largely irrelevant. I shall sketch the derivation of the strong coupling field theory here, highlighting the differences, but for full details on the derivation of the field theory and the following renormalisation group equations, see Chapter 3.

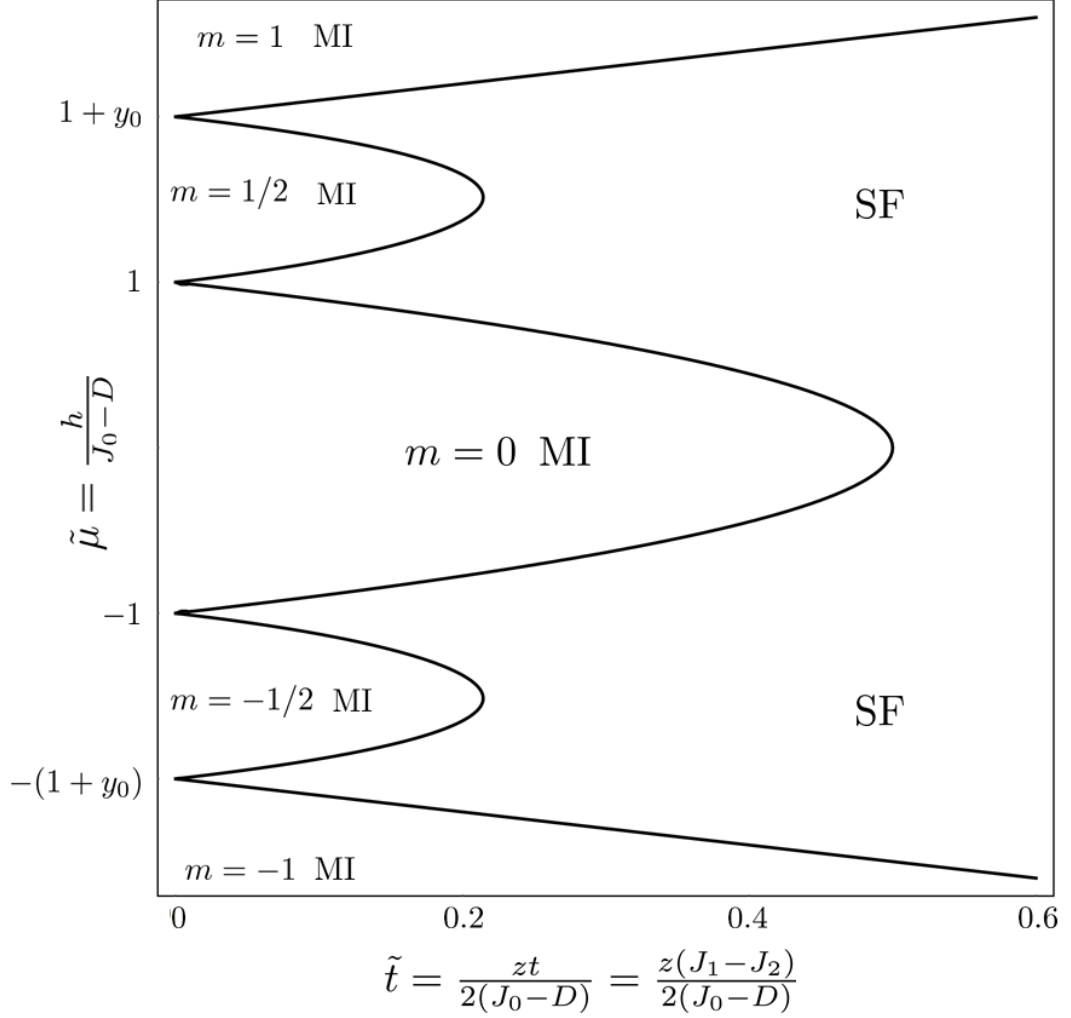We can again write down the action as a path integral over coherent states

Figure 4.2 *The mean-field phase diagram of Eq. 4.1 recast in the language of a hard-core lattice gas of bosons, as in Eq. 4.11, with $y_0 = (J_1 + J_2)/[2(J_0 - D)] \equiv 1$ for the purposes of the figure. There are five Mott insulating (MI) regions corresponding to the bosonic vacuum ($m = 0$), half-filled ($m = \pm 1/2$) and fully polarised ($m = \pm 1$). The superfluid (SF) phase refers to the presence of long-range x-y order, corresponding to canted antiferromagnetism for all $h \neq 0$ and perfect Néel order for $h = 0$.*

in the usual way:

$$S = \int_0^\beta \mathrm{d}\tau \left[\sum_{i,\sigma=\pm} (\phi_{i\sigma}^*(\partial_\tau - \mu_{i\sigma})\phi_{i\sigma} + V \sum_{\langle i,j \rangle} \sum_{\sigma_1,\sigma_2=\pm} \sigma_1\sigma_2 |\phi_{i\sigma_1}|^2 |\phi_{j\sigma_2}|^2 \right.$$
$$\left. + t \sum_{\langle i,j \rangle} \left((\phi_{i-}^* - \phi_{i+})(\phi_{j-} - \phi_{j+}^*) + h.c. \right) \right]. \qquad (4.24)$$

After a Hubbard-Stratonovich transform to decouple the hopping term:

$$S_{\phi\psi} = \int_0^\beta \sum_{ij} \psi_i^* T_{ij}^{-1} \psi_j + \sum_i (\psi_i^*(\phi_{i-}^* - \phi_{i+}) + \psi_i(\phi_{i-} - \phi_{i+}^*)). \qquad (4.25)$$

Because of the structure of the initial action, we only require a single Hubbard-Stratonovich field to perform the decoupling, despite the presence of two bosonic fields in the initial action. This is because we choose the Hubbard-Stratonovich field ψ_i to be conjugate to the linear combination $(\phi_{i-} - \phi_{i+}^*)$. The Hubbard-Stratonovich transform is not unique (see Appendix A), however this form gives a simple physical description of the expected ordering transitions.

We can now trace out the original fields to obtain the partition function in terms of the new fields, denoting this as $\mathcal{Z} = \mathcal{Z}_0 \int \mathcal{D}[\psi^*, \psi] \mathrm{e}^{-(S_\psi^{(0)} + S_\psi')}$ where

$$S_\psi^{(0)} = \int_0^\beta \mathrm{d}\tau \sum_{\langle ij \rangle} \psi_i^* T_{ij}^{-1} \psi_j, \qquad (4.26)$$

$$S_\psi' = -\ln \left\langle T_\tau \exp\left[\int_0^\beta \mathrm{d}\tau \sum_i (\psi_i^*(\tau)(\hat{t}_{i-}(\tau) - \hat{t}_{i+}^\dagger(\tau)) + h.c.) \right] \right\rangle_0. \qquad (4.27)$$

Note that as this problem is strongly interacting, Wick's theorem cannot be employed and the averages must be taken with respect to the Hamiltonian in the zero-hopping limit. At $T = 0$ we know precisely what phases are in the model and with that knowledge we can calculate the averages directly.

Single particle Green's functions

Expanding up to quadratic order:

$$S_\psi^{(2)} = \int_0^\beta \sum_i \left(K_1 \psi_i^* \partial_\tau \psi_i + K_2 |\partial_\tau \psi_i|^2 + G|\psi_i|^2 \right), \qquad (4.28)$$

with

$$G = -\int_{-\beta}^\beta \mathrm{d}\tau \left\langle T_\tau [\hat{t}_{i-}(\tau)\hat{t}_{i-}^\dagger(0) + \hat{t}_{i+}^\dagger(\tau)\hat{t}_{i+}(0)] \right\rangle_0, \qquad (4.29)$$

$$K_1 = \int_{-\beta}^\beta \mathrm{d}\tau\, \tau \left\langle T_\tau [\hat{t}_{i-}(\tau)\hat{t}_{i-}^\dagger(0) + \hat{t}_{i+}^\dagger(\tau)\hat{t}_{i+}(0)] \right\rangle_0, \qquad (4.30)$$

$$K_2 = -\frac{1}{2}\int_{-\beta}^\beta \mathrm{d}\tau\, \tau^2 \left\langle T_\tau [\hat{t}_{i-}(\tau)\hat{t}_{i-}^\dagger(0) + \hat{t}_{i+}^\dagger(\tau)\hat{t}_{i+}(0)] \right\rangle_0, \qquad (4.31)$$

where $G = G_m(\omega = 0)$ is the zero-frequency piece of the single-particle Green's function and the frequency coefficients K_1 and K_2 can be written in terms of derivatives of G just as in Chapter 3.

We may calculate the Green's functions as before, but this time we must distinguish the uniformly-filled states from the checkerboard states and calculate G separately in each region. In calculating the Green's functions for the $m = \pm 1/2$ lobes we must take into account that there are two possible checkerboard configurations (corresponding to each of the two sublattices being either filled or empty) and sum over both accordingly, but aside from this minor extension the calculation proceeds exactly as in Chapter 3. The Green's functions $G_m(\omega)$ in each of the five regions are given by:

$$G_{m=\pm 1}(\omega) = -\frac{1}{\mu_\pm - zV + i\omega},$$
$$G_{m=\pm 1/2}(\omega) = -\frac{1}{\mu_\pm + i\omega} - \frac{1}{-\mu_\pm + zV - i\omega} + \frac{1}{\mu_\mp + zV - i\omega},$$
$$G_{m=0}(\omega) = \sum_{\sigma=\pm} \frac{1}{\mu_\sigma + \sigma i\omega}. \qquad (4.32)$$

Hopping terms

The hopping term $S_\psi^{(0)}$ can be treated using a continuum expansion in the long-wavelength limit.

$$\int_\tau \mathrm{d}\tau \sum_{ij} T_{ij}\phi_i^*\phi_j \approx \int_\omega \int_{\mathbf{k}} t(z - k^2)\phi^*(\mathbf{k},\omega)\phi(\mathbf{k},\omega), \qquad (4.33)$$

where $z = 2d$ is the lattice co-ordination number. The corresponding term in the dual action is:

$$S_\psi^{(0)} \approx \int_k \left(\frac{1}{zt} + \frac{1}{z^2 t} k^2 \right) \psi^*(k) \psi(k), \qquad (4.34)$$

with $k = (\mathbf{k}, \omega)$ for brevity. Combining with the local terms leads to an effective quadratic action of the form:

$$S_\psi^{(2)} = \int_{\mathbf{r},\tau} \left[R|\psi(\mathbf{r},\tau)|^2 + K_1 \psi^*(\mathbf{r},\tau) \partial_\tau \psi(\mathbf{r},\tau) + K_2 |\partial_\tau \psi(\mathbf{r},\tau)|^2 + K_3 |\nabla \psi(\mathbf{r},\tau)|^2 \right], \qquad (4.35)$$

where $K_3 = 1/z^2 t$, $R = R_m = 1/zt + G_m(\omega = 0)$ and the temporal gradient coefficients K_1 and K_2 are defined as before.

Calculating the interaction vertex

Repeating the above procedure to quartic order, we can also calculate the coefficient of the interaction vertex $U = U_m$, which is related to the local two-particle Green's function.

$$S^{(4)} = \frac{1}{8} \left\langle T_\tau \left[\int_0^\beta d\tau \sum_i \left[\psi_i^*(\tau)(\hat{t}_{i-}(\tau) - \hat{t}_{i+}^\dagger(\tau))(\hat{t}_{i-}(\tau) - \hat{t}_{i+}^\dagger(\tau)) \phi_i(\tau) \right] \right]^2 \right\rangle_0^2$$
$$- \frac{1}{4!} \left\langle T_\tau \left[\int_0^\beta d\tau \sum_i \left[\psi_i^*(\tau)(\hat{t}_{i-}(\tau) - \hat{t}_{i+}^\dagger(\tau))(\hat{t}_{i-}(\tau) - \hat{t}_{i+}^\dagger(\tau)) \phi_i(\tau) \right] \right]^4 \right\rangle_0, \qquad (4.36)$$

which may again be written as $U \int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4$ after expanding to quartic order.

Calculation of this coefficient $U = U_m$ is significantly simpler than calculation of the corresponding vertex in Chapter 3 due to the hard-core constraint, though it is still lengthy so I present only the final answer here:

$$U_{\pm 1} = \frac{1}{(\mu_\pm - 2dV)^3} - \frac{1}{2} \frac{1}{(\mu_\pm - 2dV)^2(\mu_\pm - \mu_\mp - 4dV)}, \qquad (4.37)$$

$$U_{\pm 1/2} = \frac{1}{\mu_\pm^3} + \frac{1}{(2dV - \mu_\pm)^3} - \frac{1}{(\mu_\mp + 2dV)^3}$$
$$- \frac{\mu_\pm + \mu_\mp}{(\mu_\pm - 2dV)^2(\mu_\mp + 2dV)^2} - \frac{1}{2} \frac{1}{\mu_\pm^2(\mu_\pm - \mu_\mp)}, \qquad (4.38)$$

$$U_0 = -\frac{1}{\mu_+^3} - \frac{1}{\mu_-^3} - \frac{\mu_+ + \mu_-}{\mu_+^2 \mu_-^2}. \qquad (4.39)$$

Each of these expressions is always positive. There are suggestions that the transition in the vicinity of the tips of the $m = \pm 1/2$ lobes could be first-order

[210] or that there could even be a supersolid phase present [204–211], though other work suggests that the checkerboard $m = \pm 1/2$ phase is thermodynamically unstable [212]. The positivity of U here suggests that either the transition remains second-order, or that effects not captured by our analysis are responsible for turning it first order.

As in Chapter 3, we obtain the final strong coupling field theory by performing a rescaling, employing dimensionless units for imaginary time and length such that $\tau \to V\tau$ and $\mathbf{r} \to \Lambda\mathbf{r}$ where Λ is the momentum cutoff. We then write the other coefficients as $r = R/K_3$, $\gamma_1 = (K_1/K_3)V$, $\gamma_2 = (K_2/K_3)V^2$ and $u = U/K_3$ to obtain the final action.

$$S = \int_{\mathbf{k},\omega} (k^2 - i\gamma_1\omega + \gamma_2\omega^2 + r)\psi^\dagger(\mathbf{k},\omega)\psi(\mathbf{k},\omega) + u\int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4 + ...$$

(4.40)

This expression has precisely the same form as Eq. 3.37 and differs only in the microscopic details. In the context of the renormalisation group calculation to follow, this expression differs from Eq. 3.37 only in the ultraviolet-scale starting values of the flow parameters. The universal behaviour of the phase transitions is the same as that of the conventional Bose-Hubbard model, with the specific differences due to the addition of finite-range interactions being taken into account by the Green's functions.

4.6   Clean Renormalisation Group

Because Eq. 4.40 has the same form as Eq. 3.37 in Chapter 3, the derivation of the renormalisation group equations is the same as for the conventional Bose-Hubbard model and they are again given by:

$$\frac{\mathrm{d}I_0}{\mathrm{d}l} = -2I_0 + 2I_0^2 - 2I_1 I_0^2 u,$$

(4.41)

$$\frac{\mathrm{d}\gamma_1}{\mathrm{d}l} = (2 - d_z)\gamma_1,$$

(4.42)

$$\frac{\mathrm{d}\gamma_2}{\mathrm{d}l} = (2 - 2d_z)\gamma_2,$$

(4.43)

$$\frac{\mathrm{d}u}{\mathrm{d}l} = (4 - d - d_z)u - (I_2 + 4I_3)u^2,$$

(4.44)

where, as before, $I_0 = 1/(1+r)$, d_z is the dynamical critical exponent, $\Omega(\mathbf{k},\omega) = (k^2 - i\gamma_1\omega + \gamma_2\omega^2 + r)$, $I_1 = \int_\omega \Omega(\mathbf{k}=1,\omega) = 1/\sqrt{\gamma_1^2 + 4\gamma_2(1+r)}$, $I_2 = \int_\omega \Omega(\mathbf{k}=1,\omega)^2 = 2\gamma_2 I_1^3$ and $I_3 = \int_\omega |\Omega(\mathbf{k}=1,\omega)|^2 = (1/2)I_0 I_1$.

As previously, we can retain the interaction vertex u and calculate the effect

it has. The clean phase diagram differs only quantitatively from the Gutzwiller mean-field phase diagram in Fig. 4.2 and the interaction vertex leads only to a small quantitative shift of the phase boundaries, as shown in Fig. 4.3. Again, we may safely neglect the interaction vertex from here on in.

4.7 Disordered Renormalisation Group

We can perform the disorder average exactly as in Chapter 3 and obtain a replica field theory of precisely the same form:

$$S = \sum_\alpha \int_{\mathbf{k},\omega} (k^2 - i\overline{\gamma_1}\omega + \overline{\gamma_2}\omega^2 + \overline{r})\psi_\alpha^\dagger(\mathbf{k},\omega)\psi_\alpha(\mathbf{k},\omega) \qquad (4.45)$$

$$- \frac{1}{2}\sum_{\alpha,\beta} g_{\alpha\beta} \int'_{\mathbf{k}_i,\omega,\omega'} \psi_\alpha^\dagger(\mathbf{k}_1,\omega)\psi_\alpha(\mathbf{k}_2,\omega)\psi_\beta^\dagger(\mathbf{k}_3,\omega')\psi_\beta(\mathbf{k}_4,\omega') + \ldots$$

which leads to the same one-step replica symmetry breaking solution as we saw previously. The RG equations, almost unchanged from Chapter 3, are:

$$\frac{\mathrm{d}I_0}{\mathrm{d}\ell} = (3/2\lambda_1 - 2)I_0 + 2I_0^2, \qquad (4.46)$$

$$\frac{\mathrm{d}\lambda_1}{\mathrm{d}\ell} = (4I_0 - d)\lambda_1 + 9\lambda_1^2, \qquad (4.47)$$

$$\frac{\mathrm{d}\overline{\gamma}_1}{\mathrm{d}\ell} = (2 - d_z + \lambda_1)\overline{\gamma}_1, \qquad (4.48)$$

$$\frac{\mathrm{d}\overline{\gamma}_2}{\mathrm{d}\ell} = (2 - 2d_z + \lambda_1)\overline{\gamma}_2 + \lambda_1 I_0 \overline{\gamma}_1^2. \qquad (4.49)$$

Here we retain the frequency terms, as it turns out they are crucial to getting the compressibility correct (see next section). We do, however, make one important change - now that we are retaining the frequency terms it is clear that $d_z = 2$ is not the correct choice of critical exponent in the disordered case. To ensure that γ_1 does not scale under the RG, we require $d_z = 2 + \lambda_1$ where λ_1 is also a running variable under the RG scheme. This does not affect the structure of the phase diagram (neither here nor in Chapter 3) since the phase diagram is determined entirely from the flow of I_0 and λ_1 and the frequency terms play no role in the setting the phase boundaries. The frequency terms do, however, play a significant role in calculating thermodynamic quantities and so it is important that we use the correct d_z taking into account its dependence on disorder.

The RG flow diagram is identical to Fig. 3.10, so I do not reproduce it here. The phase diagram may be determined in the same manner as before, characterising the glassy phases by the divergence of the relative disorder variance λ_1, resulting in boundaries shown in Fig. 4.4. Similarly to the
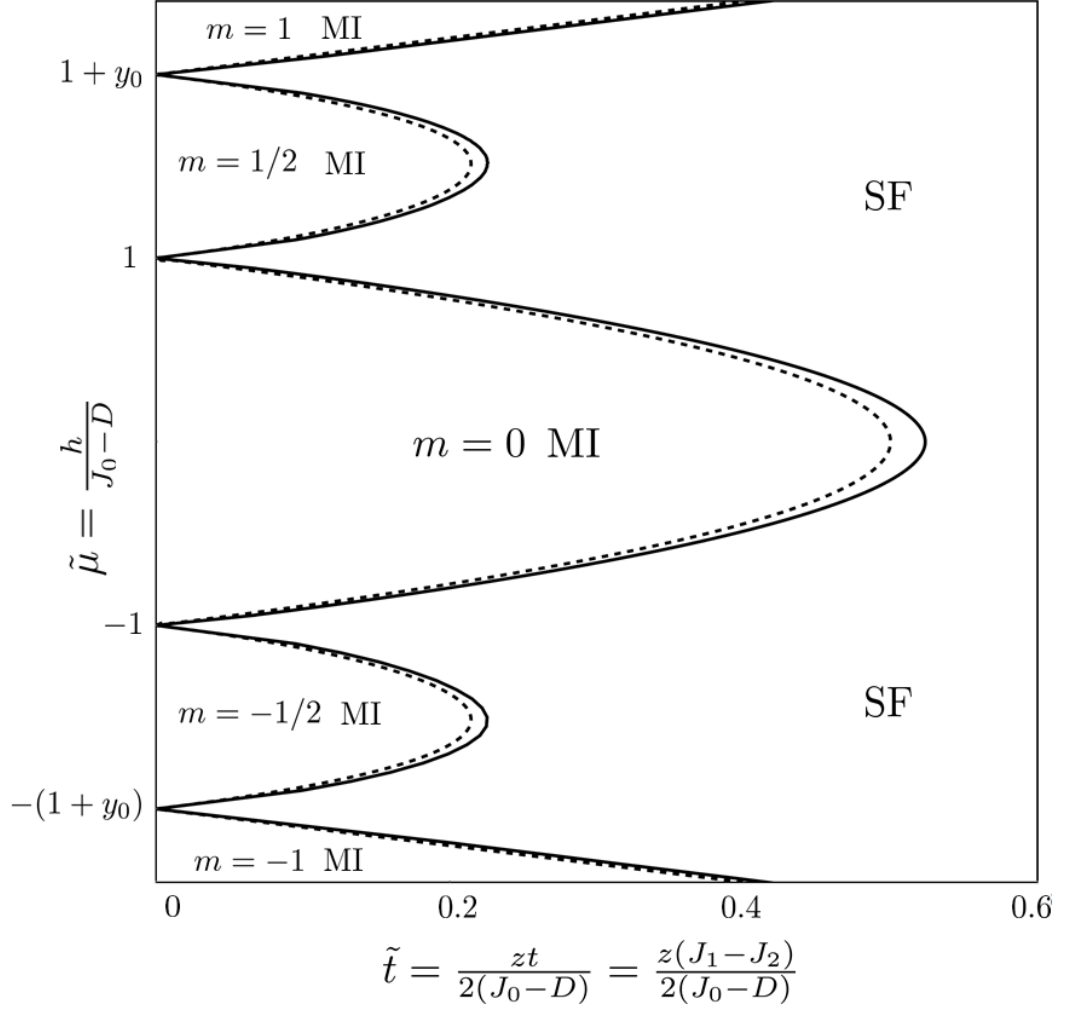
Figure 4.3 *The phase boundaries determined by the renormalisation group calculation (solid lines) plotted against the Gutzwiller mean-field phase boundaries (dashed lines). The mean-field solution overestimates the phase coherence and generates Mott lobes which are slightly too small, though there is no qualitative difference in the topology of the phase diagram.*
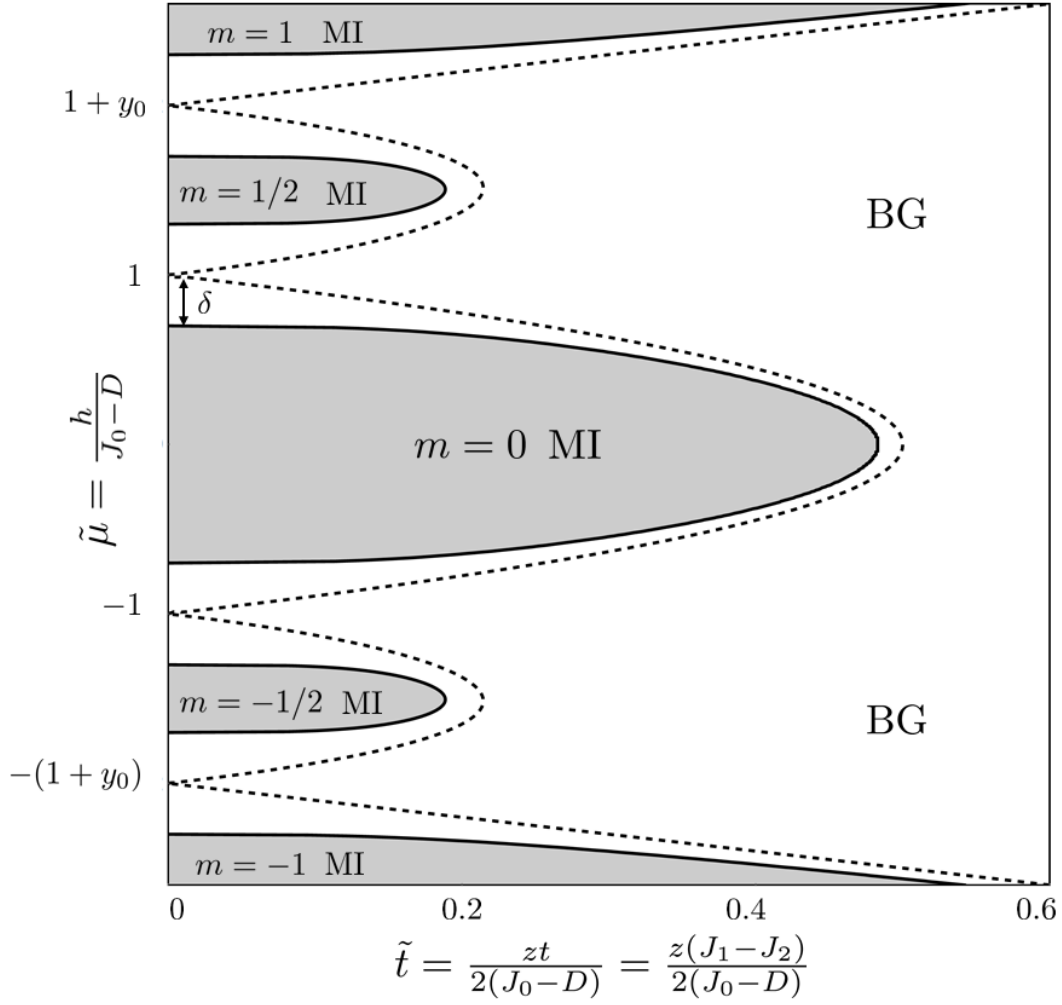
Figure 4.4 *The MI/BG phase boundary determined by the renormalisation group calculation (solid black lines) plotted against the clean MI/SF phase boundary determined by the Gutzwiller mean-field theory (dashed lines). The shaded areas are the Mott insulating regions. As in Chapter 3, we see a shrinkage of the MI regions in the presence of disorder and the emergence of a (magnetic) Bose glass phase between the MI and the SF. As in Chapter 3, the BG-SF transition is not captured by this strong-coupling field theory and is not shown.*

regular Bose-Hubbard model from Chapter 3, we see that the MI regions shrink commensurately with the strength of disorder. Note that in this case, the $m = \pm 1/2$ lobes appear significantly less stable to the effects of disorder than the central $m = 0$ lobe or the fully polarised $m = \pm 1$ regions, though this may be in part a cosmetic feature due to our choice of $y_0 = 1$ used in plotting the figure.

4.8 Compressibility

Motivated by experimental results on the dimerised quantum antiferromagnet DTN [171] and various numerical simulations [150, 152, 153, 201] of Bose-Hubbard models which have reported zero or anomalously low compressibility in the vicinity of the $d_z = 1$ tips of the Mott lobes, we wish to examine in detail these particle-hole symmetric points where the universality class changes. Due to complications arising from the possibility of first-order or supersolid behaviour in the vicinity of the $m = \pm 1/2$ lobes, we restrict our analysis to the central $m = 0$ lobe; however, our conclusions are valid for every Bose-Hubbard-like system with a continuous phase transition at the lobe tips.

The compressibility was first calculated in the context of this RG framework in Ref. [136] for the conventional soft-core Bose-Hubbard model where it was found to take the form:

$$\kappa \approx \frac{1}{\overline{r}^2} \left(1 + \frac{4g}{\overline{r}^2} \right). \tag{4.50}$$

This expression was derived from an asymptotic approximation in which only the $\overline{r}$ and g terms were retained in the action and gives rise to a compressible phase whenever the relative variance $\lambda \approx g/\overline{r}^2$ diverges. Close to the transition, this expression for the compressibility has a scaling form of $\kappa \sim (x - x_c)^{\gamma}$ with $\gamma = 4/D - 1$, where $x - x_c$ is the distance to the transition.

In light of the results suggesting a small or zero compressibility at the tips of the Mott lobes where the linear frequency coefficient γ_1 vanishes, it seemed likely that any approximation which neglected the frequency dependence would not be sensitive to any such change in behaviour at the tip of the lobes. In order to examine the change in behaviour at the tips of the lobes, here we derive the compressibility retaining all frequency terms.

The compressibility is typically written as the derivative of the average occupation with respect to chemical potential, and is related to number fluctuations through the formula: $\kappa_i = \partial \langle \hat{n}_i \rangle / \partial \mu = \beta \sum_j (\langle \hat{n}_i \hat{n}_j \rangle - \langle \hat{n}_i \rangle \langle \hat{n}_j \rangle)$. In the spin (triplon) case, the derivative is instead with respect to the dimensionless magnetic field which plays the role of chemical potential $\tilde{\mu} = h/(J_0 - D)$ and the compressibility may be directly calculated from the replica field theory via the expression

$$\kappa = \sum_i \kappa_i = -\frac{\partial^2 \overline{\mathcal{F}}}{\partial \tilde{\mu}^2} = T \lim_{n \to 0} \frac{1}{n} \left(\frac{\partial^2 \overline{\mathcal{Z}^n}}{\partial \tilde{\mu}^2} \right). \tag{4.51}$$

We can calculate this by expanding out the disorder and interaction vertices to obtain a total compressibility in terms of three contributions: the mean-field κ_0,

the interaction term κ_u and the disorder-induced compressibility κ_g.

$$\kappa \approx T \lim_{n \to 0} \frac{1}{n} \frac{\partial^2}{\partial \tilde{\mu}^2} \left[\int \mathcal{D}[\psi^*, \psi] e^{-S_0} (1 - S_u - S_g) \right] \tag{4.52}$$

$$= \kappa_0 + \kappa_u + \kappa_g, \tag{4.53}$$

where

$$S_0 = \sum_\alpha \int_{\mathbf{k},\omega} (k^2 - i\gamma_1 \omega + \gamma_2 \omega^2 + r)\psi_\alpha^\dagger(\mathbf{k},\omega)\psi_\alpha(\mathbf{k},\omega), \tag{4.54}$$

$$S_u = u \int_{\mathbf{r},\tau} |\psi(\mathbf{r},\tau)|^4, \tag{4.55}$$

$$S_g = -\sum_{\alpha\beta} \frac{g_{\alpha\beta}}{2} \int_{\mathbf{r},\tau,\tau'} |\psi_\alpha(\mathbf{r},\tau)|^2 |\psi_\beta(\mathbf{r},\tau')|^2. \tag{4.56}$$

Explicit expressions for the three contributions are as follows:

$$\kappa_0 = T \lim_{n \to 0} \frac{1}{n} \left[2\overline{\gamma}_2 \sum_\alpha \int_{\mathbf{k},\omega} \langle |\psi_\alpha(\mathbf{k},\omega)|^2 \rangle_0 \right.$$
$$\left. + \sum_{\alpha\beta} \int_{\mathbf{k},\mathbf{k}',\omega,\omega'} (\overline{\gamma}_1 + 2i\omega\overline{\gamma}_2)(\overline{\gamma}_1 + 2i\omega'\overline{\gamma}_2)\langle |\psi_\alpha(\mathbf{k},\omega)|^2 |\psi_\beta(\mathbf{k}',\omega')|^2 \rangle_0 \right], \tag{4.57}$$

$$\kappa_u = -T \lim_{n \to 0} \frac{u}{n} \left[2\overline{\gamma}_2 \sum_{\alpha\beta} \int_{\mathbf{k},\omega} \int_{\mathbf{r},\tau} \langle |\psi_\alpha(\mathbf{k},\omega)|^2 |\psi_\beta(\mathbf{r},\tau)|^4 \rangle_0 \right.$$
$$\left. + \int_{\mathbf{k},\mathbf{k}',\omega,\omega'} \int_{\mathbf{r},\tau'} \sum_{\alpha\beta\gamma} (\overline{\gamma}_1 + 2i\omega\overline{\gamma}_2)(\overline{\gamma}_1 + 2i\omega'\overline{\gamma}_2)\langle |\psi_\alpha(\mathbf{k},\omega)|^2 |\psi_\beta(\mathbf{k}',\omega')|^2 |\psi_\gamma(\mathbf{r},\tau)|^4 \rangle_0 \right], \tag{4.58}$$

$$\kappa_g = \kappa_{g1} + \kappa_{g2}$$
$$= T \lim_{n \to 0} \frac{1}{n} \left[\overline{\gamma}_2 \sum_{\alpha\beta\gamma} \int_{\mathbf{k},\omega} \int_{\mathbf{r},\tau,\tau'} g_{\alpha\beta}\langle |\psi_\gamma(\mathbf{k},\omega)|^2 |\psi_\alpha(\mathbf{r},\tau)|^2 |\psi_\beta(\mathbf{r},\tau')|^2 \rangle_0 \right.$$
$$+ \int_{\mathbf{k},\mathbf{k}',\omega,\omega'} \int_{\mathbf{r},\tau,\tau'} \sum_{\alpha\beta\gamma\delta} \frac{g_{\alpha\beta}}{2}(\overline{\gamma}_1 + 2i\omega\overline{\gamma}_2)(\overline{\gamma}_1 + 2i\omega'\overline{\gamma}_2)$$
$$\left. \times \quad \langle |\psi_\gamma(\mathbf{k},\omega)|^2 |\psi_\delta(\mathbf{k}',\omega')|^2 |\psi_\alpha(\mathbf{r},\tau)|^2 |\psi_\beta(\mathbf{r},\tau')|^2 \rangle_0 \right], \tag{4.59}$$

where we neglect terms higher than quadratic order in frequency[1]. Now that all expectation values are taken with respect to the quadratic action S_0, Wick's theorem may be employed to evaluate the above expressions.

---

[1]The frequency terms may be retained to infinite order and the following results still hold, though this complicates the presentation slightly so I show only up to quadratic order here.

Replica Symmetric Solution

As a first approximation, we can attempt to evaluate the compressibility using the replica symmetric field theory. Based on the results of Chapter 3, it appeared that the replica symmetric theory was qualitatively correct and differed only in small quantitative ways from the full replica symmetry-broken field theory.

First, let us briefly examine κ_0. Earlier we defined the quantity

$$I_1 = \int_{k,\omega} \frac{1}{k^2 + \overline{r} - i\overline{\gamma}_1\omega + \overline{\gamma}_2\omega^2} = \frac{1}{\sqrt{\overline{\gamma}_1^2 + 4(1+\overline{r})\overline{\gamma}_2}},$$

(4.60)

where $r > 0$. Evaluating κ_0 we find that:

$$\kappa_0 = \frac{S_d}{(2\pi)^d}\left[2\overline{\gamma}_2 - \overline{\gamma}_1^2\frac{\partial}{\partial\overline{r}} + 4\overline{\gamma}_1\overline{\gamma}_2\frac{\partial}{\partial\overline{\gamma}_1} + 4\overline{\gamma}_2^2\frac{\partial}{\partial\overline{\gamma}_2}\right]I_1.$$

(4.61)

Taking these derivatives, we find that κ_0 vanishes identically. The same is true for κ_u and all the replica-symmetric contributions to κ_g, since both of these contributions are proportional to two further derivatives of the above expression.

The identical vanishing of the compressibility cannot be a correct result, as we know from many other independent methods that the Bose glass is a compressible phase. In this instance, the replica symmetric field theory leads to a completely incorrect answer. It turns out that by incorporating the replica symmetry breaking solution, however, we can repair this calculation and obtain a physically reasonable solution.

The importance of replica symmetry breaking

Let's take a closer look at the disordered term κ_g allowing for some generic form of replica symmetry breaking using the following definitions:

$$g_{\alpha\alpha} = \tilde{g},$$

(4.62)

$$\sum_{\alpha\neq\beta} g_{\alpha\beta} = \int_0^1 g(u)\mathrm{d}u \equiv \langle g\rangle.$$

(4.63)

In the case of replica symmetry, $\tilde{g} = \langle g\rangle$. In the presence of replica symmetry breaking, these quantities differ. Splitting Eq. 4.59 into two pieces κ_{g1} and κ_{g2}

and evaluating, we find:

$$\kappa_{g1} = \lim_{n \to 0} \frac{1}{n} \overline{\gamma}_2 \sum_{\alpha\beta\gamma} g_{\alpha\beta} \langle |\psi_\gamma(k,\omega)|^2 |\psi_\alpha(\mathbf{r},\tau)|^2 |\psi_\beta(\mathbf{r},\tau')|^2 \rangle_0 \qquad (4.64)$$

$$= 2\tilde{g}\overline{\gamma}_2 \int_{\mathbf{k},\omega} \Omega(\mathbf{k},\omega)^3 + 2(\tilde{g} - \langle g \rangle)\overline{\gamma}_2 \int_{k,q,\omega,\omega'} \Omega(\mathbf{k},\omega)^2 \Omega(\mathbf{q},\omega'), \qquad (4.65)$$

where $\Omega(\mathbf{k},\omega) = (k^2 - i\overline{\gamma}_1\omega + \overline{\gamma}_2\omega^2 + r)$ is as defined earlier in the chapter. The importance of RSB is already visible in Eq. 4.65. The second term in the equation - the RSB term - comes with a different frequency and momentum structure to the replica-symmetric term, meaning it no longer respects the proportionality in Eq. 4.61. Proceeding to evaluate κ_{g1} gives:

$$\kappa_{g1} = 2\overline{\gamma}_2 \frac{S_d}{(2\pi)^d} \left[\frac{6\tilde{g}\overline{\gamma}_2^2}{[\overline{\gamma}_1^2 + 4(1+r)\overline{\gamma}_2]^{5/2}} + \frac{2\overline{\gamma}_2^2(\tilde{g} - \langle g \rangle)}{[\overline{\gamma}_1^2 + 4(1+r)\overline{\gamma}_2]^2} \right], \qquad (4.66)$$

where S_d is again the surface area of a d-dimensional hypersphere. We can similarly evaluate the other contribution to the disorder-induced compressibility, κ_{g2}, to obtain:

$$\kappa_{g2} = -12\overline{\gamma}_2 \frac{S_d}{(2\pi)^d} \frac{\tilde{g}\overline{\gamma}_2^2}{[\overline{\gamma}_1^2 + 4(1+r)\overline{\gamma}_2]^{5/2}}. \qquad (4.67)$$

Here, the RSB terms vanish identically. Putting the pieces together, we find that κ_{g2} exactly cancels the replica symmetric contribution to κ_{g1}, leaving us with the remaining piece that is only non-zero in the presence of replica symmetry breaking:

$$\kappa_g = \frac{S_d}{(2\pi)^d} \frac{4\overline{\gamma}_2^2(\tilde{g} - \langle g \rangle)}{[\overline{\gamma}_1^2 + 4(1+r)\overline{\gamma}_2]^2} \qquad (4.68)$$

$$= \frac{S_d}{(2\pi)^d} \frac{4(\tilde{\lambda} - \langle \lambda \rangle)I_0^2 \overline{\gamma}_2^2}{(\overline{\gamma}_1^2 I_0^2 + \overline{\gamma}_2 I_0)^2}. \qquad (4.69)$$

The replica off-diagonal terms are the only ones which do not cancel due to the RSB diagrams having an additional internal frequency integral and consequently not following the same frequency structure as the other terms and therefore not vanishing due to the proportionality in Eq. 4.61. Without replica symmetry breaking, $\tilde{\lambda} = \langle \lambda \rangle$ and the compressibility is exactly zero everywhere in the disordered system. (In the clean system, $r < 0$ and the proportionality in Eq. 4.61 no longer holds.)

Replica symmetry breaking is therefore crucial to include if we are to obtain a compressible Bose glass state. This result is only possible because the frequency structure of the replica symmetry broken terms is different from that of the replica symmetric ones, and this result will also hold for all Bose-Hubbard-like models which display the same frequency structure.

110

Scaling form of the compressibility

Using the 1-step RSB detailed in Chapter 3 and in Ref. [120] and inserting into Eq. 4.69, we find:

$$\kappa_g = \frac{S_d}{(2\pi)^d} \frac{4u_c\tilde{\lambda}I_0^2\overline{\gamma}_2^2}{(\overline{\gamma}_1^2 I_0^2 + \overline{\gamma}_2 I_0)^2}. \tag{4.70}$$

Substituting for u_c and defining $\tilde{\lambda} \equiv \lambda$ we arrive at the final expression:

$$\kappa_g = \frac{S_d}{(2\pi)^d} \frac{I_0^2\overline{\gamma}_2^2\lambda}{(I_0^2\overline{\gamma}_1^2 + 4I_0\overline{\gamma}_2)^{3/2}}. \tag{4.71}$$

In the vicinity of $P_{MI/BG}$ we can linearise the RG equations and solve two of them analytically, yielding

$$I_0(l) = c_0 e^{-\left(2 - \frac{d}{6}\right)l}, \tag{4.72}$$

$$\lambda(l) = \frac{d}{9} + \delta\lambda e^{d\,l}. \tag{4.73}$$

From our choice of a scale-dependent $z(l)$, we ensure that $\overline{\gamma}_1(l)$ is not renormalised. The RG equation for $\overline{\gamma}_2(l)$ becomes:

$$\overline{\gamma}_2'(l) = -\left(2 + \frac{d}{9}\right)\overline{\gamma}_2 + c_0\frac{d}{9}\overline{\gamma}_1^2 \exp\left[-\left(2 - \frac{d}{6}\right)l\right]. \tag{4.74}$$

Using the ansatz $\overline{\gamma}_2(l) = f(l)\exp[-(2 + d/9)l]$, we can then solve this equation to get a final expression for the scaling of $\overline{\gamma}_2(l)$:

$$\overline{\gamma}_2(l) = \frac{2}{5}c_0\overline{\gamma}_1^2 \exp[-(2 - d/6)l] + c_1 \exp[-(2 + d/9)l]. \tag{4.75}$$

Linearising this result around $P_{MI/BG}$, we find that the compressibility vanishes as

$$\kappa = \kappa(\ell^*) = C(x - x_c)^{\frac{2}{d} - \frac{1}{6}}, \tag{4.76}$$

which gives $\kappa = C(x - x_c)^{\frac{1}{2}}$ near the transition in $d = 3$. The results from numerical integration of the renormalisation group equations are shown in Fig. 4.5 where we find excellent agreement with the scaling exponent of $1/2$ in the vicinity of the phase transition. Close to the tip of the lobe, we see an anomalous suppression of the range of this universal behaviour though it still always holds sufficiently close to the transition. We can go arbitrarily close to the tip and always find an exponent of $1/2$ close to the transition.
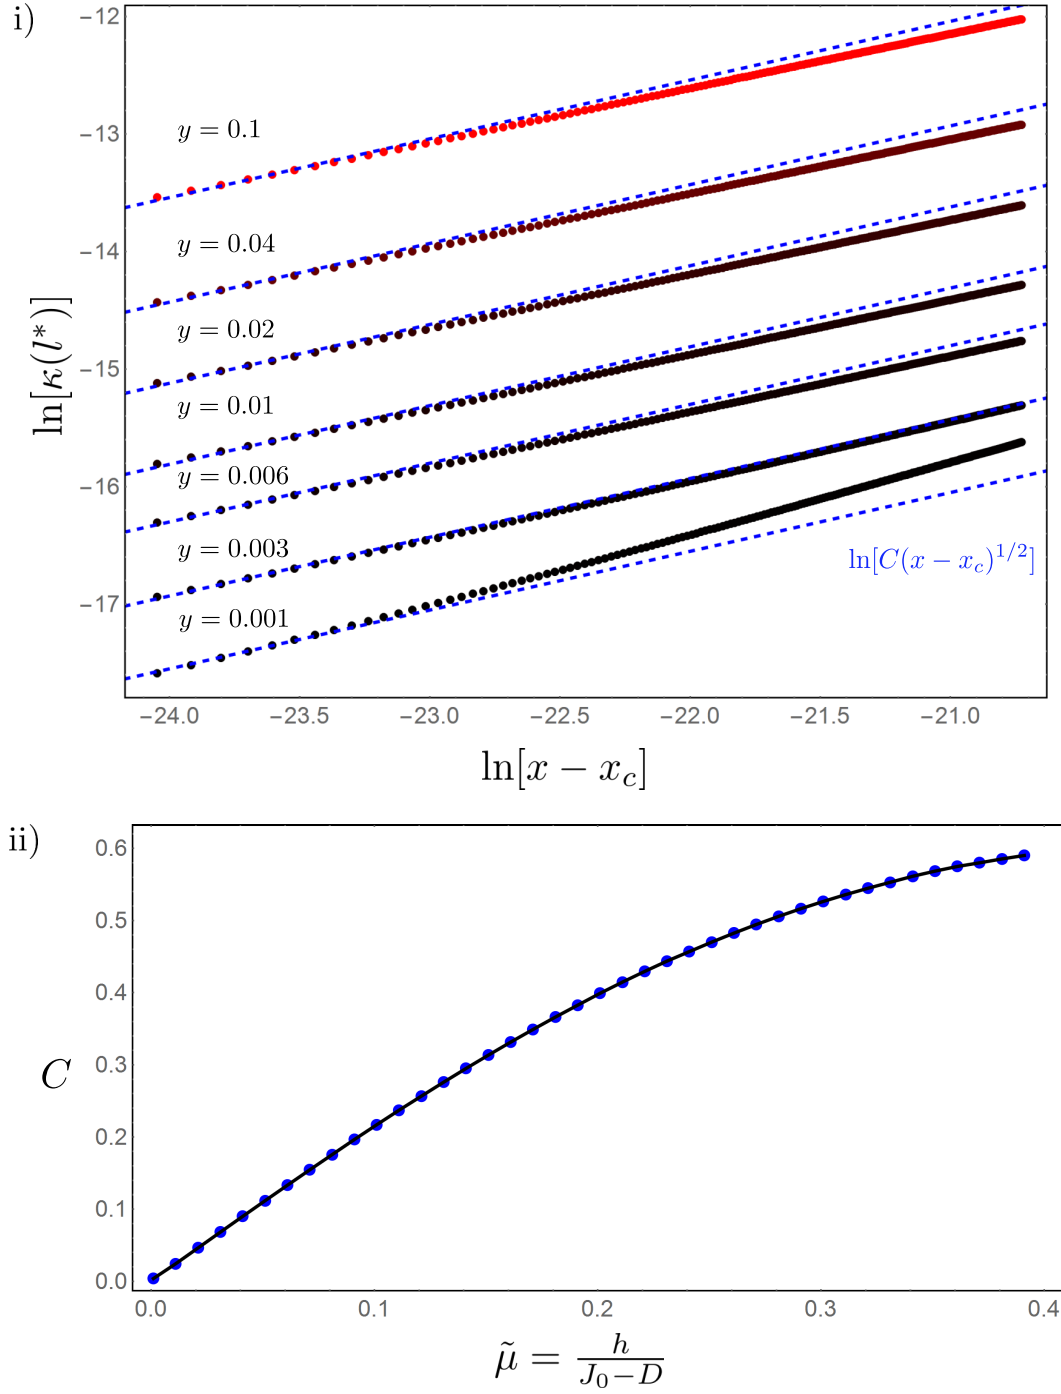
Figure 4.5 *i) Scaling behaviour of κ close to the transition for a variety of different $\tilde{\mu}$ values. The blue lines are guides to the eye with gradient 1/2. Note the anomalous suppression of the range of universal behaviour close to the tip at small $\tilde{\mu}$ values. ii) Vanishing of C on approach to the particle-hole symmetric tip of the Mott lobe.*

Vanishing compressibility

Fitting the numerical integration of the renormalisation group equations with this scaling form allows us to numerically extract the coefficient C close to the transition. On approach to the tip of the Mott lobe, we find that the coefficient C vanishes (Fig. 4.5), signifying the transition into an incompressible state that still nonetheless has a diverging relative variance λ.

This strongly suggests that the transition at the tip is in fact into an incompressible Mott glass rather than a Bose glass. This was also suggested in Ref. [153] where the authors termed the phase an 'anomalous Bose glass' as their numerics could not resolve whether the compressibility was truly zero, only that it was strongly suppressed. Although our RG is not valid strictly at the tip of the Mott lobe, our analysis provides strong evidence that the compressibility does truly vanish at this point.

The suggestion of an incompressible phase in a Bose-Hubbard model with random-mass disorder has been slightly controversial, with some claiming that it must necessarily violate the theorem of inclusions [147, 148] which appears to say that an incompressible phase with compressible inclusions must itself be compressible. This isn't true and the resolution to this objection is found in the symmetries of the multicritical point at the tip of the Mott lobe.

In zero applied field and with a symmetric disorder distribution, there is perfect symmetry between the spin-up and spin-down triplets. In the thermodynamic limit there will be an exactly equal number of each in the system. Any infinitesimal change in applied field will increase the number of one species of triplet, but due to the symmetry at this point there will be a precisely equal decrease in the occupation number of the other series of triplet. In other words, the compressibility $\partial \langle \hat{n} \rangle / \partial \tilde{\mu} = 0$ because $\langle \hat{n} \rangle$ does not change on average across the entire system. Even though the system still contains locally compressible regions, the bulk behaviour of the system is incompressible due to this perfect symmetry. This result does not violate the theorem of inclusions.

In any finite size system, this symmetry will no longer be exact and the compressibility will not vanish identically, which is why the Monte Carlo simulations in Ref. [153] do not find a compressibility of exactly zero.

This result has implications beyond the zero-field limit of the spin model, however. As the field theory used in Chapter 3 is precisely the same as the one used here, barring changes in the ultraviolet-scale starting values of the flow parameters, we should expect to see the same behaviour at the tips of the Mott lobes in the regular Bose-Hubbard model and by extension all Bose-Hubbard-like models which obey this form of field theory. How can we reconcile this with the theorem of inclusions?

The answer is that this form of field theory predicts an emergent particle-hole symmetric point at the tip of every Mott lobe, even in the case of chemical

potential disorder. At these points, there is always a mapping one can do such that the integer filling of the Mott lobe is the 'vacuum' with excitations above this vacuum behaving like particles and excitations below it behaving like holes. Precisely at the tip in the disordered case, there is always an equal number of particles and holes and there is always a perfect cancellation between the change in number of particles and change in number of holes as the chemical potential is varied. The Mott glass is therefore generic to all Bose-Hubbard-like models which display particle-hole symmetric points, emergent or otherwise.

4.9 Discussion

There are two key results from this section, both of which build upon the work in Chapter 3.

- **The importance of replica symmetry breaking** - Replica symmetry breaking, contrary to even our own intuition (see Ref. [120] where we stated there to be no relation between compressibility and RSB) turns out to be crucial for the thermodynamics of the system.

- **The Mott glass** - An incompressible Mott glass phase was found to exist in this model, and by extension must also be present in the regular Bose-Hubbard model as the field theories are identical. This phase has been suggested to exist in two and three dimensions by numerics and in three-dimensional antiferromagnets by a single experiment [171], but this is the first analytic prediction of a Mott glass in anything other than a $d = 1$ fermionic system or rotor model [149, 162, 198–200].

One key question remains: is replica symmetry breaking a real, physical effect? RSB was first invoked in the renormalisation of disordered systems [41, 163, 164] to take into account the possibility of there being a large number of metastable states, but also because the entire procedure underlying renormalisation begins to break down as soon as the possibility of a 'glassy' or non-ergodic free energy landscape is allowed. Here, however, we required replica symmetry breaking to obtain a physical result before we even attempted the renormalisation group procedure. In this calculation, it was required in order to correctly compute the thermodynamic averages involved in calculating the compressibility.

Replica symmetry breaking arises here because blindly averaging across all minima in the free energy landscape is akin to averaging the 'up' and 'down' ground states of a homogeneous ferromagnet. Without allowing for a symmetry breaking field to restrict the averaging procedure to particular regions of the free energy landscape, the averaging procedure washes out all of the physics we are interested in. In a glassy system where the ultrametricity is not caused by

114

any symmetry of the Hamiltonian, we do not know the form of the symmetry-breaking field we need to apply to correctly restrict our averages. By allowing the system the freedom to spontaneously break replica symmetry, we allow it to self-consistently restrict the thermodynamic averages to the correct regions of the free energy landscape. RSB plays the role of the external field in the ferromagnetic example.

This work shows that incorporating RSB into the field theory is a requirement in order to obtain physically meaningful predictions and is a good indicator that RSB is a real physical effect. It doesn't, however, definitively prove that RSB is not just a trick required to make a potentially incorrect model work. Experimental measurements of the Edwards-Anderson order parameters and/or independent confirmation from another theoretical method will be required in order to know for sure. Quantum Monte Carlo, for example, would be an ideal tool to numerically test for replica symmetry breaking. In light of previous works, however, and the resemblance of our results to the experimental findings of Ref. [171], it seems very likely that RSB is real and that the nature of the Bose glass is somewhat more complex than was previously thought.

Direct measurement of the Edwards-Anderson order parameters is still complicated in a spin system, however, due to the lack of single-site addressability in the relevant experimental materials. While the Edwards-Anderson order parameters are measurable in certain bulk measurements [126], and other glassy effects are certainly measurable, a method which allows us to directly probe the boson number occupancy on a single-site level could be enormously illuminating and allow the Edwards-Anderson order parameters to be directly measured. Recent developments in quantum gas microscopes suggest that this is now possible, and this will be the subject of Chapter 5.

Chapter 5

Imaging Glassy Phases Using Quantum Gas Microscopes

In Chapters 3 and 4 I showed how density fluctuations are linked to the compressibility of disordered lattice gases of bosons, giving rise to replica symmetry breaking and non-zero Edwards-Anderson order parameters. All currently available experimental measurements of properties of the Bose glass, both in ultracold atomic gases and related magnetic systems, are based on inferences drawn from measurements of bulk properties, but it is the local effects of disorder that are important for the glassiness in this model [119]. The development of optical lattice technology [213, 214] and the advent of single-atom resolved fluorescence imaging [215–217] mean that it should be entirely possible to directly image the number fluctuations in the glassy phases and extract the Edwards-Anderson order parameters.

Here, I outline a proposal to directly image the Bose glass, show that this proposal is within reach of current experimental technology and present numerical calculations that serve as predictions for the experiments. I show that at mean-field level, the Edwards-Anderson order parameter is already a good indicator of the Bose glass and that direct measurement of this parameter is well within the capability of current-generation quantum gas microscopes.

This work was conducted collaboratively with Dr Graham Bruce and much of the numerical data presented here was gathered by Liam Walker as part of his final-year MPhys project, with input from project co-ordinator Dr Jonathan Keeling and programming assistance from Tiffany Harte of the University of Oxford.

The results of this work may be found in *"Measuring the Edwards-Anderson order parameter of the Bose glass: a quantum gas microscope approach"*, S. J. Thomson, L. S. Walker, T. L. Harte and G. D. Bruce, arXiv:1607:05254 [218].

5.1 Why quantum gas microscopes?

In the previous chapter we looked at the effects of disorder on the bulk properties of bosonic systems, however we are no closer to being able to measure the Edwards-Anderson order parameters introduced in Chapter 3. Although measuring two-time correlation functions is unlikely to be possible in ultracold atomic gas measurements (as the measurements are typically destructive), the thermally-averaged order parameter q only requires knowledge of the local properties of a system at any given moment in time:

$$q = \overline{\langle \hat{n}_i \rangle \langle \hat{n}_i \rangle} - \overline{\langle \hat{n}_i \rangle} \, \overline{\langle \hat{n}_i \rangle}. \tag{5.1}$$

This complete knowledge of the exact quantum state of each individual lattice site is a tall order and all but impossible in condensed matter systems, however quantum gas microscopes are ideally suited for just this sort of measurement. Single atom resolution fluorescence imaging of ultracold gases of bosons was achieved in 2009 [215–217] and of fermions in 2015 [219–223]. Although disorder has been investigated in numerous experiments on ultracold atomic gases [156, 157, 224–226], these measurements have all inferred the presence of a Bose glass from bulk measurements. Never before has anyone used a single-atom resolution quantum gas microscope to directly image disordered phases and from there extract information about the density fluctuations which characterise them.

The basic idea behind the quantum gas microscope is that ultracold atoms are trapped in an optical lattice generated by counter-propagating laser beams. During the imaging procedure the lattice depth is suddenly ramped up, pinning the atoms onto their respective lattice sites and allowing their density to be imaged using fluorescence imaging techniques. This is most commonly done in two spatial dimensions, and this is the case we will consider here, however one- and three-dimensional lattices are also possible [227].

Disorder can be introduced either by superimposing a speckle potential on top of the lattice [155–158], or by using a spatial light modulator (SLM) to modulate the lattice depth from site to site creating an effective chemical potential disorder [159–161]. The advantage of the SLM is that it allows the generation of arbitrary disorder, e.g. bimodal distributions or spatially correlated disorder, however the imperfect device response of the SLM can introduce an additional unintended disorder on top of the intended pattern [228].

In a real experimental system, adding any form of disorder will simultaneously modify not only the on-site chemical potential, but also the hopping amplitude and the strength of the on-site repulsion [156]. In the following it will be assumed that the dominant effects can be modelled by looking at local chemical potential disorder alone and that any additional effects will lead to small quantitative changes but no new physics.

All of the work presented in this chapter is mean-field in nature and aimed

at simulating the experimental quantum gas microscope setup and showing that current generation experiments can resolve q. The results in this chapter will be presented in two parts. First, we shall look at characterising the various phases in the model in small lattices at a uniform chemical potential and we will see that the Edwards-Anderson order parameter q exists as a mean-field order parameter which is capable of distinguishing the Bose glass from the Mott insulator and superfluid phases. Then, we shall take into account the harmonic trapping potential present in experiments and show that current generation quantum gas microscopes have sufficient resolution to measure q and unambiguously identify the Bose glass.

5.2   Experimental setup

Here I will briefly summarise the experimental setup from the point of view of the minimum detail required in order to theoretically model the system. For a more thorough review of the details of the experimental apparatus, see Refs. [229–231].

In an optical dipole trap, the oscillating electric field of the laser light induces an oscillating dipole moment in the atom. This oscillating dipole moment then interacts with the laser light to generate a trapping potential proportional to the intensity of the applied laser light, $V \propto |E|^2/(\omega - \omega_0)$ where ω is the frequency of the laser and ω_0 is the resonance frequency of the atom. In a real atom, there will be many resonance frequencies and the trapping potential generated will depend on the specific state of the atom. The precise details of optical dipole trapping may be found in Ref. [229], but for our purposes we will assume the atoms to be two-level systems with only one transition of frequency ω_0.

The laser light is typically far from resonance with the atomic transition, meaning that there are no couplings between the field and the atom other than the dipole moment and the potential can be treated as a regular conservative potential. For laser light with a frequency greater than the resonance frequency of the atoms, the resulting potential will be repulsive. Conversely for light with a frequency less than the resonance frequency of the atoms, the resulting potential will be attractive. The advantage of using a purely optical trapping mechanism is that it is completely controllable and the trapping geometry can be changed by changing the intensity profile of the laser beams.

In quantum gas microscopes, the optical confinement takes the form of an optical lattice generated by the interference of counter-propagating laser beams, with the resulting potential shown in Fig. 5.1. Each local minimum acts as a lattice site which the atoms are tightly bound to. Optical lattices can be generated in any dimensionality by setting up counter-propagating laser beams in each dimension. The resulting lattice potential takes the form of a sine wave with a lattice spacing given by half the wavelength of the laser beams used to
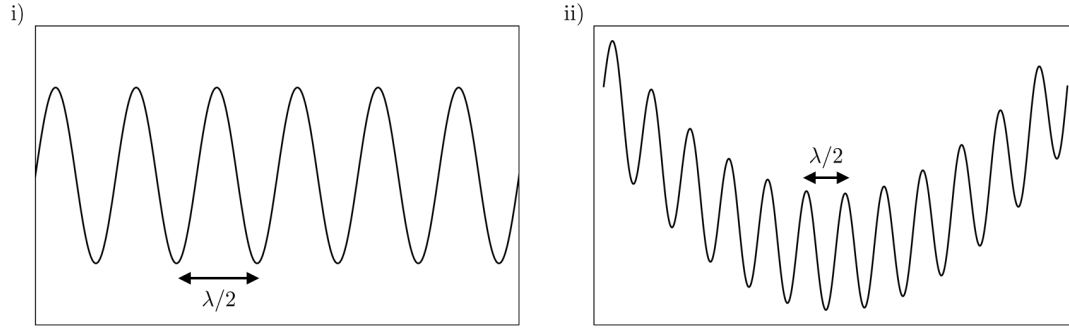
Figure 5.1 *i) The lattice pattern generated by interfering two counter-propagating beams of uniform intensity and wavelength λ. The resulting lattice is homogeneous, with a lattice spacing of λ/2. ii) An exaggerated cartoon of the shape of a real optical lattice made from two orthogonal pairs of counter-propagating beams - a corrugated lattice with spacing λ/2 with a slowly varying background harmonic potential on top resulting in an inhomogeneous lattice. In reality the harmonic contribution is much more slowly varying than shown in this image, but still significant enough to cause measurable effects.*
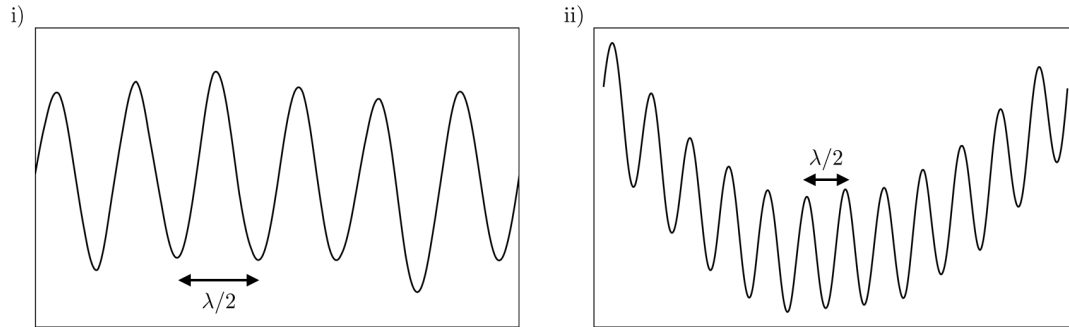

Figure 5.2 *The same two patterns as in Fig. 5.1 except now with added on-site disorder. This disorder will modify the local chemical potential as well as the tunnelling (hopping) amplitude, however in the following we will assume that the main effects can be modelled as chemical potential disorder only.*

generate it. Due to the Gaussian profile of the laser beams (i.e. they are more intense at the centre and less so at the edges), the optical lattice is not entirely uniform and will have a weakly varying harmonic potential in addition to the corrugated standing wave pattern.

This inhomogeneity across the lattice is visible in experiments and results in the lattice 'filling up' with atoms from the centre where the local chemical potential is largest. Our model must therefore take this into account if we are to properly compare with experimental images. Provided the potential varies slowly enough, we can use the local density approximation and assume that the atoms at any given local chemical potential act as though they were part of a bulk system uniformly at that chemical potential. In our model, we add disorder by superimposing another randomly varying pattern over the top of the optical lattice beams (see Fig. 5.2), resulting in small local deformations of the optical lattice which we can identify with on-site chemical potential disorder.

Upon exposure to near-resonant laser light, the atoms will fluoresce and emit light. Detection of this light allows the position of each individual atom to be located. This requires extremely high resolution optics with a diffraction limit comparable to the lattice spacing. Though it is possible to image three dimensional optical lattice systems, most experiments are performed in two dimensions. Depending on the experimental setup, the atoms can either be loaded into a two dimensional optical lattice from the beginning or loaded into a three-dimensional lattice, and then all the atoms not in the imaging plane can be removed. This is typically done by the application of a magnetic field gradient and a microwave frequency sweep which simultaneously flips the spins of all of the atoms except the plane of interest, a technique known as spatially selective magnetic resonance. The remaining atoms can then be removed by a resonant laser pulse.

Throughout the imaging procedure, in order to keep the atoms confined the trap depth must first be ramped up by a factor of several hundred, freezing the atoms in place on their respective lattice sites and prevented any further hopping or movement while the image is taken. Essentially, the atoms are probabilistically projected onto an independent Fock state on each optical lattice site and this forms the basis of the method which we use to model the imaging process. The imaging light is slightly detuned from resonance to reduce any additional heating from the imaging process, preventing the atoms from gaining enough energy to move from their respective lattice sites.

The most important feature of the imaging procedure which we must capture in any attempt at modelling the experimental setup is parity projection. Should there be more than a single atom per optical lattice site, pairs of atoms can undergo a light-assisted collision during imaging and will be removed from the trap. In practice, this means that any lattice site occupied by an even number of particles will appear empty under fluorescence imaging, and any lattice site occupied by an odd number of particles will appear as though occupied by a

single particle. The imaging technique can only distinguish between even and odd occupancies, in other words, and though there are some recent techniques which can go beyond this parity constraint [232, 233], in all of the following we assume it to be an important experimental constraint and model it accordingly.

One key difference between most theoretical methods and the experimental setup is the way the particle number enters. In most theories we work in the grand canonical ensemble where the particle number is not fixed and we can change it by varying the chemical potential μ. In experiments, the number of particles in the trap varies slightly from run to run but throughout the course of an individual experiment it is fixed, subject to minor fluctuations induced by three body losses and other forms of collisions in the trap. Essentially, the way around this is to treat the central region of the harmonic trap (where the potential varies mostly slowly) as the 'core' region and treat the region around the edges of the harmonic trap as a particle reservoir similarly to Ref. [234], allowing us to use calculations in the grand canonical ensemble to model the region of interest at the centre of the trap.

With this brief survey of relevant experimental details that must be taken into account in our analysis, now we can move on to look at the theoretical analysis itself.

5.3   Theoretical Framework

Again, we start from the typical Bose-Hubbard Hamiltonian, which is a good description of spinless bosons hopping on a hypercubic lattice:

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle} (\hat{b}_i^\dagger \hat{b}_j + \hat{b}_j^\dagger \hat{b}_i) - \sum_i \mu_i \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i(\hat{n}_i - 1), \qquad (5.2)$$

where all the symbols have the usual meanings as defined elsewhere in this thesis and $\mu_i = \mu + \varepsilon_i$ represents the chemical potential disorder, where ε_i is in the following drawn from a box distribution of width 2Δ. In all of the following we use $\Delta = 0.3U$.

The conventional mean-field theory we looked at in Chapter 3 works well for calculating phase boundaries and bulk properties of the Mott insulator and superfluid, but it does not capture the properties of the Bose glass. In calculating the phase diagram using this method, we eliminated the possibility of local variations in chemical potential, and the mean-field decoupling used no longer makes a great deal of sense in the presence of disorder. Instead, we need to use an alternative form of mean-field theory which preserves information about the local properties of the system and will ultimately allow us to extract those for comparison with experiment.

Analogously to Refs. [145, 146, 235] we use a local mean-field theory based

122

around the Gutzwiller ansatz [236] wavefunction:

$$|\Psi_G\rangle = \prod_i |\psi_i\rangle = \prod_i \sum_{n_i} \left(f_{0,i} + f_{1,i}\hat{b}_i^\dagger + ... f_{n_i,i}\frac{(\hat{b}^\dagger)_i^n}{\sqrt{n_i!}} \right) |0\rangle \qquad (5.3)$$

$$= \prod_i \sum_{n_i} f_{n_i,i} |n\rangle, \qquad (5.4)$$

from which we can calculate a mean-field energy and minimise it with respect to the variational parameters $f_{n,i}$. The wavefunction is normalised such that $\sum_{n_i} |f_{n_i,i}|^2 = 1 \,\forall i$. The Gutzwiller ansatz is a mean-field wavefunction which becomes exact in the limits of infinite dimensionality, zero hopping or infinite-range hopping (equivalent to infinite dimensionality). Here we use it to simulate experiments in two spatial dimensions where it is known to be a reliable and accurate method other than at the multicritical points at the tips of the Mott lobes where mean-field theory breaks down.

The usual criticism of the Gutzwiller ansatz is that while it is a very capable method of modelling local properties, it often gets bulk properties wrong in the thermodynamic limit. Here, we will make this work for us, simulating small experimentally feasible system sizes and deliberately only looking at the same local properties accessible to the experiments, thus keeping the Gutzwiller ansatz well within its regime of validity. After obtaining the values of the variational parameters $f_{n,i}$, we can then probabilistically calculate the occupancy of each lattice site and generate simulated 'snapshot' images that mimic those produced by quantum gas microscopes, using $\langle \hat{n}_i \rangle = \sum_{n_i} n_i |f_{n_i,i}|^2$.

From the typical theoretical perspective, this manoeuvre is unusual. In probabilistically generating snapshots, we are deliberately throwing away information about the system contained in the Gutzwiller coefficients. There is method in this madness: rather than using these exact values, we instead wish to generate snapshots equivalent to what the experiment would see in order to determine whether the measurements we propose are experimentally feasible. Where possible, we follow closely the experimental setup of Ref. [217] to ensure our mean-field theory takes into account as realistic a description of current experimental capabilities and limitations as possible.

The mean-field energy expression is given by taking the expectation value of the Hamiltonian with respect to the Gutzwiller ansatz wavefunction. The on-site terms give the contribution $E^{(0)}$ as follows:

$$E^{(0)} = \langle \Psi_G | \left(-\sum_i \mu_i \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i(\hat{n}_i - 1) \right) |\Psi_G\rangle \qquad (5.5)$$

$$= \sum_i \sum_{n_i} \left(\frac{U}{2} n_i(n_i - 1) - \mu_i n_i \right) |f_{n_i,i}|^2. \qquad (5.6)$$

The hopping term is slightly more complicated, involving neighbouring sites, and

gives:

$$E^{(t)} = -t \sum_{\langle i,j \rangle} \langle \Psi_G | \, \hat{b}_i^\dagger \hat{b}_j + \hat{b}_j^\dagger \hat{b}_i \, | \Psi_G \rangle \,, \tag{5.7}$$

where the expectation values can be computed like so:

$$\langle \Psi_G | \, \hat{b}_i^\dagger \hat{b}_j \, | \Psi_G \rangle = \prod_{a,b} \sum_{m,n} f_{m_a,a}^* f_{n_b,b} \, \langle m_a | \, \hat{b}_i^\dagger \hat{b}_j \, | n_b \rangle \tag{5.8}$$

$$= \left(\sum_{ni,m_i} \langle m_i | \, \hat{b}_i^\dagger \, | n_i \rangle \right) \left(\sum_{nj,m_j} \langle m_j | \, \hat{b}_j^\dagger \, | n_j \rangle \right) f_{m_i,i}^* f_{n_i,i} f_{m_j,j}^* f_{n_j,j} \tag{5.9}$$

$$= \left(\sum_{n_i=0}^{\infty} \sqrt{n_i+1} f_{n_i+1,i}^* f_{n,i} \right) \left(\sum_{n_j=0}^{\infty} \sqrt{n_j+1} f_{n_j,i}^* f_{n+1,j} \right) \,, \tag{5.10}$$

since any terms containing products of operators on different sites always factorise and can be dealt with separately. In the clean case where all sites are identical, the energy expression becomes site-diagonal and reduces to:

$$E = \langle \Psi_G | \hat{\mathcal{H}} | \Psi_G \rangle = \sum_i E_i, \tag{5.11}$$

$$E_i = \sum_n \left(\frac{U}{2} n(n-1) - \mu n \right) |f_{n,i}|^2 - zt \left| \sum_n \sqrt{n+1} f_{n,i}^* f_{n+1,i} \right|^2 \,, \tag{5.12}$$

subject to the normalisation constraint $\sum_n |f_{n,i}|^2 = 1 \, \forall i$.

In practice, it is numerically more efficient to write the wavefunction $|\Psi_g\rangle$ to explicitly include normalisation which allows us to minimise the resulting energy expression using unconstrained minimisation algorithms, as opposed to implementing a computationally more costly constrained minimisation algorithm which attempts to enforce normalisation upon the Gutzwiller coefficients.

We used a conjugate gradient algorithm to minimise the variational wavefunction [237]. The wavefunction itself was truncated at 6 particles per site, at which point the coefficient of the sixth Fock state $f_{6,i}$ was effectively zero for every site, giving us confidence that this truncation was a sensible one. As a consistency check, the clean case simulations were performed using the disorder code with disorder strength set to zero and were found to identically match the results from using the site-diagonal clean-case expression and exactly reproduce the usual mean-field phase boundary previously described in Chapter 3.

Unless otherwise stated, data presented in all of the following was generated from Python code written by Liam Walker as part of his MPhys project [238]. The majority of the data presented here was taken on the CM-CDT EPSRC Computer Cluster.

5.4   Snapshot Generation

After minimising the mean-field energy expression, occupancy of lattice sites can be calculated probabilistically. Taking into account the parity of the imaging procedure, the probability of imaging a given site i as empty is $\sum_n^{even} |f_{n,i}|^2$ and the probability of imaging the site as occupied is simply $\sum_n^{odd} |f_{n,i}|^2$. In all of the following we simulate two-dimensional lattices for comparison with the experiments.

A simulated snapshot of a variety of 25×25 lattices at different points in the phase diagram is shown in Fig. 5.3. Sites occupied by an odd number of particles are shown in orange and sites with an even occupancy are shown in white. This image differs from the usual 'snapshot' images of Bose glass phases generated by computational methods (see, e.g. Ref. [146]) in that a probabilistic density is imaged, rather than an exact numerical mean occupancy or superfluid order parameter. This corresponds to what one would see in a single experimental image. Larger snapshots can also be used - up to lattice sizes of 100×100 we find quantitative changes on the order of a few percent at the expense of a large increase in computation time.

All of the following analysis is performed upon these snapshot images, with the intention that we do not perform any analysis using information not accessible to the measurements themselves (i.e. exact Gutzwiller coefficients of the Fock states).
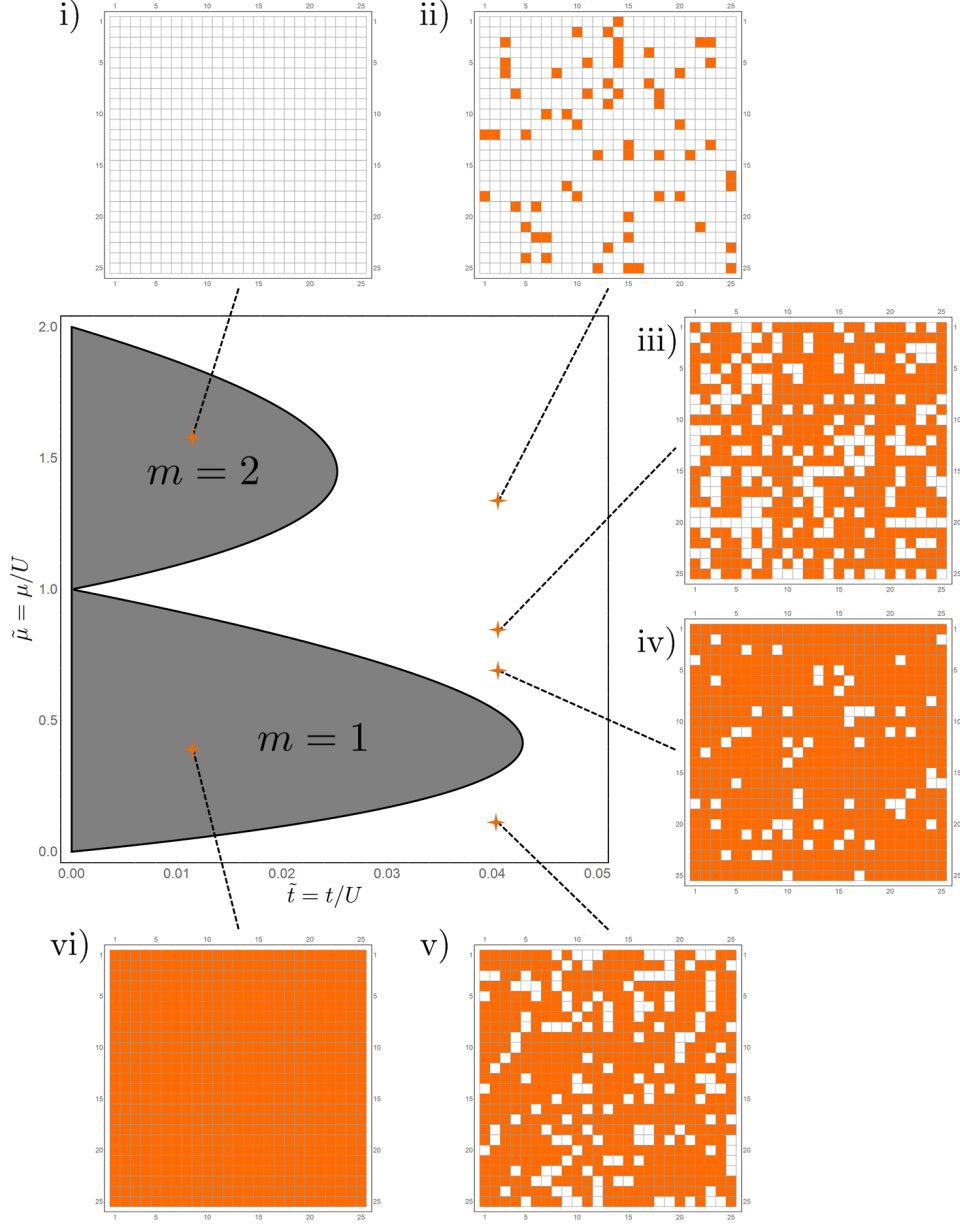
Figure 5.3 *Simulated snapshots in a 25×25 flat lattice at various points in the phase diagram, produced in Wolfram Mathematica using Eq.5.12. Orange indicates a site with a single particle, while white indicates either one or two particles. The precise values of dimensionless hopping and chemical potential at which the snapshots were taken are: i) $\tilde{t} = 0.01, \tilde{\mu} = 1.6$, ii) $\tilde{t} = 0.04, \tilde{\mu} = 1.25$, iii) $\tilde{t} = 0.04, \tilde{\mu} = 0.9$, iv) $\tilde{t} = 0.04, \tilde{\mu} = 0.7$, v) $\tilde{t} = 0.04, \tilde{\mu} = 0.1$, vi) $t = 0.01, \tilde{\mu} = 0.4$. Snapshots i) and vi) are taken in the Mott insulator phases where the filling is uniform. Note that panel i) appears blank due to parity - actually every site is occupied by exactly two particles. The other four snapshots were taken in the superfluid phase and show a non-integer average number of particles per site.*

5.5 Edwards-Anderson as a mean-field order parameter

To make contact with the experimental results, we will ultimately need to run simulations that take into account the presence of the background harmonic trapping potential responsible for the 'wedding cake' structure of the experimental fluorescence images. First, however, we shall show that the Edwards-Anderson order parameter already exists at mean-field level and is capable of distinguishing the Bose glass from the Mott insulator and superfluid. In all of the following, we use 25×25 lattices.

Compressibility

One of the first quantities measured in quantum gas microscopes was the compressibility [217], used to distinguish the incompressible Mott insulator from the compressible superfluid. In the zero-temperature Bose-Hubbard model, we may define a dimensionless local compressibility:

$$
\begin{aligned}
\kappa &= \frac{1}{\beta} \frac{\partial \langle \hat{n}_i \rangle}{\partial \mu} = \frac{1}{\beta} \frac{\partial}{\partial \mu} \left[\frac{\sum_{n_i} n_i \exp[-\beta E_{n_i}]}{\mathcal{Z}} \right] \\
&= \sum_{ij} \left(\langle \hat{n}_i \hat{n}_j \rangle - \langle \hat{n}_i \rangle \langle \hat{n}_j \rangle \right),
\end{aligned} \tag{5.13}
$$

which is equivalent to the variance of our snapshot images. Due to the parity limitation inherent in most quantum gas microscopes, the variance will saturate at a maximum of $\kappa_{max} = 0.25$ (i.e. the maximum variance that any binary distribution can have). In a superfluid with a Poisson distribution (i.e. deep in the superfluid phase), the measured compressibility will saturate at this maximum value, as shown in the experimental results of Ref. [217].

By averaging the variance over an appropriate number of snapshots (typically around ten) we can construct a phase diagram in the $(t/U, \mu/U)$ plane. The results in both the clean and disordered cases are shown in Fig. 5.4. In the clean case, the boundary of zero compressibility matches up exactly with the conventional mean-field phase boundary, as expected, confirming that even a 25×25 lattice is large enough to give reliable results. The phase diagram can be constructed from larger lattices with a commensurate increase in computation time, however there is only a negligible difference in simulated quantities.

In the disordered case we see that the Mott lobes shrink, as expected, and a highly compressible phase encroaches upon what was previously a Mott insulator. The MI-BG phase boundary as determined from the replica symmetric renormalisation group developed in Chapter 3 is shown overlaid[1] - there is a minor

---

[1]The replica symmetric RG is used here as the interaction vertex W is relevant in two dimensions. The one-step RSB result was derived in three dimensions where W is irrelevant
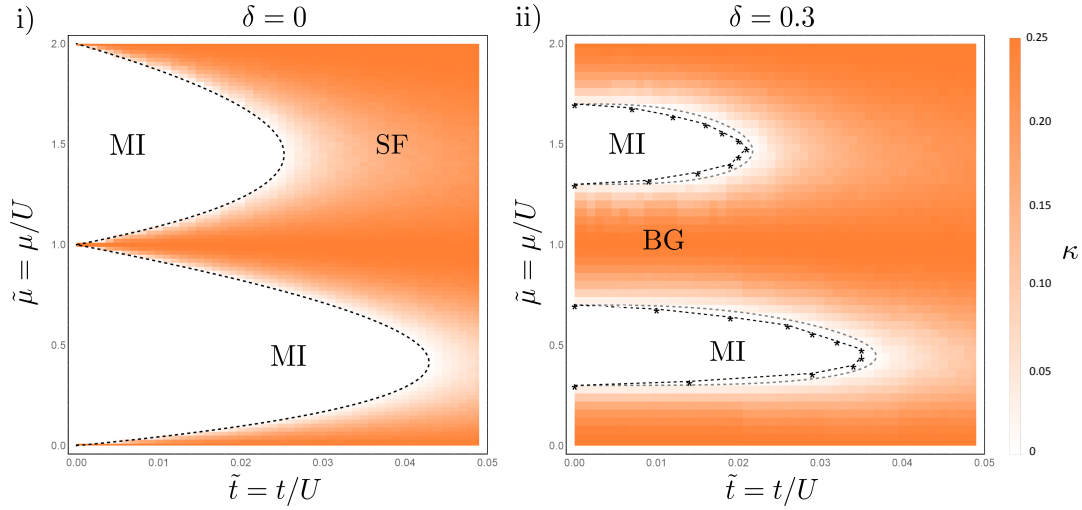
Figure 5.4 *Phase diagram constructed by extracting the variance (compressibility) of a 25×25 lattice for a grid of values of t and μ. i) The clean case, showing incompressible Mott lobes in excellent agreement with the conventional mean-field theory (dashed line). ii) After addition of disorder, the incompressible Mott lobes shrink and are now completely surrounded by a compressible phase. This plot was produced by averaging over ten disorder realisations. The black data points are the $\kappa = 0$ boundary determined from the numerics, joined with an interpolating function. The grey lines show the MI-BG boundary as determined from the renormalisation group calculation in Chapter 3. As before, the mean-field theory produces qualitatively correct results but slightly underestimates the size of the Mott lobes.*

quantitative difference, as expected, but the Gutzwiller ansatz performs well and gets the main features correct.

Even in the clean case, we see a strong suppression of compressibility in the vicinity of the tip of the lobe where the mean-field approach is furthest from its region of validity (in agreement with other models, such as Refs. [127]), though in neither case does the compressibility strictly vanish as predicted in Chapter 4. This is not a surprise, as the Mott glass calculation relies on effects not captured at a mean-field level and which were not expected to be present in a Gutzwiller ansatz calculation.

Already, we see that a measurement of compressibility in a disordered system should produce a quantitative difference large enough to be picked up in experimental measurements, however this alone is not sufficient to uniquely identify the Bose glass. A finite compressibility could simply be a superfluid phase - let us now test for that.

---

and thus was not included in the calculation. Given the negligible effect of RSB on the location of phase boundaries and the fact that we cannot be sure RSB even exists in two dimensions, the use of the replica symmetric RG is appropriate here.
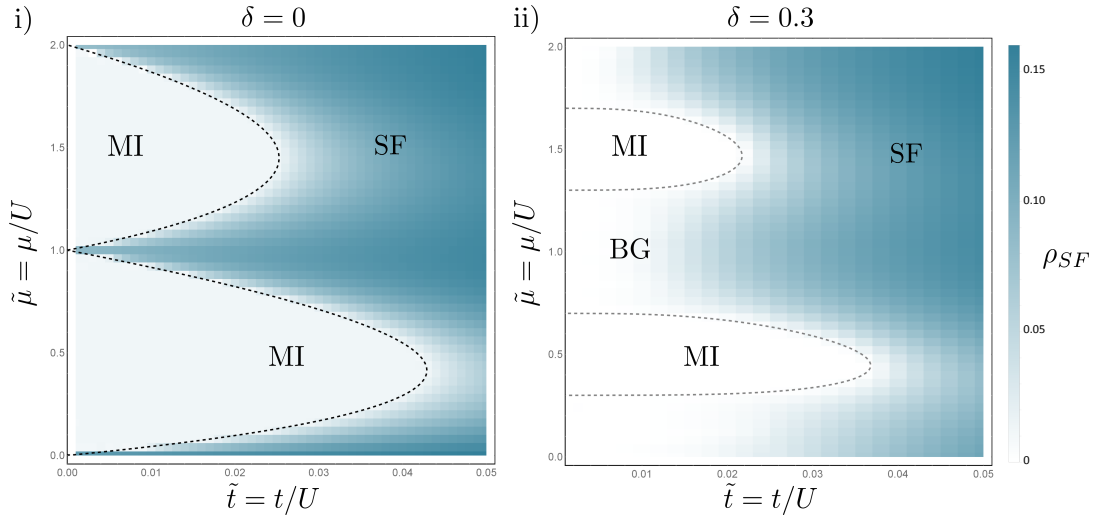
Figure 5.5 *Phase diagram constructed by extracting the superfluid stiffness of a 25 × 25 lattice for a grid of values of t and μ. i) The clean case, showing a clear separation between the Mott lobes and the superfluid. ii) Upon adding disorder, the region of zero superfluid stiffness increases in size, in contrast to what happened to the compressibility, indicating the presence of a compressible non-superfluid phase in the regions where the effects of disorder are strongest. This figure was produced from a single disorder realisation.*

Superfluid stiffness

Although not experimentally measurable in any straightforward way, the superfluid stiffness is a quantity which characterises the response of the system to a phase twist and is only non-zero in a phase which exhibits long-range phase coherence. In this model, only the superfluid exhibits such long-range phase coherence and consequently the superfluid stiffness ρ_{SF} is expected to be zero in both the Mott insulator and the Bose glass. In our framework, it is given by [145, 235]:

$$\rho_{SF} = \frac{E(\phi) - E(0)}{Nt\phi^2}. \tag{5.14}$$

Essentially this tests for the increase in energy when a superfluid with long-range phase coherence is forced to have a spatially varying phase, and is exactly analogous to the spin stiffness in a magnetic material. In our model, following Refs. [145, 235], we implement the phase twist by adding a direction-dependent complex tunnelling amplitude along one axis of the lattice $t \to t \exp[\pm i\phi]$ where the $\pm$ refers to left- and right-moving particles respectively.

We can again map out the phase diagram in terms of superfluid stiffness, here using a value $\phi = 0.05$, shown in Fig. 5.5. Strictly, the expression shown in Eq. 5.14 should be a derivative, implying the need to take the limit $\phi \to 0$,

however it can be well-approximated using a small finite value of the phase twist, as in Ref. [235]. For our purposes, as we are interested in determining the region where $\rho_{SF} \approx 0$ rather than in extracting the precise numerical value of ρ_{SF}, this approximation is sufficient.

The clean-case result performs as expected, with a clear boundary between the MI and SF phases visible in superfluid stiffness. The disordered case is less sharp (due to finite-size effects) but there is still a clear region for small hopping values around integer values of μ where the superfluid stiffness is arbitrarily small, and then a smooth increase in ρ_{SF} when the tunnelling amplitude is increased.

The superfluid stiffness is not strictly zero except at zero hopping, which is likely a finite-size effect. In a one-dimensional system [145, 235] calculation of the superfluid stiffness using the Gutzwiller ansatz results in a region of exactly zero superfluid stiffness which is associated with the Bose glass, however in one dimension even a single site with gapped excitations will destroy long-range phase coherence and result in zero superfluid stiffness. In two dimensions there are far more possible paths from one side of the lattice to the other and as a result one must go to intractably large lattices to obtain a strictly zero superfluid stiffness in this region.

Between the compressibility measurement (experimentally possible) and the superfluid stiffness (experimentally not), we see that the region surrounding the shrunken Mott lobes is highly compressible and has a vanishingly small superfluid stiffness. Together, these quantities are good indications that the phase in this region is indeed the Bose glass, however identifying it on the basis of two variables, one of which is not experimentally measurable, is unsatisfactory and not sufficient for a quantum gas microscope to be able to distinguish the BG phase. Now that we have confirmed that the Bose glass exists in this mean-field model, we are equipped to look for a single order parameter capable of unambiguously identifying it based purely on the local number fluctuations that a quantum gas microscope can measure.

Edwards-Anderson Order Parameter

Here we return to the Edwards-Anderson order parameter, previously discussed in Chapters 1 and 3. In the Bose-Hubbard model, the thermally averaged order parameter q takes the form:

$$q = \overline{\langle \hat{n}_i \rangle \langle \hat{n}_i \rangle} - \overline{\langle \hat{n}_i \rangle} \ \overline{\langle \hat{n}_i \rangle}, \tag{5.15}$$

which depends only upon local properties which are measurable by a quantum gas microscope. In Ref. [126] the authors identify q with q_{EA}; however, in keeping with the convention established earlier in this thesis, I restrict the usage of the symbol q_{EA} to mean the long-time correlation function rather than the thermally-averaged correlation function q.
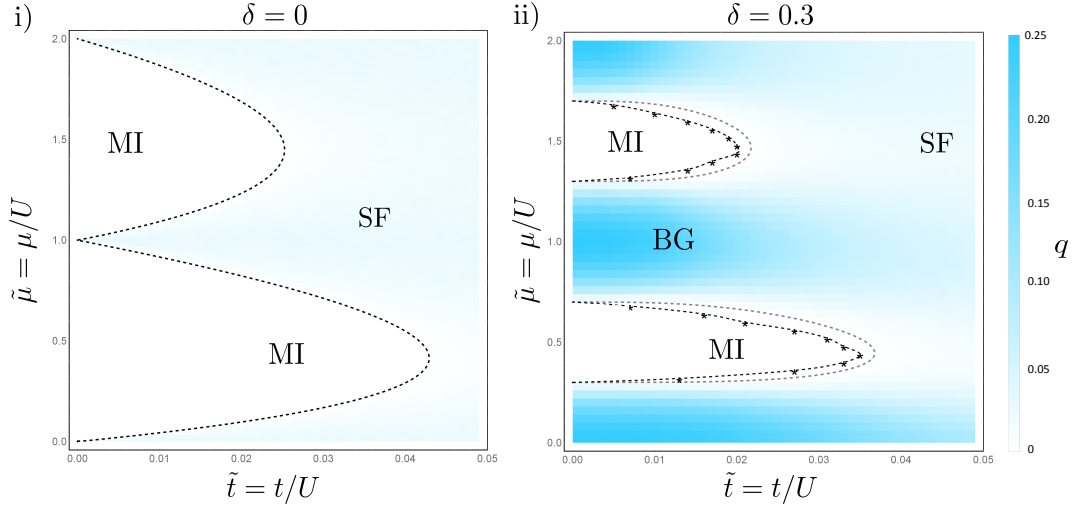
Figure 5.6 *Phase diagram showing the value of q across the $\tilde{t}, \tilde{\mu}$ plane. i) In the clean case, $q \approx 0$ everywhere with only lingering statistical fluctuations causing non-zero q in regions of the superfluid. The black dashed line is the mean-field MI-SF boundary. ii) In the disordered case, q is strongly visible in the Bose glass regions at small hopping, with a continuous crossover to the superfluid for larger values of hopping as predicted in Ref. [126]. This plot was produced by averaging over ten disorder realisations. Similarly to Fig. 5.4, the black data points are the $q = 0$ boundary extracted from the numerics and joined with an interpolating function, while the grey lines are the results of the renormalisation group calculation in Chapter 3. All following plots showing q values use the same colour scale.*

Strictly, q is the same on every single site and one could measure $\hat{n}_i$ repeatedly solely on site i within a single disorder realisation to obtain $\langle \hat{n}_i \rangle$, and then perform the disorder average of this over multiple disorder realisations without measuring any other lattice site. In practice, it is far more efficient to take into account that every lattice site has a different (random) disorder strength and so by averaging over the sites within a single snapshot, we essentially already perform a disorder average. For a sufficiently large system, there is no benefit to be gained by averaging over multiple disorder realisations and in fact the measurement can be performed with a single disorder realisation, with multiple runs needed only to extract the thermal averages.

Phase diagrams produced using q are shown in Fig. 5.6. The clean case is uninteresting, as $q = 0$ across the entire phase diagram other than small statistical fluctuations. In the disordered case, however, we see a clear and unmistakable region of $q \neq 0$ in precisely the regions where we expected to see it. As in the case of compressibility, due to parity the value of q saturates at $q_{max} = 0.25$. In fact, numerically it appears that q is always bounded from above by the value of the compressibility κ.

For larger values of hopping, q smoothly goes to zero making it impossible to identify a BG-SF transition using q. This is because even after the establishment of global superfluid coherence, the disorder still causes lingering inhomogeneities in the superfluid that are visible in the density snapshots. The same result was also found in Ref. [126]. In practice, one must define a cutoff q_{crit} where we identify $q < q_{crit}$ as being characteristic of an inhomogeneous superfluid and $q > q_{crit}$ as characteristic of a Bose glass.

The choice of cutoff is rather arbitrary, but for $q_{crit} \approx 0.2$ or greater we can certainly be confident we are in the Bose glass phase. In principle, we could use a different disorder distribution such as the bimodal hopping disorder looked at in Ref. [201] which might produce a sharper transition from the glass to the superfluid. However, since it is not currently possible to experimentally implement hopping disorder alone in a quantum gas microscope setup and the implementation of bimodal disorder would require a spatial light modulator (whereas on-site chemical potential disorder can be added to existing experimental setups relatively easily) we do not investigate this possibility any further.

Although we know that the transition from the Mott insulator to the superfluid is always via the Bose glass in the disordered case, the phase diagram in Fig. 5.6 shows a strong suppression of q in the vicinity of the tips of the lobes. The strong suppression of the compressibility at the tips of the lobes even in the clean case carries over to the disordered model as well (Fig. 5.4) and as κ acts as an upper bound on q, we find that q is also strongly suppressed in the vicinity of the tip. Importantly, it is not zero (as indicated in Fig. 5.6) but likely too small for experiments to reliably detect here. As we saw in Chapter 3, the universality class changes at the tip of the lobe and consequently the Gutzwiller mean-field approach breaks down in the vicinity of it and is less accurate. It may be that it is simply not a good description of the transition at the tip in two dimensions.

Whatever the reason for the suppression of q at the tips of the Mott lobes, the physics of this point is beyond the scope of our simple mean-field model and would require a full quantum Monte Carlo treatment to simulate accurately. In the regions where our model can be trusted to make accurate predictions (which coincide with the experimentally relevant values of hopping), we find a well-defined Bose glass phase and have shown that at mean-field level q is capable of clearly distinguishing the Bose glass from the Mott insulator and the superfluid.

At the time this work was performed, we believed it to be the first identification of an Edwards-Anderson order parameter as a mean-field order parameter for the Bose glass. After preparing the manuscript that became Ref. [218], we were made aware of previous attempts to define similar order parameters in mean-field models [239, 240], though neither the order parameters nor the techniques used are the same as those presented here. The method presented in this chapter more directly relates the Bose glass order parameter to density fluctuations and the idea of extracting it from quantum gas microscope images

remains unique to our work in Ref. [218]. While it goes without saying that the Gutzwiller approach will not get the critical behaviour correct, the ability to uniquely identify the Bose glass based on local properties means that a mean-field theory structured around such local properties is ideally suited for such an identification. Within the context of this thesis, however, it has one other important limitation.

5.6   Replica Symmetry Breaking

As in Ref. [126], we can define an 'overlap' order parameter in replica space:

$$q_{\alpha\beta} = \overline{\langle \hat{n}_{i,\alpha} \rangle \langle \hat{n}_{i,\beta} \rangle} - \overline{\langle \hat{n}_{i,\alpha} \rangle} \; \overline{\langle \hat{n}_{i,\beta} \rangle}, \qquad (5.16)$$

where α and β label different independent copies of the system with the same disorder realisation. In principle, by keeping the same disorder distribution between experimental runs and taking multiple measurements, an experiment can explore the entire space of replicas and reconstruct this order parameter. If $q_{\alpha\beta}$ varies with the choice of α and β, this would constitute experimental observation of replica symmetry breaking, a feat previously achieved in random lasers [241] but never in any condensed matter system.

Our mean-field numerics presented here always converge to the same minimum and as such, every 'replica' configuration is exactly the same. This ostensibly contradicts the results of Chapters 3 and 4, however this is not too surprising. While the density fluctuations which lead to the Bose glass are accessible by this mean-field description, it is unlikely that this variational wavefunction is able to fully explore the true configuration space of the model[2]. Quantum Monte Carlo simulations may be able to better probe this replica landscape.

Should experiments or Monte Carlo be able to distinguish replica symmetry breaking effects, this may also provide a way to more clearly distinguish the Bose glass from inhomogeneous superfluid regions with a small but non-zero value of q. The superfluid should truly thermalise and display no glassy freezing, inhomogeneous or not, therefore only the Bose glass will exhibit replica symmetry breaking. If $q_{\alpha\beta}$ is non-zero but independent of α and β, the phase is likely to be a superfluid. If $q_{\alpha\beta}$ depends on the choice of replica indices than the phase breaks ergodicity and must be a glassy phase.

---

[2]Alternatively, it could be that there is no RSB in two dimensions and that mean-field theory in three dimensions may find evidence for RSB, however this was outwith the scope of the work presented here.

5.7 Harmonic Trap

So far we have only considered 'flat' lattices, i.e. lattices at a uniform average chemical potential. In real quantum gas microscope experiments, the Gaussian profile of the laser beams results in a background harmonic confining potential.

Following Ref. [217] as closely as possible, we look at a 60×60 lattice. The variation of chemical potential in the presence of the harmonic trap is:

$$\mu_{loc}(r) = \mu_0 - \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2), \qquad (5.17)$$

where m is the mass of the atoms, in this case rubidium ^{87}Rb. There is a slight asymmetry present in the experimental results due to $\omega_x \neq \omega_y$. In our simulations we use a trap frequency $\omega \equiv \sqrt{\omega_x \omega_y} = 2\pi \times 77.3$Hz which is the the geometric mean of the two trapping frequencies used in the experiment

Clean case

In the zero-temperature clean limit, we reproduce the 'wedding-cake' structure seen in the experimental images of Ref. [217], minus the thermal fluctuations which are present in the experiment. Simulated snapshots at the same values of μ_0 as the experiments are shown in Fig. 5.7 and all harmonic trap snapshots are taken at $t/U = 1/300$ unless otherwise stated.

Disordered case

Once disorder is added into the model, we again see the same structure in Fig. 5.7 but now with density fluctuations due to the appearance of superfluid regions in the trap. Our simulations in the presence of disorder now look qualitatively very much like the finite-temperature experimental results from Ref. [217]. In our zero temperature simulation, the areas which display density fluctuations must be regions of Bose glass, but in a real experiment there will be both disorder and thermal fluctuations and it is important that we distinguish between them.

Finite temperature effects

As the density fluctuations due to disorder bear such a strong cosmetic resemblance to the experimentally observed fluctuations due to temperature, we need to check that q is still zero even in a clean, finite-temperature system. To that end, we move to a simplified version of the model in the limit of zero hopping where the thermal occupation of the lattice sites can be described by Boltzmann-weighted Fock states (as used in Ref. [217] to model the temperature of their
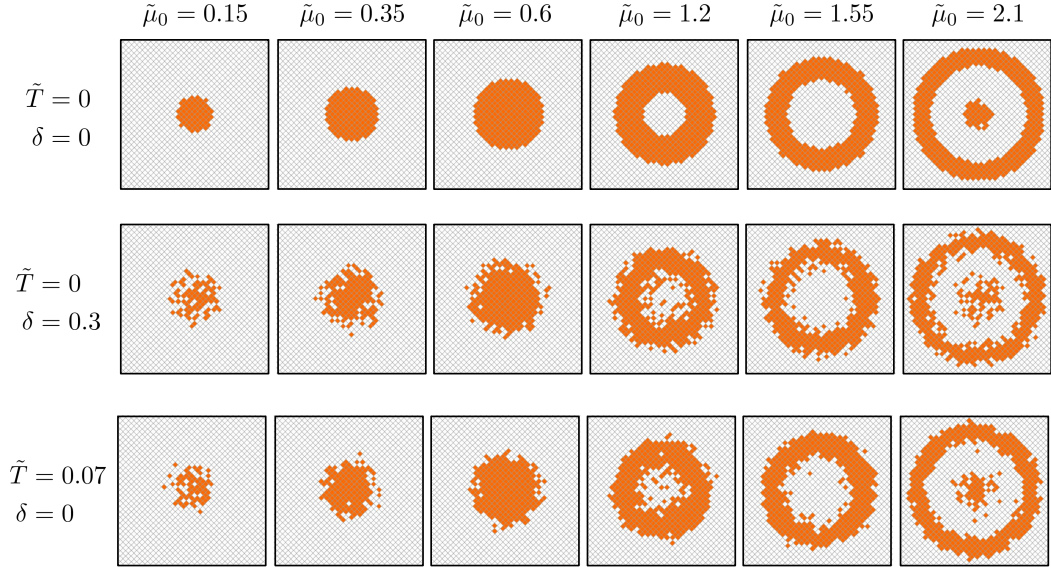
Figure 5.7 *Simulated snapshots of the density distributions in the harmonic trap, comparing the the effects of temperature $\tilde{T} = k_B T/U$. and disorder $\delta = \Delta/U$. In the clean case (top row) there are no fluctuations and the Mott insulating rings are almost perfectly ordered. Cosmetically, the disorder-induced density fluctuations (middle row) look a lot like the temperature-induced density fluctuations (bottom row), so it is important that our order parameter q is able to distinguish thermal fluctuations from true disorder-induced glassiness. The data for this figure, and all harmonic trap figures to follow, was taken by Liam Walker [238].*

system) according to the probability function:

$$P(n_i) = \frac{\exp[\beta(\mu_{loc}(r) - E_{n_i}^{(0)})]}{\mathcal{Z}(r)},$$ (5.18)

where $E_{n_i}^{(0)} = \left(\frac{U}{2}n_i(n_i - 1) - \mu_i n_i\right)|f_{n_i,i}|^2$ and $\mathcal{Z}(r) = \sum_n \exp[\beta(\mu(r) - E_n]$ is the partition function for a homogeneous system with chemical potential $\mu(r)$ and inverse temperature $\beta = 1/\tilde{T}$.

The results of the finite-temperature simulations are shown in Fig. 5.7. Cosmetically, these simulated snapshots bear a close resemblance both to the disordered zero-temperature snapshots but also to the thermal defects seen experimentally.

As discussed in Refs. [235, 242, 243], strictly there is no Mott insulator at finite temperature and it is replaced with a compressible non-superfluid phase, however the compressibility of this phase is so small that for all intents and purposes the non-superfluid phase may still be considered a (quasi-) Mott insulator. As the experiments we are interested in comparing with see no signs indicative of normal fluid behaviour [217], we do not consider it any further.

Edwards-Anderson Order Parameter

From the density snapshots in the presence of the harmonic trap, we can now calculate the regions in the trap which have non-zero values of q, as shown in Fig. 5.8. Upon the addition of disorder, we would now expect to see alternating Mott insulating and Bose glass rings, which is exactly what our simulations show.

We can extract q in multiple different ways, either using an angular average over contours of equal chemical potential or by calculating it on each lattice site individually. In the first case, the angular average acts as a disorder average as each site has the same bare chemical potential $\mu(r)$ but a different random offset from it. This allows an experiment to detect regions of Bose glass using a single disorder realisation with multiple runs required only in order to build up thermal averages. This is convenient as it resetting the disorder distribution must be done manually, so being able to resolve the Bose glass with the fewest possible measurements is advantageous.

That said, there are inherent problems in taking angular averages on a square grid. The resulting plot of q is blocky and this procedure can lead to strange artefacts even in the clean case where the 'ring' being averaged over contains regions of two different MI states, which our order parameter picks up as a non-thermal density fluctuation and gives a false positive result. A much cleaner, if experimentally more challenging, method is to construct q on each lattice site individually using multiple disorder distributions. This results in a much more accurate numerical value of q.
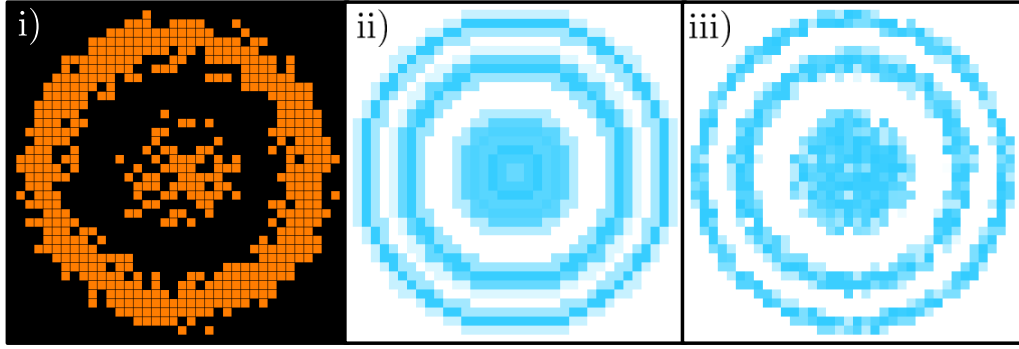
Figure 5.8 *A comparison of the site-averaged q with the azimuthally averaged q. i) A simulated density snapshot taken at $\tilde{\mu}_0 = 2.1$. ii) Azimuthally averaged q produced from 10 thermal averages over a single disorder realisation showing clear regions of Bose glass. The advantage to this is that it requires less measurements as the angular average over sites of constant chemical potential essentially does the disorder average for us. The disadvantage is that the inner regions are averaged over fewer sites than the outer regions and the overall appearance of the plot is rather blocky. iii) Site averaged q produced from 10 thermal averages over 10 disorder realisations. Experimentally more challenging as it requires 10 different disorder distributions, this one is free from the artefacts of the azimuthally averaged figure. Data for this figure was taken by Liam Walker.*

Both of these methods are shown in Fig. 5.8, with panel ii) showing the angular-averaged result of a single disorder distribution and ten thermal averages and panel iii) showing the site-averaged result of ten disorder distributions each with ten thermal averages. Even the large number of measurements required in the latter case remains within reach of experiments. This is a clear indication that current-generation quantum gas microscopes are capable of detecting the Bose glass in an experimentally feasible number of runs.

Before we can conclude that q is a good experimental order parameter, we need to see how it looks in the case of the finite-temperature simulations. Fig. 5.9 shows the site-averaged q for a variety of combinations of temperature and disorder strength. Though by construction q is identically zero in the clean system even at finite-temperature, over a finite number of measurements we still see a small non-zero value due to statistical fluctuations. By comparison with the disordered case, however, the signal from thermal fluctuations is at least an order of magnitude smaller across all lattice sites, confirming that q is capable of distinguishing thermally-induced density fluctuations from true disorder-induced glassiness.

This was the goal of our simulations and is our main result - proof that q can be directly measured in a quantum gas microscope. The Edwards-Anderson
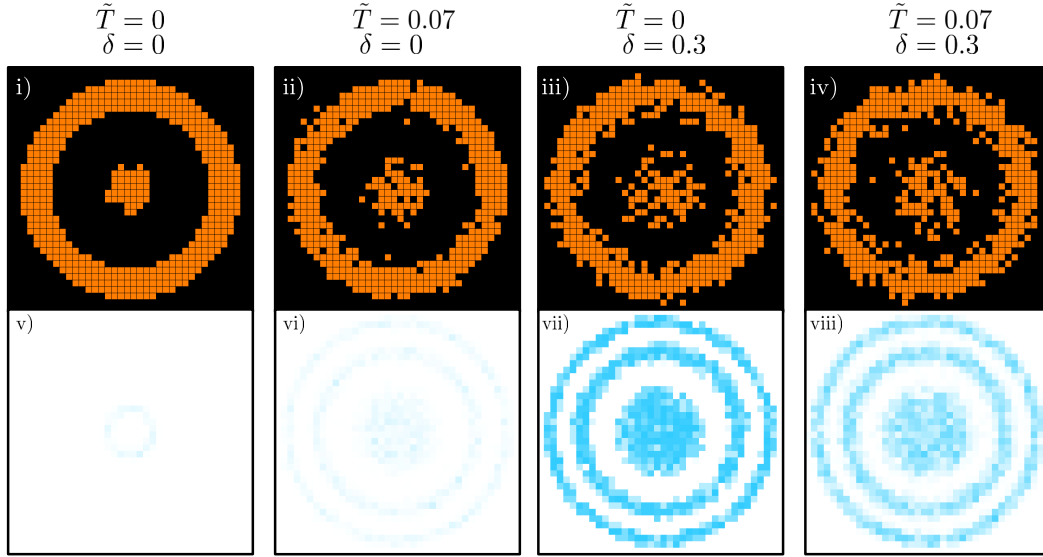
Figure 5.9   *A comparison of thermal fluctuations and disorder-induced fluctuations. Panels v) to viii) were generated using ten thermal averages and ten disorder averages, where applicable. Panels i) to iv) show simulated density snapshots, as in Fig. 5.7, at a range of temperature and disorder strengths. In particular, panels ii), iii) and iv) are cosmetically very similar and it is not possible to tell whether the fluctuations are caused by temperature or disorder. Panels v) to viii) show plots of site-averaged q to illustrate which regions of the trap are Bose glass. Panel v) is entirely blank as there are no fluctuations of any kind. Panel vi) is almost blank aside from statistical fluctuations that lead to a very small non-zero value of q. Panel vii) shows clearly that $q \neq 0$ in the Bose glass at zero temperature, where all of the fluctuations are disorder-induced and panel viii) puts both thermal and disorder-induced fluctuations together and shows that temperature reduces the magnitude of q. Eventually the Bose glass will 'melt' and $q \to 0$ as the temperature is increased. Data for this figure was taken by Liam Walker.*

order parameters have never been directly measured in any strongly-interacting quantum system to date, and the possibility of directly measuring q in ultracold atomic gases with designer Hamiltonians will open up an entirely new tool with which to probe the behaviour of disordered quantum systems.

In the remainder of this chapter we will quantify a few experimental limitations and suggest further ways to increase the visibility of q in real physical systems. In particular, panel viii) of Fig. 5.9 shows that the combination of both disorder and finite-temperature leads to a suppression of q over the zero-temperature case - just how damaging are the effects of temperature to the measurements proposed here?
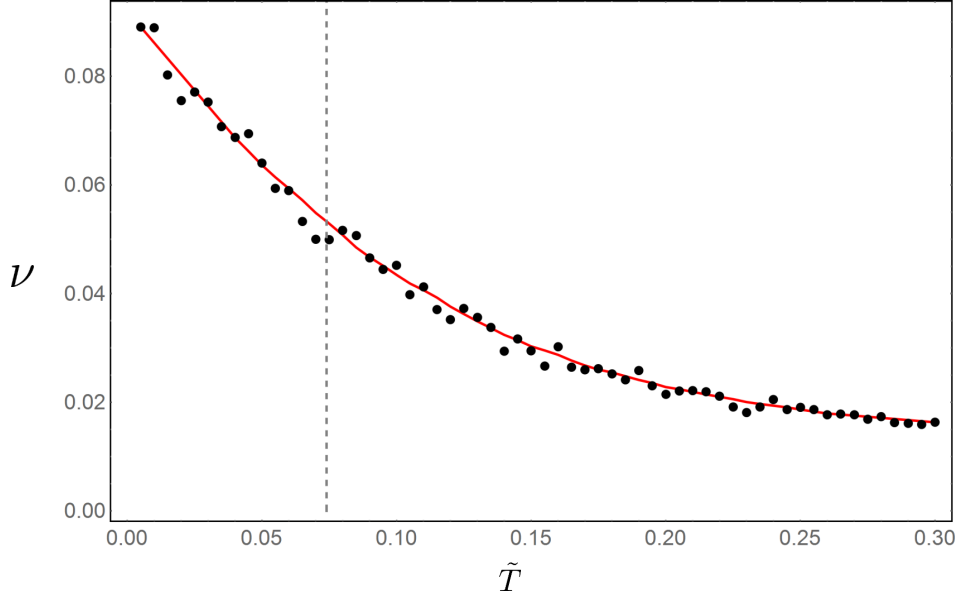
Figure 5.10 *The visibility ν decreases as the temperature $\tilde{T}$ increases and the Bose glass 'melts'. The black points indicate the value of the visibility obtained from averaging ten snapshots of q as shown in Fig. 5.8ii) and the red line is a guide to the eye fit of these points. The dashed grey line indicates the lowest experimentally achieved temperature shown in Ref. [217]. Figure produced using analysis code written by Liam Walker.*

Visibility

Remaining with the simplified zero-hopping model, we can calculate q in the disordered, finite-T model and see how the Bose glass 'melts', i.e. how the visible signal of q is diluted with temperature and eventually the signal vanishes. Using a standard metric of visibility:

$$\nu = \frac{q_{RMS}}{\langle q \rangle}, \tag{5.19}$$

where q_{RMS} is the root-mean-square value of q and ν effectively measures the contrast within an entire image, we can extract how distinct the finite-q regions are from the zero-q regions. Fig. 5.10 shows a plot of visibility versus temperature. As the temperature is increased, thermal fluctuations grow in magnitude and eventually dominate over the disorder-induced density fluctuations, washing out the glassy order parameter almost entirely. Importantly, this result is not crucially dependent on the visibility metric we use - various other metrics all give qualitatively similar results.

In practice, beyond a temperature of around $\tilde{T} = 0.1$ the measured q is too

faint to be reliably distinguished. The lowest temperature reported in Ref. [217] was $\tilde{T} = 0.074$ and so even the very early single-atom resolved quantum gas microscope experiments were already well within the region where the Bose glass can be distinguished, with more recent experiments now capable of even lower temperatures.

5.8 Alternative trap geometries

The biggest problem with the experiments as they are conventionally set up is not the temperature, but rather the fact that the Bose glass only occurs in very narrow rings due to the relatively steep gradient of the harmonic trapping potential, limiting the region over which they can be measured. This would rule out measurements of longer length-scale properties such as the spacing of superfluid regions or direct observation of the glass to superfluid transition, still the subject of intense theoretical debate. Thus far we have shown what the current experiments would be able to see and how well they will be able to see it - now we deviate from previously performed experiments and briefly suggest a way to make the Bose glass visible over a larger region.

Using different combinations of laser beams, it is possible to change the trapping geometry. With spatial light modulators and other such apparatus, any arbitrary trapping potential can be generated. Here we consider a slowly-varying 'shallow' trap with hard walls at either side acting as boundaries. This could be generated by superimposing a blue-detuned beam on top of the confining potential to cancel out most of the curvature, followed by applying a Laguerre-Gauss beam to generate the boundaries at the edges. The advantage to this trapping potential is that the gradient is much smaller, meaning the features of interest are more spread out across the physical system rather than being confined to rings of only a few lattice sites' width.

The results of our simulations are shown in Fig. 5.11. Instead of the wedding-cake ringed structure, we now have a large region of approximately homogeneous phase. This now opens up the possibility of looking experimentally at relatively large segments of Bose glass phase and extracting properties such as the spacing of the rare superfluid regions or the scaling of the compressibility close to the phase transition.

5.9 Conclusion

The work presented in this chapter is a proof of concept study to show that current generation quantum gas microscopes are capable of spatially resolving the Bose glass phase. Using a mean-field wavefunction in two dimensions, we find the clean case calculations are in good agreement with the available experimental
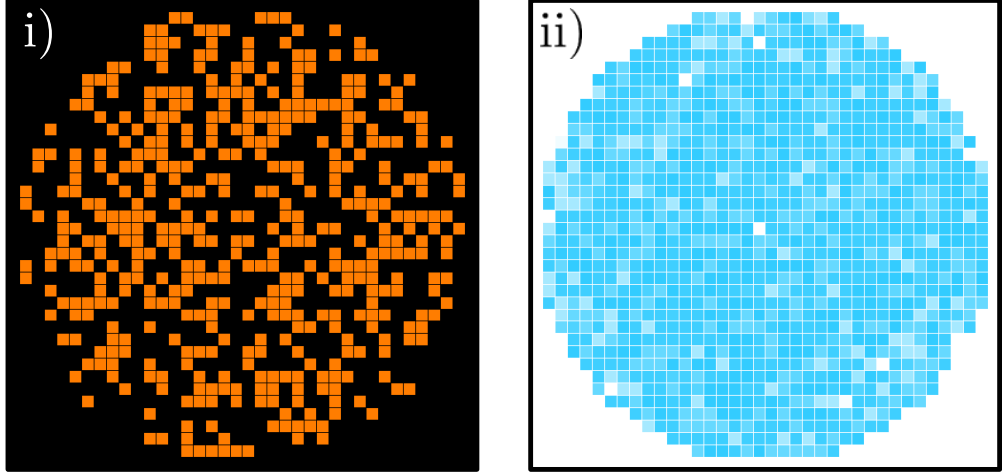
140

Figure 5.11   *The weakly varying trap produced by setting $\omega \to \omega/5$ allows for larger regions of the Bose glass phase than a typical harmonic trap. i) A single simulated density snapshot, with $\tilde{\mu}_0 = 2.1$. ii) The Edwards-Anderson order parameter of this new trap geometry, generated from ten disorder realisations each with ten thermal averages. This shows that the entire region in panel i) is Bose glass. With this identification, we could now begin to look into longer wavelength properties of the Bose glass in more detail. Data for this figure was taken by Liam Walker.*

data [217] and this allows us to state with confidence that the density fluctuations predicted by the disordered calculation are realistic and that we expect reasonable quantitative agreement with experiments.

Here, we have shown that quantum gas microscopes can distinguish the Bose glass from the MI and SF phases and that spatially resolved experiments with single-site addressability allow for the Edwards-Anderson order parameter to be precisely measured. The proposed order parameter q takes a large non-zero value in the Bose glass phase, as well as a smaller non-zero value in inhomogeneous regions of the superfluid.

The ability to measure the Edwards-Anderson order parameter also opens up the possibility for experiments to explicitly test for the presence of replica symmetry breaking in the model, a feat achieved on random laser systems comparatively recently though outwith the scope of the mean-field numerics presented here. The controllability and single-site addressability afforded by quantum gas microscopes allows a far more detailed probe of glassiness and the replica structure of strongly correlated quantum systems than bulk measurements on spin glasses and other solid state systems.

The result from this work is a call to arms for experimentalists to test these

predictions with the aim of achieving the first unambiguous observation of the Bose glass phase, but also the establishment of a new tool to probe disorder in strongly correlated quantum systems. Similar single-site addressability is currently coming on line in STM systems and may soon allow this sort of measurement to be performed on solid state systems, however STMs as they currently stand are not yet capable of performing the measurements suggested here. For the foreseeable future, only ultracold atomic gases can offer the full control and detail required for these measurements.

Chapter 6

Conclusion

The problem of disorder in strongly interacting quantum systems has been experiencing something of a renaissance in recent years, spurred on by new advances in both experimental and theoretical techniques. We are moving from the age of passively understanding quantum phenomena in materials to actively manipulating them. From high-temperature superconductors to quantum computers, disorder and impurities play important roles in many of the most exciting quantum technologies of our time. If we can harness the effects of disorder in quantum materials, we may be able to create new designer materials with properties carefully tailored to enhance the quantum mechanical behaviour we are most interested in making use of.

The recent enthusiasm for investigating disorder is, a touch ironically, a testament to the precision and purity achieved by experiments. Experimenters can now create almost perfectly clean systems to which disorder and impurities can be added in a highly controlled way, even at the level of single lattice sites or atoms, allowing us to study the role of disorder in unprecedented detail. As one experimenter remarked to me recently, *"I've spent my entire career trying to remove disorder from my experiments, and now you're telling me to put it back in?"*.

On the theoretical side, the discovery of many-body localisation (MBL) [63, 244] has led to a great deal of interest in the interplay between disorder and interactions in many-body systems. MBL is an extreme form of disorder-induced localisation in interacting quantum systems where every degree of freedom becomes non-ergodic, in contrast to typical glassy systems where only some degrees of freedom break ergodicity. The question of precisely how disorder can lead to the breaking of ergodicity in strongly interacting quantum systems has therefore become highly relevant.

There is still a great deal about disordered quantum systems which we do not understand, even in quantum glass phases which have been investigated for years. In this thesis, I have studied a number of different aspects of quantum

glass phases in order to help put our understanding of them on a firmer footing. Here, I'd like to summarise my work in the wider context of the field as a whole and discuss possible extensions of this work and future directions of interest.

Quantum critical fluctuations

In **Chapter 2** we saw disorder that had no effects at mean-field level could nonetheless lead to significant changes in the quantum critical behaviour of an itinerant ferromagnet. Quenched charge disorder, which had no effect on the nature of the paramagnetic or ferromagnetic phases, led to local changes in the pitch of the fluctuation-induced spiral phase which in turn caused the magnetic correlations to become short-range. The resulting 'melted spiral' phase was named a helical glass and is characterised by the highly unusual temperature-dependence of the correlation length and spiral ordering wavevector.

This is a particularly interesting example of the role of disorder in a quantum system in that the disorder only leads to changes in the fluctuation-induced phase due to the way that it disorder couples to the Goldstone modes, as opposed to introducing more conventional disorder effects such as spin-glass-like frustration. In candidate material CeFePo [84], signs of glassy behaviour were detected but there appeared to be an additional long-range ordered phase in between the ferromagnetic and putative glass phase. This could well be the helical glass in a region of the phase diagram where its correlation length is larger than the sample size however no follow-up work has been done to test for this.

This was the only work I did during my PhD which involved fermionic quantum order-by-disorder and it is largely unconnected to the rest of the work in the thesis. Further work on the order-by-disorder front has been performed since my involvement [245, 246], though it has been confined to clean systems. Perhaps the biggest outstanding challenge in the field of fermionic order-by-disorder is the extension to antiferromagnetic quantum critical points, for which the fluctuation corrections are extremely difficult to compute.

In the broader context of itinerant ferromagnetic quantum criticality, disorder remains an active area of interest. A comprehensive review may be found in Ref. [247]. Of the outstanding questions in the field, two of the most interesting ones from the point of view of disorder are i) Griffiths effects in ferromagnetic metals and ii) the potential role of disorder in suppressing the soft modes associated with the fluctuation-induced behaviour, in contrast to the disorder considered in Chapter 2. The work in this thesis was a perturbative treatment of weak disorder facilitated by the order-by-disorder framework, however future dedicated studies of these strong disorder effects could perhaps lead to a more complete description of the interplay between disorder and critical fluctuations.

Glassiness and non-ergodic effects

In **Chapter 3**, I showed the presence of one-step replica symmetry breaking in the Bose glass phase, a hallmark of 'true' glassy freezing behaviour and a sign of the non-trivial breaking of ergodicity. It is thought that the disordered Bose-Hubbard model exhibits many-body localisation at non-zero temperatures, where thermal excitations are allowed to occur.

The crossover between quantum glass and many-body localised phases is largely unexplored and is an extremely interesting open question. Understanding the physics behind the Bose glass and studying its finite temperature behaviour could lead to further understanding of the MBL-to-glass transition, and indeed some very recent experiments have begun to find evidence for many-body localisation in disordered Bose-Hubbard models [248].

By itself, the work in Chapter 3 simply showed that the Bose glass was more exotic in nature than previously realised. The real importance of this result was not realised until the work presented in **Chapter 4** in which I showed that replica symmetry breaking was vital to a correct description of the thermodynamics of the Bose glass and its magnetic analogues. This result acts as a strong signature of the non-trivial nature of the ergodicity breaking in this system and a warning sign to any calculations of the thermodynamic properties of quantum glasses that did not allow for the possibility of replica symmetry breaking.

This is particularly interesting in the context of the MBL-to-glass crossover. In an MBL phase, all degrees of freedom are non-ergodic whereas in the Bose glass this is likely only true for some degrees of freedom. This leads on to enough new questions to fill a whole other thesis. How does a material move from a quantum glass to an MBL phase? Do the degrees of freedom successively freeze out, or is there a sharp transition at which point the system moves from glassy dynamics to a true MBL phase? Is MBL perhaps just an extreme case of a quantum glass phase that occurs under conditions which freeze out every degree of freedom, or is it really qualitatively different in some way?

More generally, given the implications of replica symmetry breaking for thermodynamic quantities, what signature would it have on other properties such as specific heat capacity, a quantity easily accessible to quantum Monte Carlo simulations on such systems? Within our framework, finite temperature poses significant problems for the calculation of the Green's functions, meaning that a numerical approach is likely required here.

In light of these questions and the potential link with MBL, it would be interesting to calculate the entanglement entropy of disordered Bose-Hubbard-like systems at both zero and non-zero temperature to look for any qualitative changes. In a system which displays MBL, the entanglement entropy follows a distinctive area-law scaling form. Analytic calculations of the Renyi entropy have been performed in similar models [249], though never with disorder, and there is no fundamental reason why they could not be conducted to study these

glassy phases. The biggest problems here would likely come from the extension to non-zero temperatures, but while this is a significant technical challenge it isn't fundamentally impossible.

The Mott glass

The strong indications we see in **Chapter 4** for the presence of a novel incompressible Mott glass phase at the multicritical points at the tip of the Mott lobes is intriguing and deserved of further study. It was previously thought that the Mott glass can only exist under very specific conditions, with numerous studies each claiming a slightly different set of different conditions to be necessary. Our work here suggests that the Mott glass is not so fragile as has previously been proposed and that the emergent symmetries of the Bose-Hubbard model at the multicritical points are sufficient to stabilise it.

This also provides a potential explanation for the lasting controversy as to whether there was a direct Mott insulator to superfluid transition at the tip of the Mott lobes. Though it is now known there is no direct transition [147, 148], any technique which used compressibility as the metric to signify the presence of a glassy phase would have missed the Mott glass and may have erroneously identified the MG-SF transition as an MI-SF transition, leading to the wrong conclusion.

There are various future directions in which this work could be taken. The obvious one would be a more detailed functional renormalisation group study specific to the $z = 1$ multicritical point in the model, again allowing for full Parisi replica symmetry breaking in its most general form but this time including the full effects of the interaction vertex on the replica structure. It could be that there is no new physics here and that the Mott glass comes about purely through the statistical emergence of particle-hole symmetry at the tip of the Mott lobes, or it could be that the restoration of this symmetry leads to qualitatively new physics not captured by the analysis in Chapter 4. This is a question that would also be interesting for future quantum Monte Carlo studies.

To really understand the Mott glass phase itself, a comprehensive review of all Mott glass works will be necessary, as it has appeared in such a wide variety of models. In particular, it will be important for future work to ascertain whether they really are all driven by the same underlying physics or whether the name 'Mott glass' is just a convenient name which is currently being used to describe a number of different effects leading to incompressible Griffiths phases in disordered systems.

New experimental probes of disordered quantum gases

The work in **Chapter 5** has suggested a way for future experiments to investigate the local properties of strongly interacting quantum systems, motivated by the experimental feasibility of directly measuring the Edwards-Anderson order parameter, a feat never achieved in any condensed matter system to date. Though the work in this chapter was explicitly concerned with a bosonic system, the mechanics of the experimental setup are easily transferable to disordered fermionic systems. This could lead to a sophisticated and highly configurable setup with which to probe the effects of disorder on model Hamiltonians such as those thought to be important for high-temperature superconductivity.

The numerical approach presented here could be extended in many ways. Different types of disorder could be investigated, such as hopping disorder or disorder with a bimodal (rather than a continuous) distribution. The bimodal distribution may allow for a clearer onset of q and a more visible experimental signature. Pure hopping disorder would eliminate the density inhomogeneities which lead to a non-zero q in the superfluid, meaning that $q \neq 0$ only in the Bose glass phase. This would allow q to very clearly delineate the boundary of the Bose glass phase and better distinguish it from the MI and the SF.

While some very early testing not included in this thesis has shown promising results, the problem with hopping disorder is that it predominantly affects the area around the tip of the Mott lobe where our mean-field approach is least accurate and the value of q is the smallest. There is still work to be done in determining precisely what form of disorder would lead to the most visible experimental signature of the Bose glass and precisely what experimental setup one would have to construct to realise this type of disorder. The work presented here was confined to the types of disorder which could be realised relatively easily by current experiments, but now that we have shown that the Edwards-Anderson order parameter can be measured in these setups, future work can go on to suggest different and better experimental systems in which to improve upon the measurements suggested in this thesis.

Additionally, finite temperatures could be more fully included in our numerics beyond the zero-hopping limit and we could even add time-dependence to the Gutzwiller coefficients. This would allow us to investigate some non-equilibrium measurements which quantum gas microscopes are capable of performing, such as transport measurements. Importantly, however, the Gutzwiller approach does not fully take into account interactions or any possibility of entanglement and therefore cannot capture any many-body localisation effects, which make it ill-suited for investigating some of the most interesting open questions.

The goal of these numerics, however, was never to formally study the disordered Bose-Hubbard model in two dimensions. The point of this calculation was simply to show that quantum gas microscopes are capable of more than

that which they have been used for thus far, with the hope of motivating future experiments on highly controllable disordered systems which offer a new window into disordered strongly correlated quantum systems.

Were we to pursue the numerical analysis angle, a technique that goes beyond this mean-field analysis would be required. In one dimension, density matrix renormalisation group and matrix product states are well-established techniques and these would be much more suitable methods to look in more detail at the non-ergodic effects. In two dimensions, the best available numerical tool is still quantum Monte Carlo. If we want a more quantitatively accurate tool with which to simulate quantum gas microscopes and look for more complex physics, quantum Monte Carlo is currently the most appropriate option available.

Outlook

My hope is that in the future, the work in this thesis will be built upon by myself and others and that the necessary experimental and numerical investigations will be conducted to either support or negate the conclusions presented here. In particular, I would be very keen to see quantum gas microscopes begin to be used in the near future to conduct some of the experiments suggested in Chapter 5.

Longer term, I think the question of the glass-to-MBL transition is one of the most exciting open questions in the field. Our lexicon of localisation phenomena in disordered systems has become much richer with the discovery of many-body localisation, and with that increase in knowledge has come an increase in the number of interesting questions we may ask.

This body of work represents the bulk of my original research conducted during the four years of my PhD studies. Taken together, the unifying theme of all of my work has been that the addition of disorder leads to rich new physics with sometimes counter-intuitive effects. Using a mixture of perturbative techniques, renormalisation group and mean-field numerics I've investigated the behaviour of an itinerant ferromagnet, lattice gases of bosons and insulating dimerised antiferromagnets. I've shown how disorder can lead to the formation of new phases, I've studied the properties of several of these phases and throughout I've maintained contact with experiments, building upon results previously known and suggesting new ways in which future experiments could probe the glassy phases studied in this thesis.

The systematic study of randomness in strongly interacting quantum systems will always be a difficult problem. The work presented in this thesis represents my contribution thus far to the efforts to solve this problem.

148

Appendix A

Hubbard-Stratonovich Transform

The Hubbard-Stratonovich transform is a powerful technique that allows us to decouple a term in the action at the cost of introducing an auxiliary field. It was first used by Stratonovich [88] and later popularised by Hubbard [89]. The technique relies on a deceptively simple result:

$$
\frac{\pi^N}{\det \mathbf{A}} e^{\phi^\dagger \mathbf{A}^{-1} \phi} = \int \mathrm{d}\psi^\dagger \mathrm{d}\psi e^{-\psi^\dagger \mathbf{A} \psi + \psi^\dagger \phi + \phi^\dagger \psi}. \tag{A.1}
$$

This is simply a Gaussian integral. By performing this 'in reverse', so to speak, we may use it to decouple a quartic interaction so the resulting expression is only quadratic, as used repeatedly in the main body of the thesis. This then lets us integrate out the original fields, leaving a partition function entirely in terms of the new auxiliary fields. An alternative representation may be found in Ref. [40] in the context of an Ising model.

In Chapter 2, I employ this transformation to decouple the hopping term, but in Chapters 3 and 4 I use it to decouple the interaction term. This transform is not a blunt instrument and can be used in many different ways. It is not unique [9] and there are multiple inequivalent choices of auxiliary fields corresponding to the different decouplings one may choose. As discussed in more detail in Ref. [9], the Hubbard-Stratonovich transform is mathematically exact, however different choices of decoupling dictate the form of the low-energy field theory obtained in the end. Different decouplings will be more suited to solving certain problems, and so the decoupling must be performed with an eye for the problem one wishes to solve (e.g. decoupling in the pairing channel to study superconductivity problems). An unphysical decoupling will result in a meaningless (or at best, difficult to interpret) low-energy effective field theory.

Appendix B

The Replica Trick

We take the disorder average using the replica trick:

$$\overline{\mathcal{F}} = -T\overline{\ln(\mathcal{Z})} = \lim_{n \to 0} \frac{\overline{\mathcal{Z}^n} - 1}{n}. \tag{B.1}$$

The cumulant expansion used repeatedly in the main body of the thesis is as follows:

$$\mathcal{Z} = \int \mathcal{D}_\psi \exp[-S(\psi)] \tag{B.2}$$

$$\overline{\mathcal{Z}^n} = \int \mathcal{D}_{\psi_i} \overline{\exp\left[-\sum_i^n S(\psi_i)\right]} \tag{B.3}$$

$$\approx \int \mathcal{D}_{\psi_i} \left[1 + \overline{\sum_i^n S(\psi_i)} + \frac{1}{2}\overline{\sum_{i,j}^n S(\psi_i)S(\psi_j)}\right] \tag{B.4}$$

$$\approx \int \mathcal{D}_{\psi_i} \exp\left[\overline{\sum_i^n S(\psi_i)} + \frac{1}{2}\overline{\sum_{i,j}^n S(\psi_i)S(\psi_j)} - \frac{1}{2}\overline{\sum_i^n S(\psi_i)}^2\right], \tag{B.5}$$

By inserting the appropriate expression for the action, performing the averages and discarding any irrelevant terms, the disorder-averaged actions in Chapters 2, 3 and 4 may be calculated in precisely this manner.

Appendix C

Parisi Matrix Algebra

Consider an $n \times n$ matrix $\mathbf{P} = (P)_{\alpha,\beta}$ in replica space. To parameterise k-step replica symmetry breaking, we define the sequence:

$$n = m_0 > m_1 > m_2 > \ldots > m_k > m_{k+1} = 1, \tag{C.1}$$

with $m_i/m_{i+1} \, \epsilon \, \mathbb{Z}_{\geq 2}$ such that all blocks are of integer length.

A Parisi matrix can be defined as:

$$\mathbf{P} = \tilde{p}\mathbb{1} + \sum_{i=0}^{k} p_i \tilde{\mathbf{A}}_i, \tag{C.2}$$

with auxiliary $n \times n$ square matrices $\tilde{\mathbf{A}}_i$ are blockwise diagonal with $m_i \times m_i$ blocks defined as:

$$\tilde{\mathbf{A}}_i = \mathbf{A}_i - \mathbf{A}_{i+1}. \tag{C.3}$$

Subtracting the smaller matrix ensures the Parisi form is adhered to. These matrices describe the shaded areas shown in Fig. 1: the matrix $p_0\tilde{\mathbf{A}}_0$ has entries p_0 in the light grey regions and is zero elsewhere. All other off-diagonal shaded regions can be described similarly in terms of the corresponding $p_i\tilde{\mathbf{A}}_i$ matrices.

The Parisi matrices form a closed group (under multiplication and inversion), and their algebra can be derived from that of the block matrices $\mathbf{A}_i$.

$$\mathbf{A}_i\mathbf{A}_j = m_{\max(i,j)}\mathbf{A}_{\min(i,j)}. \tag{C.4}$$

This leads to:

$$\tilde{\mathbf{A}}_i\tilde{\mathbf{A}}_j = \begin{cases} (m_j - m_{j+1})\tilde{\mathbf{A}}_i & \text{for} \quad i < j \\ (m_i - m_{i+1})\mathbf{A}_i - m_{i+1}\tilde{\mathbf{A}}_i & \text{for} \quad i = j \\ (m_i - m_{i+1})\tilde{\mathbf{A}}_j & \text{for} \quad i > j \end{cases}. \tag{C.5}$$

We can show the Parisi matrices form a closed group under multiplication, as

well as derive the form of the inverse matrix $\mathbf{Q}$ via:

$$\mathbf{R} = \mathbf{PQ} = \left(\tilde{p}\mathbb{1} + \sum_{i=0}^{k} p_i \tilde{\mathbf{A}}_i \right) \left(\tilde{q}\mathbb{1} + \sum_{j=0}^{k} q_j \tilde{\mathbf{A}}_j \right) \tag{C.6}$$

$$= \tilde{r}\mathbb{1} + \sum_{i=0}^{k} r_i \tilde{\mathbf{A}}_i, \tag{C.7}$$

where it can be shown that:

$$\tilde{r} = \tilde{p}\tilde{q} + \sum_{j=0}^{k} p_j q_j (m_i - m_{j+1}), \tag{C.8}$$

$$r_i = \tilde{p}q_i + p_i\tilde{q} - p_i q_i m_{i+1} + \sum_{j=i+1}^{k} (p_i q_j + q_i p_j)(m_j - m_{j+1})$$

$$+ \sum_{j=0}^{i} p_j q_j (m_j - m_{j+1}). \tag{C.9}$$

Therefore the product of two Parisi matrices remains a Parisi matrix, and the group is closed under multiplication.

Analytic Continuation

We can describe the elements of a Parisi matrix by the piecewise continuous function $p(u) = p_i$ for $m_{i+1} \le n < m_i$. We can take the continuum limit of this as follows:

$$\sum_j (m_j - m_{j+1}) = (n - 1) \to \int_1^n dv, \tag{C.10}$$

$$\sum_{j \le 1} (m_j - m_{j+1}) \to \int_u^n dv, \tag{C.11}$$

$$\sum_{j > 1} (m_j - m_{j+1}) \to \int_1^u dv. \tag{C.12}$$

This results in the following expressions for the matrix elements of the inverse Parisi matrix $\mathbf{R}$:

$$\tilde{r} = \tilde{p}\tilde{q} + \int_1^n dv\, p(v)q(v), \tag{C.13}$$

$$r(u) = \tilde{p}q(u) + p(u)\tilde{q} - u p(u)q(u)$$

$$+ \int_1^u dv\, [p(u)q(v) + p(v)q(u)] + \int_u^n dv\, p(v)q(v). \tag{C.14}$$

154

Replica Limit

Taking the replica limit $n \to 0$ turns the previous expressions into the following:

$$\tilde{r} = \tilde{p}\tilde{q} - \langle qp \rangle, \tag{C.15}$$

$$r(u) = (\tilde{p} - \langle p \rangle)q(u) + (\tilde{q} - \langle q \rangle)p(u) - \int_0^u \mathrm{d}v(p(v) - p(u))(q(v) - q(u)), \tag{C.16}$$

with $\langle p \rangle = \int_0^1 \mathrm{d}u\, p(u)$.

Computing the Elements of the Inverse

To compute the matrix elements of the inverse matrix $\mathbf{Q}$ explicitly, we set $\mathbf{R} = \mathbf{PQ} = \mathbb{1}$, thereby obtaining $r(u) = 0$ and $\tilde{r} = 1$. We then have to solve the following two equations:

$$\tilde{r} = \tilde{p}\tilde{q} - \langle pq \rangle = 1, \tag{C.17}$$

$$r(u) = 0 = (\tilde{p} - \langle p \rangle)q(u) + (\tilde{q} - \langle q \rangle)p(u)$$
$$- \int_0^u \mathrm{d}v(p(v) - p(u))(q(v) - q(u)). \tag{C.18}$$

To solve these for $\tilde{q}$ and $q(u)$, we take the derivative of the second equation and define $[p](u) = \int_0^u \mathrm{d}v(p(u) - p(v))$ and $[p]'(u) = up'(u)$ such that:

$$0 = (\tilde{p} - [p](u) - \langle p \rangle)\, q'(u) + (\tilde{q} - [q](u) - \langle q \rangle)\, p'(u). \tag{C.19}$$

This can be rewritten as:

$$r'(u) = 0 = -\frac{1}{u}\frac{\mathrm{d}}{\mathrm{d}u}\left[(\tilde{p} - [p](u) - \langle p \rangle) \times (\tilde{q} - [q](u) - \langle q \rangle)\right]. \tag{C.20}$$

Or in other words, we see that:

$$[(\tilde{p} - [p](u) - \langle p \rangle)\, (\tilde{q} - [q](u) - \langle q \rangle)] = c. \tag{C.21}$$

where c is some constant which we can proceed to determine. With the knowledge that $r(u) = 0 \forall u$, we can compute $r(1)$ and note that the above expression appears in $r(1)$, leading us to the condition that $0 = 1 - c$, or $c = 1$. With this in mind, setting $u = 0$ in the above equation gives us the condition $(\tilde{p} - <p>)(\tilde{q} - <q>) = 1$, which we can use to rearrange the above equation for $q[u]$ to obtain:

$$[q](u) = -\frac{1}{\tilde{p} - \langle p \rangle}\frac{[p](u)}{\tilde{p} - \langle p \rangle - [p](u)}. \tag{C.22}$$

155

Now that we've found $[q](u)$, we can obtain $q(u)$ through differentiating the above, recalling that $[q]'(u) = uq'(u)$ and integrating appropriately:

$$q(u) = q(0) - \frac{1}{\tilde{p} - \langle p \rangle} \left[\frac{1}{u} \frac{[p](u)}{\tilde{p} - \langle p \rangle - [p](u)} + \int_0^u \frac{\mathrm{d}v}{v^2} \frac{[p](v)}{\tilde{p} - \langle p \rangle - [p](v)} \right].$$

(C.23)

We can determine $q(0)$ by setting $u = 0$ in Eq. (18), resulting in the final expression for $q(u)$:

$$q(u) = -\frac{1}{\tilde{p} - \langle p \rangle >} \left[\frac{p(0)}{\tilde{p} - \langle p \rangle} + \frac{1}{u} \frac{[p](u)}{\tilde{p} - \langle p \rangle - [p](u)} + \int_0^u \frac{\mathrm{d}v}{v^2} \frac{[p](v)}{\tilde{p} - \langle p \rangle - [p](v)} \right].$$

(C.24)

We can then set $u = 1$ in Eq. (18) to obtain the full expression for $\tilde{q}$:

$$\tilde{q} = \frac{1}{\tilde{p} - \langle p \rangle} \left(1 - \frac{p(0)}{\tilde{p} - \langle p \rangle} - \int_0^1 \frac{\mathrm{d}v}{v^2} \frac{[p](v)}{\tilde{p} - \langle p \rangle - [p](v)} \right).$$

(C.25)

Consistency Check: Replica-Symmetric Case

In the replica-symmetric case, all off diagonal elements take the same value, so we set all off diagonal elements to be p_0 and keep the diagonal elements as $\tilde{p}$. We then note the following simplifications:

$$p(0) = p_0,$$

(C.26)

$$\langle p \rangle = \int_0^1 \mathrm{d}v p(v) = p(0),$$

(C.27)

$$[p](u) = \int_0^u \mathrm{d}v \, (p(u) - p(v)) = 0.$$

(C.28)

Plugging these into the previous expressions, we obtain:

$$q(u) = -\frac{p_0}{(\tilde{p} - p_0)^2},$$

(C.29)

$$\tilde{q} = \frac{1}{\tilde{p} - p_0} - \frac{p_0}{(\tilde{p} - p_0)^2}.$$

(C.30)

These are exactly as we'd obtain by inverting the replica-symmetric matrix by hand.

Appendix D

Comment on Dotsenko et al (1995)

In Ref. [164], the authors conduct a renormalisation group study of the critical properties of a random-mass disordered p-component Heisenberg ferromagnet and showed that the RG flows of the system were unstable to replica symmetry breaking. Two different one-step RSB fixed points were found. In both cases the critical exponents depended on a parameter of the disorder distribution (denoted x_0) which remained undefined and led to non-universal critical exponents. The authors speculated that either x_0 could be determined to higher-order in their epsilon expansion, or that perhaps the disorder in their system led to non-universal exponents.

Here I show that x_0 can be uniquely determined from the existing equations in Ref. [164] and that once this value has been found, both of the one-step RSB fixed points the authors have found reduce to the same point which turns out to be unphysical.

Eq. 2.12 in Ref. [164] gives the fixed point equation as:

$$g(x) - 4g(x)^2 - (4 + 2p)g(x)\tilde{g} + 2pg(x)\int_0^1 \mathrm{d}y g(y) + p\int_0^x \mathrm{d}y\,[g(x) - g(y)]^2 = 0.$$

$$(\mathrm{D}.1)$$

The authors conclude, from "taking the derivative over x twice" that $g'(x) = 0$ and that therefore $g(x)$ is constant or step-like. For $g(x)$ to have a step-like structure, there has to be a value of x for which the gradient $g'(x)$ is not required to be zero, denoted x_0. This position went undetermined in Ref. [164].

Repeating their analysis and taking one derivative with respect to x yields:

$$g'(x)\left(1 - 8g(x) - (4 + 2p)\tilde{g} + 2p\int_0^1 \mathrm{d}yg(y) + 2xpg(x) - 2\int_0^x \mathrm{d}yg(y)\right) = 0,$$
(D.2)

for which $g'(x) = 0$ is one solution, as is $(...) = 0$. Taking this second bracketed term and performing a further derivative with respect to x gives:

$$g'(x)(2px - 8) = 0,$$ (D.3)

$$\to x_0 = \frac{4}{p}.$$ (D.4)

So the value of x_0 which determines the position of the step in the step function has to be $x_0 = 4/p$. This appears to be an oversight as there is no physical reason to dismiss this solution [165, 166].

Plugging this into their values for the step heights at the first fixed point, which the authors later find to be unphysical, gives the results:

$$g_0 = -\frac{4 - px_0}{16(px_0 - 1)} = 0,$$ (D.5)

$$g_1 = \tilde{g} = \frac{px_0}{16(px_0 - 1)} = \frac{1}{12}.$$ (D.6)

And for the second fixed point, considered by the authors to be the physical one:

$$g_0 \approx 0,$$ (D.7)

$$g_1 = -\frac{4 - p}{16(p - 1) - px_0(8 + p)} = \frac{1}{12},$$ (D.8)

$$\tilde{g} = \frac{p(1 - x_0)}{16(p - 1) - px_0(8 + p)} = \frac{1}{12}.$$ (D.9)

The values at both fixed points are entirely equal, therefore they really only have a single one-step RSB fixed point which does not depend on the undetermined x_0 and consequently leads to universal exponents.

However, there is now a further problem. These fixed points only have 1-step RSB solutions for $0 \leq x_0 \leq 1$, however this requires $p \geq 4$. The analysis of Ref. [164] shows that the clean system is in fact stable to disorder for $p > 4$, so the only disordered fixed point remaining with putative RSB is the one where $x_0 = 4/p = 1$.

The authors also state there are "several other one-step RSB solutions....which we do not reproduce here because they are always unstable," however checking this reveals a family of solutions, all of which are either unphysical or reduce back to the replica symmetric results. The case of $p = 4$ is somewhat anomalous as the $g(x) = 0$ everywhere except $x = 1$ which corresponds to the replica diagonal

entries of the Parisi matrix where it can never equal zero.

I make no further comment except to note that Ref. [250] in their study of quantum (anti-)ferromagnets find no evidence of RSB for $1 < p < 4$ and instead find their fixed points are replica symmetric and stable to RSB perturbations. The authors contrast their result with Ref. [164] and suggest this difference must stem from a physical difference between quantum and classical ferromagnetic systems, In light of the above analysis I instead suggest that stability to RSB perturbations is the correct result and that the critical behaviour of these systems is entirely replica symmetric.

Bibliography

[1] M. Reibold, P. Paufler, A. A. Levin, W. Kochmann, N. Pätzke, and D. C. Meyer, Nature **444**, 286 (2006).

[2] J. G. Bednorz and K. A. Müller, Zeitschrift für Physik B Condensed Matter **64**, 189 (1986).

[3] S. Chakravarty, R. B. Laughlin, D. K. Morr, and C. Nayak, Phys. Rev. B **63**, 094503 (2001).

[4] J. A. Mydosh and P. M. Oppeneer, Rev. Mod. Phys. **83**, 1301 (2011).

[5] H.-H. Kung, R. E. Baumbach, E. D. Bauer, V. K. Thorsmølle, W.-L. Zhang, K. Haule, J. A. Mydosh, and G. Blumberg, Science **347**, 1339 (2015).

[6] J. M. Kosterlitz and D. J. Thouless, Journal of Physics C: Solid State Physics **6**, 1181 (1973).

[7] G. Jaeger, Archive for History of Exact Sciences **53**, 51 (1998).

[8] S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar, Rev. Mod. Phys. **69**, 315 (1997).

[9] A. Altland and B. D. Simons, *Condensed Matter Field Theory* (Cambridge University Press, 2006).

[10] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, 1999).

[11] M. V. Berry, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **392**, 45 (1984).

[12] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).

[13] Clay Mathematics Institute Millenium Prize, `http://www.claymath.org/millennium-problems/p-vs-np-problem`, last accessed: 14/05/2016.

[14] A. J. Schofield, Contemporary Physics **40**, 95 (1999).

[15] C. Pfleiderer, S. R. Julian, and G. G. Lonzarich, Nature **414**, 427 (2001).

[16] Y. J. Uemura, T. Goko, I. M. Gat-Malureanu, J. P. Carlo, P. L. Russo, A. T. Savici, A. Aczel, G. J. MacDougall, J. A. Rodriguez, G. M. Luke, S. R. Dunsiger, A. McCollam, J. Arai, C. Pfleiderer, P. Böni, K. Yoshimura, E. Baggio-Saitovitch, M. B. Fontes, J. Larrea, Y. V. Sushko, and J. Sereni, Nat. Phys. **3**, 29 (2007).

[17] M. Otero-Leal, F. Rivadulla, M. García-Hernández, A. Piñeiro, V. Pardo, D. Baldomir, and J. Rivas, Phys. Rev. B **78**, 180415(R) (2008).

[18] V. Taufour, D. Aoki, G. Knebel, and J. Flouquet, Phys. Rev. Lett. **105**, 217201 (2010).

[19] E. A. Yelland, J. M. Barraclough, W. Wang, K. V. Kamenev, and A. D. Huxley, Nat. Phys. **7**, 890 (2011).

[20] A. Huxley, E. Ressouche, B. Grenier, D. Aoki, J. Flouquet, and C. Pfleiderer, Journal of Physics: Condensed Matter **15**, S1945 (2003).

[21] V. Taufour, D. Aoki, G. Knebel, and J. Flouquet, Phys. Rev. Lett. **105**, 217201 (2010).

[22] V. Taufour, A. Villaume, D. Aoki, G. Knebel, and J. Flouquet, Journal of Physics: Conference Series **273**, 012017 (2011).

[23] H. Kotegawa, V. Taufour, D. Aoki, G. Knebel, and J. Flouquet, Journal of the Physical Society of Japan **80**, 083703 (2011).

[24] S. A. Grigera, P. Gegenwart, R. A. Borzi, F. Weickert, A. J. Schofield, R. S. Perry, T. Tayama, T. Sakakibara, Y. Maeno, A. G. Green, and A. P. Mackenzie, Science **306**, 1154 (2004).

[25] R. A. Borzi, S. A. Grigera, J. Farrell, R. S. Perry, S. J. S. Lister, S. L. Lee, D. A. Tennant, Y. Maeno, and A. P. Mackenzie, Science **315**, 214 (2007).

[26] A. W. Rost, R. S. Perry, J.-F. Mercure, A. P. Mackenzie, and S. A. Grigera, Science **325**, 1360 (2009).

[27] A. W. Rost, S. A. Grigera, J. A. N. Bruin, R. S. Perry, D. Tian, S. Raghu, S. A. Kivelson, and A. P. Mackenzie, Proceedings of the National Academy of Sciences **108**, 16549 (2011).

[28] N. Doiron-Leyraud, I. R. Walker, L. Taillefer, M. J. Steiner, S. R. Julian, and G. G. Lonzarich, Nature **425**, 595 (2003).

[29] C. Pfleiderer, D. Reznik, L. Pintschovius, H. v. Löhneysen, M. Garst, and A. Rosch, Nature **427**, 227 (2004).

[30] G. Abdul-Jabbar, D. A. Sokolov, C. D. ONeill, C. Stock, D. Wermeille, F. Demmel, F. Krüger, A. G. Green, F. Lévy-Bertrand, B. Grenier, and A. Huxley, Nature Physics **11**, 321 (2015).

162

[31] J. D. Plumhof, T. Stöferle, L. Mai, U. Scherf, and R. F. Mahrt, Nature Materials **13**, 247 (2014).

[32] H. Kleinert, Phys. Rev. D **60**, 085001 (1999).

[33] M. Campostrini, M. Hasenbusch, A. Pelissetto, P. Rossi, and E. Vicari, Phys. Rev. B **63**, 214503 (2001).

[34] A. Pelissetto and E. Vicari, Physics Reports **368**, 549 (2002).

[35] M. Gell-Mann and F. E. Low, Phys. Rev. **95**, 1300 (1954).

[36] K. G. Wilson, Phys. Rev. B **4**, 3174 (1971).

[37] K. G. Wilson, Phys. Rev. B **4**, 3184 (1971).

[38] M. E. Fisher, Rev. Mod. Phys. **46**, 597 (1974).

[39] R. Shankar, Rev. Mod. Phys. **66**, 129 (1994).

[40] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics* (Cambridge University Press, 1995).

[41] P. Le Doussal and T. Giamarchi, Phys. Rev. Lett. **74**, 606 (1995).

[42] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).

[43] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005).

[44] S. Sachdev, Rev. Mod. Phys. **75**, 913 (2003).

[45] P. A. Lee, N. Nagaosa, and X.-G. Wen, Rev. Mod. Phys. **78**, 17 (2006).

[46] N. P. Armitage, P. Fournier, and R. L. Greene, Rev. Mod. Phys. **82**, 2421 (2010).

[47] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, Rev. Mod. Phys. **87**, 457 (2015).

[48] Y. Kamihara, H. Hiramatsu, M. Hirano, R. Kawamura, H. Yanagi, T. Kamiya, , and H. Hosono, Journal of the American Chemical Society **128**, 10012 (2006).

[49] Y. Kamihara, T. Watanabe, M. Hirano, , and H. Hosono, Journal of the American Chemical Society **130**, 3296 (2008).

[50] G. R. Stewart, Rev. Mod. Phys. **83**, 1589 (2011).

[51] P. Dai, Rev. Mod. Phys. **87**, 855 (2015).

[52] A. P. Mackenzie and Y. Maeno, Rev. Mod. Phys. **75**, 657 (2003).

[53] P. J. Curran, S. J. Bending, W. M. Desoky, A. S. Gibbs, S. L. Lee, and A. P. Mackenzie, Phys. Rev. B **89**, 144504 (2014).

[54] C. W. Hicks, D. O. Brodsky, E. A. Yelland, A. S. Gibbs, J. A. N. Bruin, M. E. Barber, S. D. Edkins, K. Nishimura, S. Yonezawa, Y. Maeno, and A. P. Mackenzie, Science **344**, 283 (2014).

[55] C. Meldgin, U. Ray, P. Russ, D. Chen, D. M. Ceperley, and B. DeMarco, Nature Physics **12** (2016), 10.1038/nphys3695.

[56] P. W. Anderson, Physics Today **41(1)**, 9 (1988).

[57] P. W. Anderson, Physics Today **41(3)**, 9 (1988).

[58] P. W. Anderson, Physics Today **41(6)**, 9 (1988).

[59] P. W. Anderson, Physics Today **41(9)**, 9 (1988).

[60] P. W. Anderson, Physics Today **42(7)**, 9 (1989).

[61] P. W. Anderson, Physics Today **42(9)**, 9 (1989).

[62] P. W. Anderson, Physics Today **43(3)**, 9 (1990).

[63] R. Nandkishore and D. A. Huse, Annual Review of Condensed Matter Physics **6**, 15 (2015).

[64] J. Eisert, M. Friesdorf, and C. Gogolin, Nature Physics **11**, 124 (2015).

[65] M. Mézard and G. Parisi, Journal of Statistical Physics **111**, 1 (2003).

[66] C. Chamon, A. W. W. Ludwig, and C. Nayak, Phys. Rev. B **60**, 2239 (1999).

[67] M. Mézard, G. Parisi, and M.-A. Virasoro, *Spin glass theory and beyond.* (World Scientific Publishing Co., Inc., Pergamon Press, 1990).

[68] K. Fischer and J. Hertz, *Spin Glasses*, Cambridge Studies in Magnetism (Cambridge University Press, 1991).

[69] V. Dotsenko, *Introduction to the Replica Theory of Disordered Statistical Systems*, Aléa-Saclay (Cambridge University Press, 2005).

[70] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).

[71] J. Zhao, S. L. Simon, and G. B. McKenna, Nature Communications **4**, 1783 (2013).

[72] L. Boltzmann, *Vorlesungen über Gastheorie*, Vol. 1 (JA Barth, 1896).

[73] R. Palmer, Advances in Physics **31**, 669 (1982).

[74] G. Parisi, arXiv preprint, cond-mat/0301157 (2003).

[75] G. Parisi, Phys. Rev. Lett. **50**, 1946 (1983).

[76] T. Vojta, Journal of Low Temperature Physics **161**, 299 (2010).

[77] T. Vojta and J. A. Hoyos, Phys. Rev. Lett. **112**, 075702 (2014).

[78] O. Motrunich, S.-C. Mau, D. A. Huse, and D. S. Fisher, Phys. Rev. B **61**, 1160 (2000).

[79] S. J. Thomson, F. Krüger, and A. G. Green, Phys. Rev. B **87**, 224203 (2013).

[80] J. A. Hertz, Phys. Rev. B **14**, 1165 (1976).

[81] T. Moriya, *Spin Fluctuations in Itinerant Electron Magnetism* (Springer Verlag, Berlin, 1985).

[82] A. J. Millis, Phys. Rev. B **48**, 7183 (1993).

[83] G. J. Conduit, A. G. Green, and B. D. Simons, Phys. Rev. Lett. **103**, 207201 (2009).

[84] S. Lausberg, J. Spehling, A. Steppke, A. Jesche, H. Luetkens, A. Amato, C. Baines, C. Krellner, M. Brando, C. Geibel, H.-H. Klauss, and F. Steglich, Phys. Rev. Lett. **109**, 216402 (2012).

[85] E. C. Stoner, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **165**, 372 (1938).

[86] E. C. Stoner, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **169**, 339 (1939).

[87] N. W. Ashcroft and N. D. Mermin, *Solid state physics* (Holt, Rinehart and Winston, 1976).

[88] R. L. Stratonovich, Soviet Physics Doklady **2**, 416 (1957).

[89] J. Hubbard, Phys. Rev. Lett. **3**, 77 (1959).

[90] H. v. Löhneysen, A. Rosch, M. Vojta, and P. Wölfle, Rev. Mod. Phys. **79**, 1015 (2007).

[91] S. P. Ridgway and C. A. Hooley, Phys. Rev. Lett. **114**, 226404 (2015).

[92] R. B. Laughlin, G. G. Lonzarich, P. Monthoux, and D. Pines, Advances in Physics **50**, 361 (2001).

[93] D. Belitz, T. R. Kirkpatrick, and T. Vojta, Rev. Mod. Phys. **77**, 579 (2005).

[94] D. Belitz, T. R. Kirkpatrick, and T. Vojta, Phys. Rev. Lett. **82**, 4707 (1999).

[95] D. Belitz and T. R. Kirkpatrick, Phys. Rev. Lett. **89**, 247202 (2002).

[96] A. V. Chubukov and D. L. Maslov, Phys. Rev. B **68**, 155113 (2003).

[97] A. V. Chubukov, C. Pépin, and J. Rech, Phys. Rev. Lett. **92**, 147003 (2004).

[98] A. V. Chubukov, D. L. Maslov, and A. J. Millis, Phys. Rev. B **73**, 045128 (2006).

[99] J. Rech, C. Pépin, and A. V. Chubukov, Phys. Rev. B **74**, 195126 (2006).

[100] D. V. Efremov, J. J. Betouras, and A. Chubukov, Phys. Rev. B **77**, 220401 (2008).

[101] D. L. Maslov and A. V. Chubukov, Phys. Rev. B **79**, 075112 (2009).

[102] A. V. Chubukov and D. L. Maslov, Phys. Rev. Lett. **103**, 216401 (2009).

[103] S. Coleman and E. Weinberg, Phys. Rev. D **7**, 1888 (1973).

[104] U. Karahasanovic, F. Krüger, and A. G. Green, Phys. Rev. B **85**, 165111 (2012).

[105] G. J. Conduit and B. D. Simons, Phys. Rev. A **79**, 053606 (2009).

[106] F. Krüger, U. Karahasanovic, and A. G. Green, Phys. Rev. Lett. **108**, 067003 (2012).

[107] C. J. Pedder, F. Krüger, and A. G. Green, Phys. Rev. B **88**, 165109 (2013).

[108] Y. Imry and S. Ma, Phys. Rev. Lett. **35**, 1399 (1975).

[109] R. A. Pelcovits, E. Pytte, and J. Rudnick, Phys. Rev. Lett. **40**, 476 (1978).

[110] M. Aizenman and J. Wehr, Phys. Rev. Lett. **62**, 2503 (1989).

[111] D. S. Fisher, Phys. Rev. B **31**, 7233 (1985).

[112] P. Le Doussal and K. J. Wiese, Phys. Rev. Lett. **96**, 197202 (2006).

[113] M. Tissier and G. Tarjus, Phys. Rev. Lett. **96**, 087202 (2006).

[114] V. S. Dotsenko and M. V. Feigelman, J. Phys. C **16**, L803 (1983).

[115] L. B. Ioffe and M. V. Feigel'man, Sov. Phys. JETP **61**, 354 (1985).

[116] M. V. Feigel'man and L. B. Ioffe, Phys. Rev. Lett. **58**, 1157 (1987).

[117] T. R. Kirkpatrick and D. Belitz, Phys. Rev. Lett. **104**, 256404 (2010).

[118] H. A. Gersch and G. C. Knollman, Phys. Rev. **129**, 959 (1963).

[119] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, Phys. Rev. B **40**, 546 (1989).

[120] S. J. Thomson and F. Krüger, EPL (Europhysics Letters) **108**, 30002 (2014).

166

[121] M. Ma, B. I. Halperin, and P. A. Lee, Phys. Rev. B **34**, 3136 (1986).

[122] T. Giamarchi and H. J. Schulz, Phys. Rev. B **37**, 325 (1988).

[123] D. G. Angelakis, M. F. Santos, and S. Bose, Phys. Rev. A **76**, 031805 (2007).

[124] D. Rossini and R. Fazio, Phys. Rev. Lett. **99**, 186401 (2007).

[125] P. Lugan, D. Clément, P. Bouyer, A. Aspect, M. Lewenstein, and L. Sanchez-Palencia, Phys. Rev. Lett. **98**, 170403 (2007).

[126] S. Morrison, A. Kantian, A. J. Daley, H. G. Katzgraber, M. Lewenstein, H. P. Büchler, and P. Zoller, New Journal of Physics **10**, 073032 (2008).

[127] U. Bissbort and W. Hofstetter, EPL (Europhysics Letters) **86**, 50007 (2009).

[128] J. K. Freericks and H. Monien, Phys. Rev. B **53**, 2691 (1996).

[129] B. V. Svistunov, Phys. Rev. B **54**, 16131 (1996).

[130] K. Sengupta and N. Dupuis, Phys. Rev. A **71**, 033629 (2005).

[131] R. Mukhopadhyay and P. B. Weichman, Phys. Rev. Lett. **76**, 2977 (1996).

[132] I. F. Herbut, Phys. Rev. Lett. **79**, 3502 (1997).

[133] I. F. Herbut, Phys. Rev. B **57**, 13729 (1998).

[134] P. B. Weichman and R. Mukhopadhyay, Phys. Rev. B **77**, 214516 (2008).

[135] F. Krüger, J. Wu, and P. Phillips, Phys. Rev. B **80**, 094526 (2009).

[136] F. Krüger, S. Hong, and P. Phillips, Phys. Rev. B **84**, 115118 (2011).

[137] A. Hegg, F. Krüger, and P. Phillips, Phys. Rev. B **88**, 134206 (2013).

[138] R. T. Scalettar, G. G. Batrouni, and G. T. Zimanyi, Phys. Rev. Lett. **66**, 3144 (1991).

[139] W. Krauth, N. Trivedi, and D. Ceperley, Phys. Rev. Lett. **67**, 2307 (1991).

[140] J. Kisker and H. Rieger, Phys. Rev. B **55**, R11981 (1997).

[141] P. Sen, N. Trivedi, and D. M. Ceperley, Phys. Rev. Lett. **86**, 4092 (2001).

[142] J.-W. Lee, M.-C. Cha, and D. Kim, Phys. Rev. Lett. **87**, 247006 (2001).

[143] N. Prokof'ev and B. Svistunov, Phys. Rev. Lett. **92**, 015703 (2004).

[144] R. V. Pai, R. Pandit, H. R. Krishnamurthy, and S. Ramasesha, Phys. Rev. Lett. **76**, 2937 (1996).

[145] J. Stasińska, P. Massignan, M. Bishop, J. Wehr, A. Sanpera, and M. Lewenstein, New Journal of Physics **14**, 043043 (2012).

[146] A. E. Niederle and H. Rieger, New Journal of Physics **15**, 075029 (2013).

[147] L. Pollet, N. V. Prokof'ev, B. V. Svistunov, and M. Troyer, Phys. Rev. Lett. **103**, 140402 (2009).

[148] V. Gurarie, L. Pollet, N. V. Prokof'ev, B. V. Svistunov, and M. Troyer, Phys. Rev. B **80**, 214519 (2009).

[149] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael, Phys. Rev. Lett. **93**, 150402 (2004).

[150] P. Sengupta and S. Haas, Phys. Rev. Lett. **99**, 050403 (2007).

[151] S. Iyer, D. Pekker, and G. Refael, Phys. Rev. B **85**, 094202 (2012).

[152] N. Ma, A. W. Sandvik, and D.-X. Yao, Phys. Rev. B **90**, 104425 (2014).

[153] Y. Wang, W. Guo, and A. W. Sandvik, Phys. Rev. Lett. **114**, 105303 (2015).

[154] S. J. Thomson and F. Krüger, Phys. Rev. B **92**, 180201 (2015).

[155] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, and A. Aspect, Nature **453**, 891 (2008).

[156] M. White, M. Pasienski, D. McKay, S. Q. Zhou, D. Ceperley, and B. DeMarco, Phys. Rev. Lett. **102**, 055301 (2009).

[157] M. Pasienski, D. McKay, M. White, and B. DeMarco, Nat. Phys. **6**, 677 (2010).

[158] S. S. Kondov, W. R. McGehee, J. J. Zirbel, and B. DeMarco, Science **334**, 66 (2011).

[159] M. Pasienski and B. DeMarco, Opt. Express **16**, 2176 (2008).

[160] G. D. Bruce, J. Mayoh, G. Smirne, L. Torralbo-Campo, and D. Cassettari, Phys. Scr. **T143**, 014008 (2011).

[161] A. L. Gaunt and Z. Hadzibabic, Scientific Reports **2**, 721 (2012).

[162] T. Giamarchi, P. Le Doussal, and E. Orignac, Phys. Rev. B **64**, 245119 (2001).

[163] V. Dotsenko and D. E. Feldman, Journal of Physics A: Mathematical and General **28**, 5183 (1995).

[164] V. Dotsenko, A. B. Harris, D. Sherrington, and R. B. Stinchcombe, Journal of Physics A: Mathematical and General **28**, 3093 (1995).

168

[165] X. tian Wu, Physica A: Statistical Mechanics and its Applications **251**, 309 (1998).

[166] L. De Cesare and M. T. Mercaldo, Phys. Rev. B **60**, 2976 (1999).

[167] D. Mouhanna and G. Tarjus, Phys. Rev. E **81**, 051101 (2010).

[168] S. G. Magalhes, C. V. Morais, and F. D. Nobre, Journal of Statistical Mechanics: Theory and Experiment **2011**, P07014 (2011).

[169] A. Perret, Z. Ristivojevic, P. Le Doussal, G. Schehr, and K. J. Wiese, Phys. Rev. Lett. **109**, 157205 (2012).

[170] V. S. Dotsenko, V. S. Dotsenko, and M. Picco, Nuclear Physics B **520**, 633 (1998).

[171] R. Yu, Y. Liang, N. S. Sullivan, J. S. Xia, C. Huan, A. Paduan-Filho, N. F. Oliveira, S. Haas, A. Steppke, C. F. Miclea, F. Weickert, R. Movshovich, E.-D. Mun, B. L. Scott, V. S. Zapf, and T. Roscilde, Nature **489**, 379 (2012).

[172] T. Nikuni, M. Oshikawa, A. Oosawa, and H. Tanaka, Phys. Rev. Lett. **84**, 5868 (2000).

[173] C. Rüegg, N. Cavadini, A. Furrer, H.-U. Gudel, K. Krämer, H. Mutka, A. Wildes, K. Habicht, and P. Vorderwisch, Nature **423**, 62 (2003).

[174] P. Merchant, B. Normand, K. W. Krämer, M. Boehm, D. F. McMorrow, and C. Rüegg, Nature Physics **10**, 373 (2014).

[175] R. Coldea, D. A. Tennant, K. Habicht, P. Smeibidl, C. Wolters, and Z. Tylczynski, Phys. Rev. Lett. **88**, 137203 (2002).

[176] T. Radu, H. Wilhelm, V. Yushankhai, D. Kovrizhin, R. Coldea, Z. Tylczynski, T. Lühmann, and F. Steglich, Phys. Rev. Lett. **95**, 127202 (2005).

[177] S. E. Sebastian, N. Harrison, C. D. Batista, L. Balicas, M. Jaime, P. A. Sharma, N. Kawashima, and I. R. Fisher, Nature **441**, 617 (2006).

[178] C. Rüegg, D. F. McMorrow, B. Normand, H. M. Rønnow, S. E. Sebastian, I. R. Fisher, C. D. Batista, S. N. Gvasaliya, C. Niedermayer, and J. Stahn, Phys. Rev. Lett. **98**, 017202 (2007).

[179] K. Kodama, M. Takigawa, M. Horvatić, C. Berthier, H. Kageyama, Y. Ueda, S. Miyahara, F. Becca, and F. Mila, Science **298**, 395 (2002).

[180] M. Jaime, R. Daou, S. A. Crooker, F. Weickert, A. Uchida, A. E. Feiguin, C. D. Batista, H. A. Dabkowska, and B. D. Gaulin, Proceedings of the National Academy of Sciences **109**, 12404 (2012).

[181] S. A. Zvyagin, J. Wosnitza, C. D. Batista, M. Tsukamoto, N. Kawashima, J. Krzystek, V. S. Zapf, M. Jaime, N. F. Oliveira, and A. Paduan-Filho, Phys. Rev. Lett. **98**, 047205 (2007).

[182] E. C. Samulon, K. A. Al-Hassanieh, Y.-J. Jo, M. C. Shapiro, L. Balicas, C. D. Batista, and I. R. Fisher, Phys. Rev. B **81**, 104421 (2010).

[183] S. Suh, K. A. Al-Hassanieh, E. C. Samulon, I. R. Fisher, S. E. Brown, and C. D. Batista, Phys. Rev. B **84**, 054413 (2011).

[184] E. C. Samulon, M. C. Shapiro, and I. R. Fisher, Phys. Rev. B **84**, 054417 (2011).

[185] T. Matsubara and H. Matsuda, Progress of Theoretical Physics **16**, 569 (1956).

[186] T. Holstein and H. Primakoff, Phys. Rev. **58**, 1098 (1940).

[187] S. Sachdev and R. N. Bhatt, Phys. Rev. B **41**, 9323 (1990).

[188] M. Vojta and T. Ulbricht, Phys. Rev. Lett. **93**, 127002 (2004).

[189] R. L. Doretto and M. Vojta, Phys. Rev. B **85**, 104416 (2012).

[190] D. G. Joshi, K. Coester, K. P. Schmidt, and M. Vojta, Phys. Rev. B **91**, 094404 (2015).

[191] D. G. Joshi and M. Vojta, Phys. Rev. B **91**, 094405 (2015).

[192] T. Giamarchi, C. Rüegg, and O. Tchernyshynov, Nature Physics **4**, 198 (2008).

[193] V. Zapf, M. Jaime, and C. D. Batista, Rev. Mod. Phys. **86**, 563 (2014).

[194] T. Hong, A. Zheludev, H. Manaka, and L.-P. Regnault, Phys. Rev. B **81**, 060410 (2010).

[195] M. B. Stone, A. Podlesnyak, G. Ehlers, A. Huq, E. C. Samulon, M. C. Shapiro, and I. R. Fisher, Journal of Physics: Condensed Matter **23**, 416003 (2011).

[196] F. Yamada, H. Tanaka, T. Ono, and H. Nojiri, Phys. Rev. B **83**, 020409 (2011).

[197] E. Wulf, S. Mühlbauer, T. Yankova, and A. Zheludev, Phys. Rev. B **84**, 174414 (2011).

[198] E. Orignac, T. Giamarchi, and P. Le Doussal, Phys. Rev. Lett. **83**, 2378 (1999).

[199] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael, Phys. Rev. Lett. **100**, 170402 (2008).

[200] E. Altman, Y. Kafri, A. Polkovnikov, and G. Refael, Phys. Rev. B **81**, 174528 (2010).

[201] T. Roscilde and S. Haas, Phys. Rev. Lett. **99**, 047205 (2007).

[202] P. J. H. Denteneer, R. T. Scalettar, and N. Trivedi, Phys. Rev. Lett. **87**, 146401 (2001).

[203] P. B. Weichman and R. Mukhopadhyay, Phys. Rev. B **77**, 214516 (2008).

[204] A. van Otterlo, K.-H. Wagenblast, R. Baltin, C. Bruder, R. Fazio, and G. Schön, Phys. Rev. B **52**, 16176 (1995).

[205] V. W. Scarola and S. Das Sarma, Phys. Rev. Lett. **95**, 033003 (2005).

[206] V. W. Scarola, E. Demler, and S. Das Sarma, Phys. Rev. A **73**, 051601 (2006).

[207] K. Yamamoto, S. Todo, and S. Miyashita, Phys. Rev. B **79**, 094503 (2009).

[208] M. Iskin and J. K. Freericks, Phys. Rev. A **79**, 053634 (2009).

[209] M. Iskin, Phys. Rev. A **83**, 051606 (2011).

[210] D. Yamamoto, A. Masaki, and I. Danshita, Phys. Rev. B **86**, 054516 (2012).

[211] D. Yamamoto, T. Ozaki, C. A. R. Sá de Melo, and I. Danshita, Phys. Rev. A **88**, 033624 (2013).

[212] G. G. Batrouni and R. T. Scalettar, Phys. Rev. Lett. **84**, 1599 (2000).

[213] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Phys. Rev. Lett. **81**, 3108 (1998).

[214] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Nature **415**, 39 (2002).

[215] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, Nature **462**, 74 (2009).

[216] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner, Science **329**, 547 (2010).

[217] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, Nature **467**, 68 (2010).

[218] S. J. Thomson, L. S. Walker, T. L. Harte, and G. D. Bruce, arXiv preprint, arXiv:1607.05254 (2016).

[219] L. W. Cheuk, M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein, Phys. Rev. Lett. **114**, 193001 (2015).

[220] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner, Phys. Rev. Lett. **114**, 213002 (2015).

[221] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr, Nature Physics **1**, 738 (2015).

[222] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen, Phys. Rev. A **92**, 063406 (2015).

[223] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross, Phys. Rev. Lett. **115**, 263001 (2015).

[224] L. Fallani, J. E. Lye, V. Guarrera, C. Fort, and M. Inguscio, Phys. Rev. Lett. **98**, 130404 (2007).

[225] B. Gadway, D. Pertot, J. Reeves, M. Vogt, and D. Schneble, Phys. Rev. Lett. **107**, 145306 (2011).

[226] C. D'Errico, E. Lucioni, L. Tanzi, L. Gori, G. Roux, I. P. McCulloch, T. Giamarchi, M. Inguscio, and G. Modugno, Phys. Rev. Lett. **113**, 095301 (2014).

[227] K. D. Nelson, X. Li, and D. S. Weiss, Nature Physics **3**, 556 (2007).

[228] G. D. Bruce, M. Y. H. Johnson, E. Cormack, D. A. W. Richards, J. Mayoh, and D. Cassettari, Journal of Physics B: Atomic, Molecular and Optical Physics **48**, 115303 (2015).

[229] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, Advances in Atomic Molecular and Optical Physics **42**, 95 (2000).

[230] I. Bloch, Nature Physics **1**, 23 (2005).

[231] H. Ott, Reports on Progress in Physics **79**, 054401 (2016).

[232] P. Cheinet, S. Trotzky, M. Feld, U. Schnorrberger, M. Moreno-Cardoner, S. Fölling, and I. Bloch, Phys. Rev. Lett. **101**, 090404 (2008).

[233] P. M. Preiss, R. Ma, M. E. Tai, J. Simon, and M. Greiner, Phys. Rev. A **91**, 041602 (2015).

[234] T. Roscilde, Phys. Rev. A **82**, 023601 (2010).

[235] P. Buonsante, V. Penna, A. Vezzani, and P. B. Blakie, Phys. Rev. A **76**, 011602 (2007).

[236] M. C. Gutzwiller, Phys. Rev. Lett. **10**, 159 (1963).

[237] M. R. Hestenes and E. Stiefel, Journal of Research of the National Bureau of Statistics **49** (1952).

[238] L. Walker, MPhys Thesis, University of St Andrews (2016).

[239] R. Graham and A. Pelster, Int. J. Bifurcation Chaos **19**, 2745 (2009).

[240] T. Khellil and A. Pelster, J. Stat. Mech. **2016**, 063301 (2016).

[241] N. Ghofraniha, I. Viola, F. Di Maria, G. Barbarella, G. Gigli, L. Leuzzi, and C. Conti, Nat. Comm. **6**, 6058 (2015).

[242] P. Buonsante and A. Vezzani, Phys. Rev. A **70**, 033608 (2004).

[243] F. Gerbier, Phys. Rev. Lett. **99**, 120405 (2007).

[244] D. Basko, I. Aleiner, and B. Altshuler, Annals of Physics **321**, 1126 (2006).

[245] F. Krüger, C. J. Pedder, and A. G. Green, Phys. Rev. Lett. **113**, 147001 (2014).

[246] G. Hannappel, C. J. Pedder, F. Krüger, and A. G. Green, Phys. Rev. B **93**, 235105 (2016).

[247] M. Brando, D. Belitz, F. M. Grosche, and T. R. Kirkpatrick, Rev. Mod. Phys. **88**, 025006 (2016).

[248] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, Science **352**, 1547 (2016).

[249] M. A. Metlitski, C. A. Fuertes, and S. Sachdev, Phys. Rev. B **80**, 115122 (2009).

[250] R. Narayanan, T. Vojta, D. Belitz, and T. R. Kirkpatrick, Phys. Rev. B **60**, 10150 (1999).