Dissertation

submitted to the

Combined Faculty of Mathematics, Engineering and Natural Sciences
of the Ruperto-Carola-University of Heidelberg, Germany

for the degree of

Doctor of Natural Sciences

Put forward by

Marc Bauer

born in: Freudenstadt, Germany
Oral examination: November 20th, 2024

Stochastic Approaches and the Physics of Ultracold Fermions: From Lattice Simulations to Machine Learning

Referees: Prof. Dr. Jan M. Pawlowski
Prof. Dr. Maurits Haverkort

**Stochastic Approaches and the Physics of Ultracold Fermions:
From Lattice Simulations to Machine Learning**

This thesis investigates fermionic ultracold atom systems in reduced dimensions using stochastic and lattice-based methods. One particular focus of the current work lies in exploring the role trapping potentials play in these systems. Including traps in lattice simulations poses a challenge, as they break translational symmetry, and is made possible via efficient sampling techniques. Another aspect we examine is the effect of population imbalances, which lead to a sign problem in Monte Carlo simulations. To mitigate this issue, we employ both complex Langevin and reweighting methods, analyzing their effectiveness and applicability. Notably, our approach to trapped systems matches experimental and theoretical benchmark results in one dimension perfectly but allows for significantly larger particle numbers and imbalances, for which we find signs of FFLO-type pairing. In two-dimensional untrapped systems, we encounter only mild sign problems and explore the normal phase of the BEC-BCS crossover regime. We offer falsifiable predictions for thermodynamic quantities. In addition to the results on fermionic systems, we introduce a novel normalizing flow architecture for upscaling field configurations. This architecture shows promise for reducing computational complexity and tackling the problem of critical slowing down in lattice simulations of all kinds.

**Stochastische Methoden und die Physik ultrakalter Fermionen:
Von Gitter-Simulationen zu maschinellem Lernen**

Diese Arbeit untersucht fermionische ultrakalte Atomsysteme in reduzierten Dimensionen unter Verwendung stochastischer und gitterbasierter Methoden. Ein besonderer Fokus liegt dabei auf der Erforschung der Rolle von Fallenpotenzialen in derartigen Systemen. Die Einbeziehung von Fallen in Gitter-Simulationen stellt eine Herausforderung dar, da diese die Translationssymmetrie brechen, und wird durch den Einsatz effizienter Samplingalgorithmen ermöglicht. Ein weiterer Aspekt, den wir untersuchen, ist der Einfluss von Populationsungleichgewichten, die in Monte-Carlo-Simulationen zu einem Vorzeichenproblem führen. Um dieses Problem zu mildern, verwenden wir sowohl die komplexe Langevin- als auch die Umgewichtungs-Methode und analysieren deren Effektivität und Anwendbarkeit. Bemerkenswert ist, dass unser Ansatz für Systeme mit Fallen in einer Dimension die experimentellen und theoretischen Vergleichsergebnisse perfekt widerspiegelt, jedoch wesentlich größere Teilchenzahlen und Ungleichgewichte zulässt, bei denen wir Anzeichen einer FFLO-artigen Paarung finden. In zweidimensionalen Systemen ohne Fallen stoßen wir auf nur leichte Vorzeichenprobleme und erforschen die normale Phase des BEC-BCS-Crossover-Regimes. Dabei machen wir überprüfbare Vorhersagen für thermodynamische Größen. Neben den Ergebnissen zu fermionischen Systemen stellen wir eine neuartige Architektur für normalisierte Flüsse vor, die zur Skalierung von Feldkonfigurationen dient. Diese Architektur ist vielversprechend, um Rechenkomplexität sowie critical slowing down in Gitter-Simulationen zu reduzieren.

Publications

This thesis was written and compiled by the author, drawing on various publications and conference presentations, to which several collaborators contributed. The author was involved in the research for:

- *Harmonically trapped fermions in one dimension: A finite-temperature lattice Monte Carlo study* [1]
  Collaborators: Felipe Attanasio, Renzo Kapust, Jan M. Pawlowski
  Published in *Phys.Rev.A 109 (2024) 3, 033305*.

- *Searching for Yang-Lee zeros in O(N) models* [2]
  Collaborators: Felipe Attanasio, Lukas Kades, Jan M. Pawlowski
  Contribution to *Proceedings of science, Lattice 2022*.

- *Density profiles and correlations of harmonically trapped ultracold fermions via complex Langevin* [3]
  Collaborators: Felipe Attanasio, Jan M. Pawlowski
  Contribution to *Proceedings of science, Lattice 2023*.

Additionally, this thesis contains the following unpublished work:

- *Population imbalanced Fermi gases in the 2D BEC-BCS crossover* [4]
  Collaborators: Felipe Attanasio, Jan M. Pawlowski

- *Normalizing flows for inverse renormalisation group steps* [5]
  Collaborators: Renzo Kapust, Jan M. Pawlowski, Finn L. Temmen

Moreover, the author contributed to the following preprint, which was not included in the compilation of this thesis:

- *Speeding up Fermionic Lattice Calculations with Photonic Accelerated Inverters* [6]
  Collaborators: Felipe Attanasio, Jelle Dijkstra, Timoteo Lee, Jan M. Pawlowski, Wolfram Pernice.
  e-print available 2401.14200 [hep-lat] (2024).

On the author's contributions

In [1], the author performed the majority of the simulations and analysis, as well as the writing of the manuscript. The simulation code and general framework were written by the author in close collaboration with Felipe Attanasio. Ground state simulations, as well as the required modifications to the code, were conducted by Renzo Kapust during a project practical co-supervised by the author.

The author performed the computations for the one-site model in [2]. The author also contributed to general discussions regarding the project. Felipe Attanasio conducted the experiments in three dimensions and wrote the manuscript.

All computations in [3, 4], as well as large parts of the manuscripts, are the work of the author.

The unpublished work [5] partly arose from the Masters project of Renzo Kapust. It builds on a codebase developed in collaboration with Finn Temmen during his own Masters project, which the author co-supervised. The author was responsible for obtaining all simulation results presented.

Contents

Chapter 1

Introduction

The study of many-body systems has long been a central theme in physics and other natural sciences, appearing in diverse areas such as complex materials in condensed matter physics, the collective behavior of cells in biophysics, and the intricate phase structures of quantum chromodynamics (QCD). Despite the seemingly disparate nature of these systems, common features emerge when the number of constituents becomes large. One of the most striking examples is, without a doubt, the emergence of phase transitions, in which small variations in a controllable parameter lead to sudden and qualitative changes in behavior. A multitude of physical phenomena vital to the existence of human life on Earth are based on such transitions, most notably the freezing and evaporation of water and the inner workings of basic rice cookers[1].

Describing many-body systems theoretically is a formidable task, as analytic solutions are generally not attainable when dealing with hundreds, thousands or even more particles. Instead, idealizations and approximations are made in order to simplify the problem. For systems in thermal equilibrium, statistical physics offers one such approach in the form of collective descriptions. Single particle degrees of freedom are replaced by collective ones, such as the temperature, pressure, and density, leading to a substantial simplification.

Another idealization comes in the form of low-energy effective theories, which aim to capture only the essential physics. Instead of an entirely microscopic description of interactions, effective theories offer formulations in terms of the most relevant degrees of freedom. Prominent examples include the introduction of phonons for the description of lattice vibrations in solids or the use of mesons and baryons in models for QCD. Another staggering success of effective theories is the explanation of superconductivity via Bardeen-Cooper-Schrieffer (BCS) theory, which models the pairing of electrons and predicts the emergence of a gap in the excitation spectrum of certain metals at low temperatures.

While effective theories and collective descriptions greatly reduce the complexity of many-body systems, gaining physical insight can still be challenging. Mean-field approaches allow for analytical access and can provide good qualitative results but are often insufficient in the strongly correlated regimes of interest. For this reason, numerical methods have become an indispensable tool

[1] Simple rice cookers work by virtue of the transition between ferromagnetic and paramagnetic phases in an induced magnet and the liquid-gas transition in water. Loosely speaking, a permanent magnet and a ferromagnetic alloy keep the activated power switch in place at temperatures below 100°C. Once the water has evaporated and temperatures rise above 100°C, the alloy demagnetizes and the switch is released, shutting down the appliance.

in the study of many-body problems. Prime among them are Markov chain Monte Carlo methods, which reduce the problem of computing high-dimensional integrals to a problem of sampling a non-trivial probability distribution. Crucially, stochastic methods can be applied even in strongly interacting scenarios and allow for non-perturbative studies of systems inaccessible by conventional approaches. Today, they provide gold-standard results for many problems in condensed matter physics, nuclear physics, and quantum field theory. In the latter, lattice gauge theory simulations, which discretize spacetime into a finite grid, have been instrumental in the study of the strong force. Amongst other numerical methods, they offer the distinct advantages of a well-defined strategy to remove systematic errors and exhibit benign scaling with larger particle contents. This is in contrast to the exponential growth of the Hilbert space, which makes full diagonalization intractable for more than a few particles.

Unfortunately, several pathologies can plague Monte Carlo simulations in physically interesting regimes. In particular, the infamous sign problem, arising when one cannot identify a real and positive probability distribution, is prevalent in systems with fermions [7]. Another challenge comes in the form of critical slowing down, which occurs when the correlation length grows large in the continuum limit. In both of these examples, the computational cost of generating relevant configurations grows rapidly, often rendering the stochastic approach infeasible.

Experimentally, probing the physics of many body systems is similarly challenging. In particular, particle physics experiments require an enormous amount of effort and infrastructure in the form of particle accelerators and detectors. In contrast, the advent of cold atom experiments in the last three decades has led to a revolution in the study of many-body effects. The field was kickstarted by the realization of a Bose-Einstein condensate in 1995 [8–10], and has since been extended to fermions, mixtures and even artificial gauge fields [11–15]. Typically, these experiments use alkali atoms such as ^{86}Rb or ^{23}Na for bosons and ^{7}Li and ^{40}K for fermions, as well various other alkali, alkaline earth and rare earth atoms. Using a robust hierarchy of scales, which is made possible through laser cooling and trapping techniques, cold atom experiments allow for precise control of the interactions between particles. Hamiltonians can be engineered almost at will, and the systems can be probed with high precision. Amazingly, the use of suitable trapping potentials even opens up the possibility of confining the atoms to one or two spatial dimensions to test the effects of dimensionality. In total, cold atom experiments offer a unique opportunity to study many-body physics and can serve as simulators for other systems, as they are not restricted by the couplings provided in nature.

This meteoric rise of cold atoms experiments has also led to a gold rush in the theoretical community, both to supply falsifiable predictions and explain experimental results. The current thesis can be seen as part of this effort and focuses in large parts on the study of fermionic cold atoms systems in reduced dimensions. In particular, we investigate the effects of trapping potentials in one dimension and how they can be included in lattice Monte Carlo simulations. Additionally, we study population imbalanced Fermi gases, where non-degenerate Fermi surfaces allow for exotic pairing. Since fermionic systems can suffer both from sign problems and critical slowing down, we also explore existing and novel approaches to alleviate these issues in the form of complex Langevin simulations and normalizing flow-based machine learning architectures.

1.1 Basics of cold atom physics

This section gives a brief overview of the basic concepts behind cold atom physics. We begin by exploring the principles of trapping and cooling, along with the different scales required to create cold atomic gases, enabling the use of low-energy effective theories. Next, we introduce the concept of Feshbach resonances, which allow precise control over atomic interactions. This, in turn, leads us to the BEC-BCS crossover, which describes the transition between a Bose-Einstein condensate (BEC) and a Bardeen-Cooper-Schrieffer (BCS) superfluid in Fermi gases. We conclude by reviewing Fermi gases with imbalanced populations, where unconventional pairing may emerge.

1.1.1 Trapping and cooling atoms

The first step in creating a cold atomic gas is to trap the atoms within a confined region of space. This is typically achieved by using a combination of magnetic and optical traps [16], which can be used to create arbitrary trapping geometries in modern experiments. Due to their ease of implementation, harmonic potentials, defined by

$$V(\mathbf{r}) = \frac{m}{2}\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right), \tag{1.1}$$

are most commonly used. Here, m is the mass of the atom, and ω_i represents the trapping frequency in the i-th direction. Crucially, these frequencies can be tuned independently, allowing for anisotropic configurations. If the frequency in one direction is made grater than all other scales, excitations in that direction become energetically unfavorable. This leads to quasi-two-dimensional or even quasi-one-dimensional behavior.

The ubiquity of harmonic trapping potentials presents challenges when comparing experimental results with theoretical models. In the latter, the assumption of homogeneity is often used, which simplifies computations but is violated by a harmonic potential. This issue also affects typical lattice simulations, where translational invariance is essential for enhancing statistical accuracy. A major focus of this thesis lies in the inclusion of trapping potentials in lattice simulations and how they can be implemented efficiently to allow for direct comparison with experiments.

With the atoms trapped, the next step is to cool them to temperatures where quantum effects become relevant. This is achieved by a combination of laser cooling [17] and evaporative cooling [18]. The former is based on the principle of Doppler cooling, where the absorption and reemission of photons slows down atoms. Evaporative cooling works by modulating the trap so that the hottest atoms can spill out, thereby cooling the remaining gas. These methods can achieve temperatures as low as a few nanokelvins, enabling studies down to the ground state.

1.1.2 An interplay of scales

We now discuss the various length scales present when realizing cold atomic gases and how control over them is instrumental in the study of many-body physics.

The first important length scale is the inter-particle spacing l_n, which is directly related to the density n of the gas as

$$n = l_n^{-d}, \tag{1.2}$$

where d is the spatial dimension. This relationship assumes that the gas is homogeneous, a condition met in box-shaped trapping geometries but not in harmonic traps. In the latter case, the density and inter-particle distance become dependent on spatial position.

Secondly, there is the length scale associated with the trapping potential itself. For an isotropic harmonic trap with frequency ω, the characteristic length scale is the oscillator length given by

$$l_{\mathrm{Trap}} = \sqrt{\frac{\hbar}{m\omega}}, \tag{1.3}$$

where $\hbar$ is the reduced Planck constant. In a box potential, this scale is simply the linear size of the box. The trap length should be the largest scale present when studying many-body systems in the thermodynamic limit. This ensures minimal boundary effects and allows for the use of a local density approximation (LDA) for harmonic potentials.

Next, we have the thermal (or de Broglie) wavelength λ_{T}, given as

$$\lambda_{\mathrm{T}} = \left(\frac{2\pi\hbar^2}{m k_{\mathrm{B}} T}\right)^{1/2}, \tag{1.4}$$

where k_{B} is the Boltzmann constant, and T is the temperature. The thermal wavelength is related to the kinetic energy of particles and can be interpreted as a measure of the size of their wave packets. Together with the inter-particle spacing, the thermal wavelength determines the degeneracy of the gas. In the regime where $l_n/\lambda_T > 1$, wavepackets of different particles do not overlap significantly, and quantum statistics play no role. However, when $l_n/\lambda_T \simeq 1$, the system transitions from a classical to a quantum degenerate regime, making the gas *ultracold*.

The final initial length scales is associated with the interaction between single alkali atoms. The interaction potential between alkali atoms can be approximated using a Lennard-Jones potential, which is effectively modeled as $V(r) = \infty$ for a core $r < r_0$, with tails $V(r) = Cr^{-6}$ for $r > r_0$. Here, C is a constant. The length scale associated with the interaction is then given by the van der Waals length

$$l_{\mathrm{vdW}} = \left(\frac{mC}{\hbar^2}\right)^{1/4}. \tag{1.5}$$

This constitutes the smallest relevant length scale, which allows interactions to be described as point-like in an effective description.

A typical hierarchy of these length scales for studying the quantum degenerate but finite-temperature properties of ultracold gases is thus given by

$$l_{\mathrm{vdW}} \ll l_n \ll \lambda_{\mathrm{T}} \ll l_{\mathrm{Trap}}. \tag{1.6}$$

While the van der Waals length must always remain the smallest, the order of the others may vary without invalidating the effective description. For example, when examining the ground state properties in a harmonic trap, the thermal wavelength should be much larger than the harmonic oscillator length to prevent the population of excited states. Conversely, when considering the equation of state of the gas, one may adjust the system from a classical to a quantum regime, flipping the order of the inter-particle spacing and the thermal wavelength in the process.

Since scattering processes in cold atom experiments happen exclusively at very low energies, only the s-wave channel is relevant. This means that the associated scale, the s-wave scattering

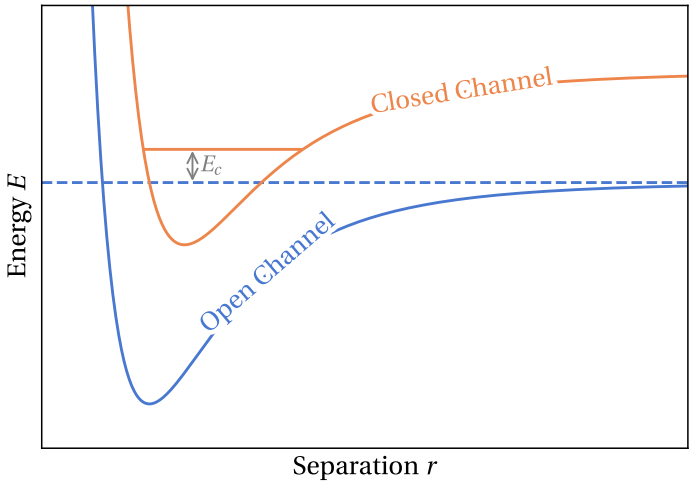

Figure 1.1: Schematic depiction of the inter-atomic potential as a function of the distance r between atoms. The straight solid line indicates the bound state permitted by the closed channel, while the dashed line indicates the asymptotic state in the open channel.

length a_s, fully characterizes the interactions between atoms. In contrast to the van der Waals length, the scattering length is not required to be small and may even be tuned to become much larger than the other scales, leading to the regime of unitarity in three dimensions. This is achieved by exploiting Feshbach resonances, which we discuss in the next section.

1.1.3 Feschbach resonances

One of the main attractions of experiments with cold atoms is the unprecedented control over particle interactions. This control is primarily achieved by exploiting the hyperfine structure of alkali atoms, which enables the tuning of the scattering length a_s via Feshbach resonances.

The concept can be illustrated using a simple two-channel model of two-body scattering. The first channel asymptotically connects two free atoms in the gas and is called *open*. On the other hand, the *closed* channel is not available as an asymptotic state. Crucially, the closed channel must support a bound state with energy E_c in the vicinity of the open channel threshold. A schematic depiction of the inter-atomic potential as a function of the distance r between atoms is shown in Figure 1.1.

When the energy of a pair of colliding atoms is close to the bound state energy E_c, they may briefly form a virtual bound state before scattering elastically. This resonance can be tuned by, for example, applying an external magnetic field, leading to a so-called *magnetic* Feshbach resonance. The scattering length a_s is then related to the external magnetic field B by the formula

$$a_s(B) = a_{\text{bg}}\left(1 - \frac{\Delta B}{B - B_0}\right), \tag{1.7}$$

where a_{bg} is the background scattering length of the open channel, B_0 is the field at which the resonance occurs and ΔB is the width of the resonance. These parameters vary depending on the specific atomic species used in the experiment.

Naturally, the situation in actual experiments is more complicated than the simple two-channel model presented here, but the basic idea remains the same. A much more detailed discussion of Feschbach resonances in cold atoms physics can be found in [19].

1.1.4 The BEC-BCS crossover

We now consider a Fermi gas of atoms with two internal states, which couple to each other through a tunable contact interaction. Even a weak attractive interaction makes it energetically favorable for fermions to form pairs. These pairs are situated around the Fermi surface, involving atoms with opposite momenta and spins, and are not localized in position space. This paired state is stable against fluctuations at low temperatures and leads to a gap in the excitation spectrum, allowing for superfluidity. This is the essence of Bardeen-Cooper-Schrieffer (BCS) theory.

The situation changes drastically when looking at strong, attractive interactions between the atom species. In that case, the pairs become tightly bound and effectively form a gas of composite bosons. The strong interactions between fermions are screened, and the system is well described by weakly interacting composites. When the temperature becomes sufficiently low, the gas undergoes a transition to a Bose-Einstein condensate (BEC) with an over-occupied zero momentum state.

One of the greatest successes of cold atoms experiments is undoubtedly the realization of the BEC-BCS crossover (see [20] and [21] for reviews). One finds a smooth transition between the two aforementioned regimes by tuning the interactions between particles. Notably, this is true in both three and two dimensions, even though the mechanisms behind the phase transitions are different. While both the BEC and BCS limits can be well described by mean-field theory, the crossover regime is characterized by strong correlations and is inherently non-perturbative, calling for advanced (numerical) methods on the theoretical side.

1.1.5 Fermi gases with imbalanced populations

So far, we have considered Fermi gases with equal populations so that each particle may find a partner to form a pair. A natural extension of the discussion above presents itself in the form of population imbalance in the gas. An imbalance may be understood as the effect of an external magnetic field that energetically favors one of the internal states. Note that this is not to be confused with the magnetic field used to tune the Feschbach resonance. Consequently, the Fermi surfaces for each spin species are no longer degenerate, disrupting conventional BCS-type pairing.

This situation has a direct influence on the occurrence of superfluidity in the gas, and we can straightforwardly identify two regimes, separated by the Chandrasekhar-Clogston limit [22, 23]. They are characterized by the competition of the pairing energy against the energy gain from aligning with the magnetic field in the normal phase. In the regime where the pairing energy dominates, the system remains in a spin-balanced superfluidity phase, and BCS theory is still applicable. Once the energy gain from aligning the spins becomes dominant, superfluid behavior is destroyed, and the system enters a spin-imbalanced normal phase. When the magnetic field is increased even further, the system eventually becomes fully polarized, and the gas turns non-interacting.

However, the situation may be more complex than the picture above. As both Fulde and Ferrel [24] as well as Larkin and Ovchinnikov [25] have pointed out, superfluidity could still be possible in the form of an intermediate phase with a spatially modulated order parameter. While the de-

tails between the proposals differ, they both predict pairing at a finite center of mass momentum, starkly contrasting conventional BCS theory.

Experimental studies of imbalanced Fermi gases have been performed most notably in three-dimensional trapped systems [26–29], where spatial phase separation between regions of an (almost) balanced superfluid and a polarized normal phase has been observed. Similar effects have also been found in two-dimensional systems [30, 31]. However, concrete evidence for the emergence of a Fulde-Ferrel-Larkin-Ovchinnikov (FFLO) phase has been elusive, likely due to the small parameter regions where it is predicted to exist.

Studying imbalanced Fermi gases poses a considerable challenge in lattice simulations. While the attractive, balanced gas allows for sign problem free formulations, no such simplifications are known in the imbalanced case. For Fermi gases in particular, the literature offers a small set of works using the complex Langevin method in three spatial dimensions [32–34], but a comprehensive study of the phase diagram has not been possible so far. This is also true for reduced dimensions, where some work in the context of the (dilute) Hubbard model exists [35], but a complete characterization of the sign problem severity is still missing.

1.2 Outline

This thesis is structured as follows:

Chapter 2 briefly introduces the general concepts of lattice field theory and stochastic sampling. We focus on non-relativistic fermions in particular, which come with some unique challenges and opportunities compared to their relativistic counterparts.

In Chapter 3, we introduce complex Langevin evolution, which allows for sampling in the presence of a sign problem and constitutes one of the main methods used in this work. We also discuss several shortcomings of the approach and explore its applicability to the study of zeros in the partition function, also known as Lee-Yang zeros.

As a final building block for successful lattice simulations, we discuss further technical details of the Determinant Quantum Monte Carlo (DQMC) approach for fermions in Chapter 4. Notably, we use a recently developed method that reduces the model space and speeds up computations considerably. This reduction is crucial when studying trapped systems, as the loss of translational invariance significantly increases the amount of samples required to obtain reliable results.

In Chapter 5, we apply the methods introduced in the previous chapters to study trapped fermions in one dimension. We show that the trapping potentials can be included straightforwardly and discuss the effects of the trap on the system. Moreover, we compare to results from experiments and other theoretical approaches.

Chapter 6 concerns the (untrapped) Fermi gas in two dimensions and its equation of state with population imbalance. The sign problem turns out to be fairly mild for a wide range of parameters, and we find excellent agreement with results obtained from a virial expansion.

Finally, in Chapter 7, we introduce a novel approach to stochastic sampling based on upscaling field configurations combined with generative machine learning models. Our method can be understood as an inversion of renormalization group transformation and has the potential to alleviate critical slowing down, which is present in any lattice simulation. While the test cases presented in this work are based on scalar fields, an extension to fermions is possible and is the subject of ongoing research.

Chapter 2

Lattice field theory

We now introduce fundamental ideas in lattice field theory that are essential for understanding the subsequent chapters. We begin with an overview of general concepts, such as the discretization of continuum theories. This is followed by a brief introduction of stochastic sampling methods and the sign problem. We then discuss lattice simulations for non-relativistic fermions, including the decoupling of fermionic interactions and projective methods for canonical computations in detail.

Since we cannot possibly cover every aspect, readers interested in a more comprehensive overview, particularly in the context of QCD, are encouraged to refer to the standard literature [36–38]. Additional resources focusing on condensed matter and cold atom systems include [39, 40].

2.1 Basics

The study of many-body quantum systems typically begins with the partition function, which is expressed in terms of the Hamiltonian $\hat{H}$ and the inverse temperature $\beta = 1/T$ as

$$Z = \mathrm{Tr}\left[e^{-\beta\hat{H}}\right], \tag{2.1}$$

where the trace is taken over the Hilbert space of the system. In this formalism, the expectation value of observables can be computed as

$$\langle\hat{O}\rangle = \frac{1}{Z}\mathrm{Tr}\left[\hat{O}\,e^{-\beta\hat{H}}\right]. \tag{2.2}$$

Another way to express the partition function is by using the Euclidean path integral formalism of quantum field theory. This approach involves Wick rotating the time coordinate from the real-time t_M to $-it$, where t now represents Euclidean time. As a result, the time evolution operator $e^{-i\hat{H}t_M}$ becomes its thermal counterpart $e^{-\hat{H}t}$, which is featured in the partition function (2.1). We may thus rewrite a given QFT into a thermal field theory using this imaginary time formalism. Given the Euclidean action S_E, the partition function can be written as

$$Z = \int \mathcal{D}\phi\, e^{-S_E[\phi]}, \tag{2.3}$$

where the field is a function of the spatial position and euclidean time, $\phi = \phi(x,t)$. The latter is restricted to the interval $[0,\beta]$. The compactness implies the need for suitable boundary conditions, which follow directly from the cyclicity of the trace in (2.1). The boundaries are periodic and

anti-periodic for bosonic and fermionic field variables, respectively. Crucially, the integral over field configurations in (2.3) now takes the form of an integral over a Boltzmann weight, allowing for stochastic sampling as long as the Euclidean action is real and the partition function finite.

2.1.1 Discretization & Cutoff

While we now have an integral over a probability distribution, the degrees of freedom are still continuous fields and not directly accessible for numerical computations on digital computers. This issue is addressed by defining the field values only on a finite lattice, which creates a regular grid in space and time. The spatial lattice comes with an associated lattice spacing a, denoting the distance between neighboring lattice sites. Moreover, it has a linear extent L, which is related to the total volume in units of the lattice spacing as $V = L^d$, where d is the number of spatial dimensions. Such a discretization introduces a cubic momentum cutoff $\Lambda = \pi/a$ to the theory, restricting momentum integrals,

$$\int \frac{d^d k}{(2\pi)^d} \to \int_{-\pi/a}^{\pi/a} \frac{d^d k}{(2\pi)^d}, \tag{2.4}$$

to the first Brillouin zone of the lattice. Furthermore, the lattice introduces a natural IR cutoff, related to its inverse linear extent $1/L$. It renders the above integral discrete, leaving a finite grid in both momentum and position.

Due to the space-time symmetry in relativistic theories, the discretization of the time direction can proceed just the same as for the spatial directions. In non-relativistic theories, however, the time direction must be treated separately, which we will discuss in Section 2.3.1.

Some care must be taken when discretizing the derivative operators in the action, as naïve discretization can lead to lattice artifacts, such as the fermion doubling problem in relativistic theories. While non-relativistic fermions are typically free of this problem, various discretization schemes are still worth considering, as they impact the approach to the continuum limit and the practical speed of computations. We will use $a = 1$ for the following discussion, which fixes the lattice unit. The central difference formula yields the naïve lattice Laplacian,

$$\Delta_{mn} = \sum_{\mu} \left[-2\delta_{m,n} + \delta_{m,n+\hat{\mu}} + \delta_{m,n-\hat{\mu}} \right], \tag{2.5}$$

where the sum runs over all lattice directions, and $\hat{\mu}$ denotes the unit vector pointing to the neighboring lattice site. As a result, free particles carry the energy dispersion relation

$$E(k) = 2 \sum_{\mu} \sin^2 (k_\mu/2), \tag{2.6}$$

where $k_\mu = \frac{2\pi n_\mu}{L}$ are the allowed momenta in the first Brillouin zone with integer n_μ. The lattice model is only a good approximation for small momenta, $k \ll \pi/a$, since the free dispersion of non-relativistic particles is given by $E(k) = k^2/2$ using $m = 1$.

Alternatively, we may introduce a discretization that reproduces the continuum dispersion exactly at all momenta. This can be achieved by simply taking the Fourier transform of the desired dispersion, yielding

$$\Delta_{m,n} = \int_{-\pi}^{\pi} d^d p \, \frac{p^2}{2} \, e^{i\boldsymbol{p}(\boldsymbol{m}-\boldsymbol{n})}, \tag{2.7}$$

where the bold letters indicate volume sized vectors. Note that the continuous Fourier transform is replaced by a discrete Fourier transform (DFT) for finite-sized lattices. The form above brings clear advantages when considering the approach to the continuum limit; see, e.g., [40–42]. However, it also comes with the disadvantage of a dense Laplacian in position space. This contrasts the simple finite difference discretization, which is sparse and thus computationally more efficient, needing $O(N)$ operations for a matrix-vector multiplication, where N is the number of lattice sites. The dense matrix naïvely requires $O(N^2)$ operations, which can be mitigated by performing Fourier transforms on the vectors at runtime, reducing the complexity to $O(N \log N)$.

2.1.2 The continuum limit

With the system now on a finite grid, it is crucial to consider how the continuum is approached as the spacing decreases and the volume increases. For lattice systems, the two aforementioned limits roughly correspond to the thermodynamic limit, $L \to \infty$, and the continuum limit, $a \to 0$. They are usually taken via extrapolations and performed subsequently, starting with the thermodynamic limit, to arrive at a well-defined theory in the continuum. Moreover, the physical theory possesses several scales, which can be made dimensionless via combinations and must be reproduced by the lattice theory in the continuum limit. To accomplish this, the bare parameters of the lattice action are adjusted to replicate the physical scales for each discretization. Given some dimensionless observable $\theta(\lambda(a), a)$, where $\lambda(a)$ denotes the set of bare parameters that depend on the lattice spacing, we can formulate the requirement as

$$\frac{d}{da}\theta(\lambda(a), a) = 0. \tag{2.8}$$

Numerous observables can be used to tune the couplings of the lattice theory, and ones where the running is known exactly a priori are most attractive. In the case of two-dimensional Fermi gases, such a parameter is given by the two-body binding energy, which can be computed exactly for arbitrary lattice parameters and will be used for scale setting in Chapter 5 and Chapter 6.

While the scales remain fixed in physical units, they must change in units of the lattice spacing, and length scales such as the correlation length ξ diverge in the continuum limit. Consequently, the continuum limit is associated with a critical point, and lattice systems must necessarily be close to criticality to reproduce continuum physics. We refer to [36] for more details on this.

In addition to UV effects induced by the finite lattice spacing, lattice simulations can also show a strong dependence on the IR cutoff due to the finite size of the system. In cold atom computations, this behavior stems from the large thermal wavelength $\lambda_T = \sqrt{2\pi\beta}$, which increases with decreasing temperature. As a result, wavefunctions of particles begin to overlap with themselves on small lattices with periodic boundaries, and larger lattices $L \gg \lambda_T$ are required to obtain reliable results.

2.2 Stochastic sampling

In this section, we introduce the notion of stochastic sampling for high dimensional integrals, which forms the basis of numerical computations in lattice field theory. After discretization, the field integral over configurations becomes a high dimensional integral over real-valued variables, that is generally not solvable analytically. In addition, the integration domain's high dimensionality

makes it infeasible to discretize the entire space for numerical integration. Instead, we may sample the integrand stochastically, drawing more samples in the region where the probability weight is large while neglecting those where the weight vanishes. The goal is to sample a finite set of configurations ϕ_n, which are distributed according to the weight in (2.3),

$$\phi_n \sim e^{-S_E[\phi]}, \tag{2.9}$$

such that expectation values of physical observables can be computed as averages over the finite set of samples as

$$\langle O \rangle \approx \frac{1}{N} \sum_{n=1}^{N} O(\phi_n). \tag{2.10}$$

Stochastic sampling also introduces statistical errors beyond the systematic finite size and difference effects discussed above. The sample size controls the errors, which scale with a factor of $1/\sqrt{N}$. Reducing uncertainties to a point where useful predictions are possible can be challenging since the sample size needs to grow quadratically compared to the reduction in error.

In practice, sampling is facilitated by a stochastic process. Markov chains are commonly used, in which subsequent configurations depend solely on the previous field values. Field configurations are updated according to a transition probability, which is chosen such that the stationary distribution of the Markov chain is the desired distribution. One simple and frequently used way to construct Markov processes is via the detailed balance condition, which is defined for a transition probability $P(\phi \to \phi')$ and stationary distribution $\pi(\phi)$ as

$$\pi(\phi) P(\phi \to \phi') = \pi(\phi') P(\phi' \to \phi). \tag{2.11}$$

This condition is the basis of the most commonly used algorithms for lattice computations, including Hybrid and Metropolis Monte Carlo, as we will briefly discuss in the following sections.

2.2.1 General recipe

The collection of samples follows a fairly generic recipe independent of the specific algorithm used. The usual steps involved in a lattice simulation are as follows:

1. Set a seed for the random number generator and initialize the field configuration ϕ, either at a random or predetermined value. The former is done to ensure reproducibility of random number streams.

2. Thermalize the Markov chain. To ensure that expectation values are not biased by the initial configurations, we perform a certain number of update steps according to $P(\phi \to \phi')$ before collecting samples. The number of steps required for thermalization depends on the system and the specific algorithm used. It can be determined by monitoring the expectation values of observables, which should stabilize after thermalization.

3. Samples are collected by stepping the stochastic process according to the transition probability. Usually, not every sample needs to be saved, as consecutive samples can be highly correlated, especially in local updating schemes.

4. (Optional) It can be desirable to compute observables based for the configurations on the fly during sampling. This is the case when much of the computational effort, such as computing matrix determinants, is the same between sampling and evaluation. On the other hand, if sampling is not the limiting factor, we may only store the observables instead of the actual configurations.

5. Save the final configuration and the state of the random number generator. This enables us to restart the computation from where it ended and construct a longer chain if we desire. In case of an interruption, it is advisable to also perform intermediate saved states during the simulation.

2.2.2 Metropolis Monte Carlo

Metropolis Monte Carlo is a simple but widely used method for sampling the weight via local updates of the field configurations. Starting from an initial configuration ϕ, a new configuration ϕ' is proposed by updating one or more of the field variables via some proposal probability $P_P(\phi \to \phi')$. The latter is usually chosen as an independent symmetric distribution on each site,

$$P_P(\phi \to \phi') = \prod_{\phi_i \in \{\phi_i\}} P_P(\phi_i' \to \phi_i), \tag{2.12}$$

where $\{\phi_i\}$ is the set of sites to be updated. In order to satisfy the detailed balance condition for the desired weight, the update is accepted with a likelihood given by the Metropolis acceptance probability

$$P_A(\phi \to \phi') = \min\left(1, \frac{\pi(\phi')\, P_P(\phi' \to \phi)}{\pi(\phi)\, P_P(\phi \to \phi')}\right). \tag{2.13}$$

If the proposal distribution is symmetric when exchanging the previous and new configuration, it cancels out, and the acceptance probability simplifies to a function of the weight ratio alone.

Assuming the target weight to be $\exp(-S_E)$ with symmetric proposals, a single step in the Metropolis algorithm proceeds as follows:

1. Select one or more lattice sites to update. This selection can be performed either randomly or systematically. In the systematic approach, we can introduce the concept of sweeps. A single sweep refers to having made a single proposal at each lattice site.

2. Propose a new configuration ϕ' according to the proposal distribution $P_P(\phi \to \phi')$.

3. Accept the new configuration with probability $P_A(\phi \to \phi') = \min\left(1, e^{-S_E(\phi') + S_E(\phi)}\right)$ by comparing the weight ratio to a random number drawn from a uniform distribution. Note that local updates can be beneficial, as they allow us to only consider the local change in energy of the action ΔS_{loc}. This reduces the cost associated with computing the full action.

4. If the update is accepted, the new configuration is added to the Markov chain. Otherwise, the previous configuration is restored, and another copy is added to the chain.

2.2.3 Global updating

Due to its simplicity, the Metropolis-Hastings algorithm is very attractive and remains widely used for condensed matter and non-relativistic systems [41, 42]. However, systems with large correlation lengths (in units of the lattice spacing) can often suffer from painfully low acceptance rates or require very narrow proposal distributions. This situation is perhaps best understood in the context of the two-dimensional Ising model close to criticality. Let us consider the likelihood of accepting a single spin flip. We will find it very low, as a large correlation length ξ implies that the surrounding spins will likely have the same orientation as the spin we are trying to flip.

To partially tackle this issue, there are several algorithms available that discard the local updates in favor of *global* ones proposed in a way that keeps acceptance rates high. Mixed forms also exist, which, for instance, perform updates for all spatial points at the same time and iterate over the slices. Such an approach is used for certain computations of ultracold fermion systems later in this work and is described in more detail in Section 4.2.2.

Some of the most successful approaches rely on updates based on the gradient of the action with respect to the fields, called drift and given by

$$K_i = -\frac{\partial S_E}{\partial \phi_i}. \tag{2.14}$$

These include the Hybrid Monte Carlo (HMC) algorithm, which is widely used in lattice QCD, and the Langevin algorithm. The former will be briefly discussed in the following, while the latter is discussed in detail in Chapter 3.

In Hybrid Monte Carlo, additional variables conjugate to the fields are introduced as momenta p_i, sampled from a unit variance Gaussian distribution with zero mean. The fields and momenta are then evolved according to the Hamiltonian equations of motion, given by

$$\frac{d\phi_i}{d\tau} = p_i, \quad \frac{dp_i}{d\tau} = -\frac{\partial S_E}{\partial \phi_i}, \tag{2.15}$$

where τ is the fictitious time. These coupled differential equations are solved using numerical integration schemes, which must be symplectic to preserve the phase space volume. If we had a perfect integrator, no additional step would be necessary, and the problem would be reduced to sampling random momenta and solving the equations of motion. However, in practice, the numerical integration introduces a systematic error, which must be corrected. This is done by inserting an additional accept/reject step based on the change in the Hamiltonian,

$$\Delta H = S_E(\phi', p') - S_E(\phi, p) + \frac{1}{2}\left(p'^2 - p^2\right), \tag{2.16}$$

which accepts the new configuration with probability $P_A = \min\left(1, e^{-\Delta H}\right)$. For more details on HMC, we refer to [36].

2.2.4 Autocorrelation & Critical slowing down

The most important characteristic for assessing the performance of different sampling approaches beyond the raw computation time is the autocorrelation time between samples in the Markov

chain. Roughly speaking, the autocorrelation time measures the number of steps required for samples to become uncorrelated. Given some observable, the time-lagged correlation function is given by

$$C_{\hat{O}}(\tau) = \langle \left(\hat{O}_i - \langle \hat{O} \rangle\right) \left(\hat{O}_{i+t} - \langle \hat{O} \rangle\right) \rangle \tag{2.17}$$

where $\hat{O}_i$ indicates the value of the observable at the i-th step in the Markov chain. The correlation function only depends on the time lag t since we are assuming an equilibrated process. This correlation function is expected to show the long time behavior

$$C_{\hat{O}}(t) \sim e^{-t/\tau_{\text{exp}}}, \tag{2.18}$$

where τ_{exp} is the exponential autocorrelation time. Note that this time generally depends on which observable we are considering, and we need to take the maximum over all observables to measure the overall autocorrelation in the system. In practice, computing the exponential autocorrelation time by fitting is often not feasible due to the large number of samples required for reliable results. Instead, the integrated autocorrelation time is used, which can be computed directly from the lagged correlation function, see, e.g., [38].

Given the estimate of the autocorrelation time, the *effective* number of samples can be computed by a rescaling as the total number of steps in the Markov chain as

$$N_{\text{eff}} = \frac{N}{2\tau_{\text{exp}}}, \tag{2.19}$$

while the estimator for the standard error of the mean is given by

$$\sigma^2 = \frac{2\tau_{\text{exp}}}{N} \left(\langle \hat{O}^2 \rangle - \langle \hat{O} \rangle^2\right). \tag{2.20}$$

Since the computation of autocorrelation times can be computationally expensive and labor-intensive, more straightforward methods are often used in practice. For instance, a simple blocking analysis can be performed, where the Markov chain is split into blocks of a specific size, for which the means are computed. The standard errors can then be computed based on the mean blocks. The block size is increased until the standard error stabilizes. Related to this approach, we may use the embarrassingly parallel nature of Markov chain Monte Carlo simulations to compute several statistically independent chains, from which the standard error can be computed. We take this approach for most of the computations of the ultracold Fermi systems in this work, where we run independent Markov chains on each hardware thread of a multi-core CPU. Each Markov chain carries a different random number seed, and sampling can only be started when the configurations between the chains are completely independent. This requires some thermalization on each of the threads, which can come with a significant overhead but is found to be unproblematic in most cases.

Notably, the autocorrelation time is not a function of the sampling algorithm alone but directly related to the physical system under study. In particular, it depends on the correlation length as

$$\tau_{\text{exp}} \sim \xi^z, \tag{2.21}$$

where z is a dynamical critical exponent set by the sampling algorithm. As the correlation length necessarily diverges in the continuum limit, the autocorrelation time will also diverge, a phenomenon

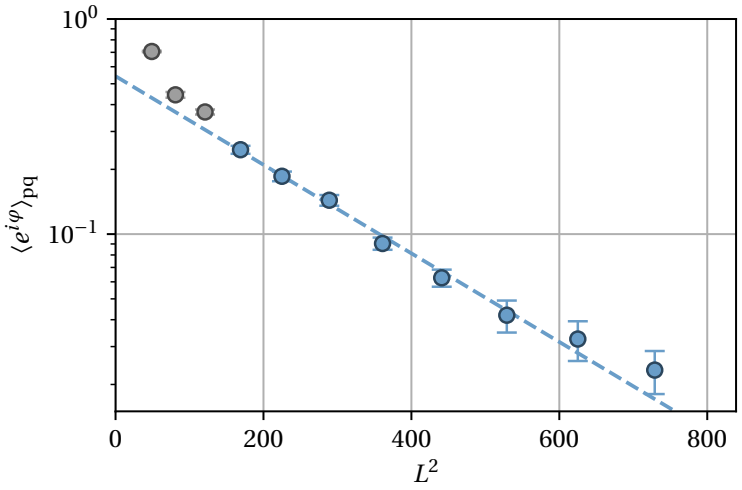

Figure 2.1: Visualization of the volume dependence of the average sign on the example of a spin imbalanced Fermi gas in two spatial dimensions. The sign is computed while keeping all physical scales except for the volume fixed at $\beta e_b = \beta\mu = \beta h = 3$. The average sign shows a clear exponential decrease as the volume is increased. The dashed line indicates a fit to the data, where only the colored points were considered.

known as *critical slowing down*. Consider a d-dimensional square lattice, where we keep the correlation length fixed in units of the linear extent of the lattice $\xi/L = \xi/V^{1/d} = const$. The above relation then implies that the autocorrelation time will scale with the volume as

$$\tau_{\mathrm{exp}} \sim V^{z/d}. \tag{2.22}$$

One of the foremost aims when developing new sampling algorithms for lattice systems is the reduction of the exponent z. Amongst the more recent attempts is the use of a machine learning method called normalizing flows [43, 44], which aims to find an invertible map from the desired distribution to a trivial one. This idea is similar to the idea of trivializing flows introduced by Lüscher [45] and can, in principle, eliminate autocorrelation times by sampling the trivial distribution precisely. However, practical implementations are still in their infancy and generally show unfavorable scaling with increasing volume [46, 47]. Further discussion of normalizing-flows can be found in Chapter 7.

2.2.5 The sign problem

So far, we have only considered the possibility of real actions, leading to a straightforward interpretation of $\exp(-S)$ as a probability weight. While this is typically the situation encountered in statistical systems, it is not generally true in quantum ones. Examples of systems with complex actions include lattice QCD at finite chemical potential, the repulsive Hubbard model away from half-filling, attractive Fermi gases at finite spin imbalance, and any real-time formulation of quantum field theories. The complex action immediately suggests a weight that is either complex or not positive definite. This disrupts the probabilistic interpretation and creates what is commonly

known as a *sign* or *phase problem*. The term '*Fermion sign problem*' often appears in the literature because complex weights frequently arise in systems with fermions. However, such behavior is not exclusive to fermions, nor do all fermionic systems exhibit a sign problem. Bosons may also suffer from sign problems, particularly when considering real-time computations where the action gives a pure phase $\exp(iS)$. On the other hand, some fermionic systems, such as QCD at zero chemical potential and spin-balanced fermions with contact interactions, are sign problem free.

Without the probabilistic interpretation, the direct application of the previously discussed approaches for stochastic sampling is not possible. Naïvely, a straightforward solution to the sign problem is to simply sample the absolute value of the weight, which we call the phase-quenched, and compute the expectation values of observables as

$$\begin{aligned}\langle O\rangle &= \frac{\int d\phi\, O(\phi)\, e^{-S}}{\int d\phi\, e^{-S}} = \frac{\int d\phi\, O(\phi)\, e^{i\varphi}\left|e^{-S}\right|}{\int d\phi\, e^{i\varphi}\left|e^{-S}\right|} \\ &= \frac{\int d\phi\, O(\phi)\, e^{i\varphi}\left|e^{-S}\right|}{\int d\phi\, e^{i\varphi}\left|e^{-S}\right|} \frac{\int d\phi\, e^{i\varphi}\left|e^{-S}\right|}{\int d\phi\, e^{i\varphi}\left|e^{-S}\right|} \\ &= \frac{\langle O e^{i\varphi}\rangle_{\mathrm{pq}}}{\langle e^{i\varphi}\rangle_{\mathrm{pq}}}. \end{aligned} \tag{2.23}$$

Here, φ is the phase of e^{-S} associated with a given lattice configuration, and $\langle\rangle_{\mathrm{pq}}$ denotes expectation values with respect to the phase quenched weight. While this procedure is mathematically sound as long as the *average sign* $\langle e^{i\varphi}\rangle_{\mathrm{pq}}$ does not vanish identically, problems arise already when it is finite but much smaller than one. In that case, samples with different phases can be understood to cancel, introducing a large amount of noise in which the signal is buried. Indeed, error estimates will heavily depend on the relative error of the average sign, which typically increases as the average sign tends to zero.

Unfortunately, the average sign decreases exponentially with the system's volume. The free energy density f is related to the partition function as

$$Z = e^{\beta V f}, \qquad Z_{\mathrm{pq}} = e^{\beta V f_{\mathrm{pq}}}, \tag{2.24}$$

for the full and phase quenched partition functions, respectively, where βV is the total spacetime volume. The average sign can be expressed as the ratio of the two partition functions

$$\langle e^{i\varphi}\rangle_{\mathrm{pq}} = \frac{Z_{\mathrm{pq}}}{Z} = e^{\beta V(f_{\mathrm{pq}} - f)}, \tag{2.25}$$

and the exponential decrease is clear, as the free energy density is an intensive quantity. This behavior is illustrated in Figure 2.1, where the average sign is computed for a spin-imbalanced Fermi gas in two spatial dimensions as a function of the spatial volume of the lattice.

Finally, we note that a severe sign problem can have different effects on different observables, depending on the correlation between the phase and the observable we are interested in. Particularly, if

$$\langle O e^{i\varphi}\rangle_{\mathrm{pq}} = \langle O\rangle_{\mathrm{pq}} \langle e^{i\varphi}\rangle_{\mathrm{pq}}, \tag{2.26}$$

the sign problem is not relevant for computations of the Observable O, and the phase-quenched expectation value becomes exact. There may also be cases where the relation (2.26) holds only approximately, and the phase-quenched result is still a good approximation to the full one. Later on, we will come across a similar situation in Chapter 6 while discussing the two-dimensional spin-imbalanced Fermi gas. Generally, the correlation must be taken into account when propagation errors for (2.23). In the cases of vanishing correlation, the error is computed via the standard formula, which neglects covariances.

2.3 Fermions on the lattice

We will now focus on fermionic systems and how they are made amenable to lattice simulations. In contrast to bosons, fermionic actions are integrals over anti-commuting Grassmann variables, which cannot be represented directly on digital computers. Moreover, in relativistic theories, the fermion doubling problem arises when discretizing the Dirac operator, leading to unwanted modes in the continuum limit. Fortunately, the latter issue is absent in the non-relativistic formulations considered in this work.

The continuum Hamiltonian for a two-component Fermi gas with contact interactions is given by

$$\hat{H} = \int d^d x \sum_\sigma \hat{\psi}^\dagger_\sigma(x) \frac{-\nabla^2}{2m} \hat{\psi}_\sigma(x) + g \int d^d x \, \hat{n}_\uparrow(x) \hat{n}_\downarrow(x), \tag{2.27}$$

where $\hat{\psi}_\sigma(x)$ is the fermionic annihilation operator for spin σ, $\hat{n}_\sigma(x) = \hat{\psi}^\dagger_\sigma(x)\hat{\psi}_\sigma(x)$ is the density operator, and g is the bare coupling constant. In practice, one tunes g to reproduce a desired scattering length to match the physical system. This is accomplished by, for example, computing the two-body binding energy exactly, see Chapter 5 and Chapter 6. To make contact with the action formalism introduced in the previous section, we may write the action as an integral over the Hamiltonian in terms of Grassmann fields ψ and $\psi^\dagger$, yielding

$$S = \int_0^\beta d\,t \int d^d x \sum_\sigma \psi^\dagger_\sigma \left(\partial_t - \frac{\nabla^2}{2m} - \mu_\sigma \right) \psi_\sigma + g \psi^\dagger_\uparrow \psi_\uparrow \psi^\dagger_\downarrow \psi_\downarrow. \tag{2.28}$$

Here, we have introduced the species-dependent chemical potential μ_σ, which controls the particle number. Moreover, we have identified the inverse temperature β with the length of the imaginary time direction.

The lattice Hamiltonian is obtained by a straightforward discretization, yielding

$$\begin{aligned} \hat{H} &= \sum_p \hat{\psi}^\dagger_{p,\sigma} \epsilon_p \hat{\psi}_{p,\sigma} + \lambda \sum_x \hat{\psi}^\dagger_{x,\uparrow} \hat{\psi}_{x,\uparrow} \hat{\psi}^\dagger_{x,\downarrow} \hat{\psi}_{x,\downarrow}, \\ &= \hat{T} + \hat{V} \end{aligned} \tag{2.29}$$

where $\hat{\psi}^\dagger_{p,\sigma}$ and $\hat{\psi}_{p,\sigma}$ are the creation and annihilation operators for a fermion with momentum p and spin σ. The lattice coupling constant λ generally differs from the continuum value g and depends on the momentum cutoff. In the second line, we have explicitly separated the Hamiltonian into kinetic $\hat{T}$ and potential $\hat{V}$ terms. The sums run over all momenta or positions on the lattice. We use the dispersion

$$\epsilon_p = \frac{p^2}{2m} \tag{2.30}$$

unless stated otherwise, which reproduces the free particle dispersion for all momenta. Using a first-order finite-difference discretization of the lattice Laplacian instead yields the Fermi-Hubbard Hamiltonian, which has also been used in the context of ultracold Fermi gases, but shows slower convergence to the continuum limit [35, 40–42, 48].

2.3.1 Time discretization

When discretizing the time direction in (2.28), it is most convenient to stick with the Hamiltonian formalism. The time evolution operator in the partition function is sliced into N_t components of length $\Delta t = \beta / N_t$, and the partition function becomes

$$Z = \mathrm{Tr}\left[e^{-\beta \mu_\sigma \hat{N}_\sigma} \prod_{N_t} e^{-\Delta t \hat{H}} \right]. \tag{2.31}$$

Here, the components of the product are propagators from one time slice to the next. We have excluded the chemical potential from the product as an aesthetic choice, since the time discretization does not affect it. Given a sufficiently small time step, the propagators can now be split into the kinetic and interaction terms via a Trotter-Suzuki decomposition [49]. This allows us to express the exponential of a sum of operators as a product of the exponentials of the individual operators, yielding to second order

$$e^{-\Delta t \hat{H}} = e^{-\Delta t \hat{T}} e^{-\Delta t \hat{V}} + \mathcal{O}(\Delta t^2). \tag{2.32}$$

A third-order equation can be obtained via a symmetric decomposition as

$$e^{-\Delta t \hat{H}} = e^{-\frac{\Delta t}{2} \hat{T}} e^{-\Delta t \hat{V}} e^{-\frac{\Delta t}{2} \hat{T}} + \mathcal{O}(\Delta t^3), \tag{2.33}$$

and is widely used in cold atom systems in the ground state and at finite temperature[32, 40, 50]. We use this formulation for ground state computations but employ the second-order one at finite temperature, as the cyclic nature of the trace makes the approaches equivalent.

2.3.2 Auxiliary field transformatios

Superficially, little has been gained by discretizing the fermionic system so far, as we still need real-valued fields that can be sampled stochastically. This can be achieved by the introduction of auxiliary fields via a Hubbard-Stratonovich transformation [51]. We will now discuss various available transformations for the Hamiltonian in (2.29) and their usefulness in different scenarios.

The Hubbard-Stratonovich transformation is a general method used to decouple density-density interactions in the Hamiltonian by replacing them with an integral over an auxiliary field. For our purposes, there will be four major variants of the transformation. We begin by introducing discrete fields, which directly couple to the total density, as originally presented in [52],

$$\begin{aligned} e^{-\lambda \Delta t \hat{n}_{x,\uparrow} \hat{n}_{x,\downarrow}} &= 1 + \sum_{n=1}^{N} \frac{(-\lambda \Delta t)^n}{n!} \hat{n}_{x,\uparrow} \hat{n}_{x,\downarrow} \\ &= 1 + (e^{-\lambda \Delta t} - 1) \hat{n}_{x,\uparrow} \hat{n}_{x,\downarrow} \\ &= \frac{1}{2} \sum_{\phi = \pm 1} \left(1 + A\phi \hat{n}_{x,\uparrow}\right)\left(1 + A\phi \hat{n}_{x,\downarrow}\right). \end{aligned} \tag{2.34}$$

In the first line, we used the series expansion of the exponential and the density operators' idempotence. This transformation is performed separately for each spatial point on each time slice, leading to a total of $N_t L^d$ auxiliary fields, which we label as $\phi_{t,x}$. The effective coupling of the auxiliary field to the density is $A = \sqrt{(e^{-\lambda\Delta t} - 1)}$. It is real and positive for $\lambda < 0$, corresponding to attractive inter-particle interactions, and purely imaginary for repulsive interactions $\lambda > 0$. The transformation above is quite attractive for Metropolis-based sampling, as the state space is large but finite and equivalent to the Ising model, which comes with a large machinery of efficient sampling algorithms. However, if we want to apply global updating schemes such as HMC or Langevin sampling, we require a continuous auxiliary field. To achieve this, a slight modification to the transformation is necessary, yielding

$$e^{-\lambda\Delta t \hat{n}_{x,\uparrow}\hat{n}_{x,\downarrow}} = \frac{1}{2\pi}\int_{-\pi}^{\pi} d\phi \left(1 + A_c\,\hat{n}_{x,\uparrow}\sin\phi\right)\left(1 + A_c\,\hat{n}_{x,\downarrow}\sin\phi\right). \tag{2.35}$$

To verify the above expression, we can integrate the right-hand side and then compare it with (2.34). The continuous auxiliary field ϕ is defined on the interval $[-\pi,\pi)$, such that the overall state space $[-\pi,\pi)^{N_t L}$ is bounded but continuous. The effective coupling constant, denoted by A_c, is modified slightly and becomes $A_c = \sqrt{2(e^{-\lambda\Delta t} - 1)}$. This value remains real for attractive interactions and imaginary for repulsive ones. We will use the continuous auxiliary field decomposition later for complex Langevin simulations.

In addition to the density-channel decomposition, we can introduce transformations that couple to the local magnetization, also called spin-channel transformations. They are related to the density-channel decompositions by a simple sign flip, yielding

$$e^{-\lambda\Delta t \hat{n}_{x,\uparrow}\hat{n}_{x,\downarrow}} = \frac{1}{2}\sum_{\phi=\pm 1}\left(1 + B\,\phi\,\hat{n}_{x,\uparrow}\right)\left(1 - B\,\phi\,\hat{n}_{x,\downarrow}\right). \tag{2.36}$$

for the discrete field and

$$e^{-\lambda\Delta t \hat{n}_{x,\uparrow}\hat{n}_{x,\downarrow}} = \frac{1}{2\pi}\int_{-\pi}^{\pi} d\phi \left(1 + B_c\,\hat{n}_{x,\uparrow}\sin\phi\right)\left(1 - B_c\,\hat{n}_{x,\downarrow}\sin\phi\right). \tag{2.37}$$

for the bounded but continuous one. The coupling parameters become $B = \sqrt{(1 - e^{-\lambda\Delta t})}$ and $B_c = \sqrt{2(1 - e^{-\lambda\Delta t})}$. They are now real for repulsive interactions and imaginary for attractive ones. This kind of formulation is often used in the context of the two-dimensional Hubbard model at half-filling, as it is free of the sign problem.

Note that the abovementioned list of transformations is neither exhaustive nor unique, and other transformations are possible. For example, the one initially proposed in [51] is commonly used and leads to a continuous and unbounded field reminiscent of scalar lattice field theories. Moreover, a transformation into the coupling channel $\psi^\dagger_\uparrow\psi_\downarrow$, physically motivated by the emergence of BCS type pairing, was recently brought forward in [53]. Unfortunately, it doubles the number of degrees of freedom. It also introduces a complex-valued weight, which increases the computational cost of simulations and requires strategies such as complex Langevin to avoid the sign problem. Initial applications are limited to the case of a $0+1$-dimensional theory, and we do not use it here.

2.3.3 Computing the weight

We are now ready to cast the partition function into a form amenable to stochastic sampling. Consider a generic auxiliary field decomposition

$$e^{-\lambda \Delta t \hat{n}_{x,\uparrow}\hat{n}_{x,\downarrow}} = \int d\phi\, \hat{V}_\uparrow(\phi)\, \hat{V}_\downarrow(\phi)\,, \tag{2.38}$$

where the potential operators are bilinear in the creation and annihilation operators. All transformations discussed in the previous section allow for this structure. The contribution of a single time slice to the time evolution operator, accurate to second order in the time step, is given by

$$\begin{aligned} e^{-\Delta t\hat{T}}\, e^{-\Delta t\hat{V}} &= e^{-\Delta t \sum_\sigma \hat{T}_\sigma} \prod_x \int d\phi_{t,x}\, \hat{V}_\uparrow(\phi_{t,x})\, \hat{V}_\downarrow(\phi_{t,x})\,, \\ &= \int \prod_x d\phi_{t,x} \otimes_\sigma e^{-\Delta t \hat{T}_\sigma}\, \hat{V}_\sigma(\phi_{t,x})\,. \end{aligned} \tag{2.39}$$

Here, the spin species are decoupled via a tensor product by virtue of the auxiliary field transformation. The full partition function now becomes a time-ordered product,

$$\begin{aligned} Z &= \int \prod_{t,x} d\phi_{t,x}\, \mathrm{Tr}\left[\prod_t \otimes_\sigma e^{-\Delta t \hat{T}_\sigma}\, \hat{V}_\sigma(\phi_{t,x}) e^{\Delta t \mu_\sigma \hat{N}_\sigma} \right], \\ &= \int \prod_{t,x} d\phi_{t,x} \prod_\sigma \det\left[\mathrm{I} + e^{\beta\mu_\sigma} \prod_t e^{-\Delta t T} V_\sigma(\phi_t) \right], \\ &= \int \prod_{t,x} d\phi_{t,x} \prod_\sigma \det\left[\mathrm{I} + e^{\beta\mu_\sigma} \prod_t B^\sigma(\phi_t) \right], \end{aligned} \tag{2.40}$$

where the trace over Fock space has been carried out in the second line, resulting in a simple representation of the weight as the determinant of matrices of spatial volume size. The matrices

$$B^\sigma(\phi_t) = B_t^\sigma = e^{-\Delta t T} V_\sigma(\phi_t) = e^{\Delta t H_{\sigma,t}}\,, \tag{2.41}$$

can be interpreted as the time evolution operator in the presence of the auxiliary field.

A proof for the final line of (2.40) can be found, for example, in [54] using the Baker–Campbell–Hausdorff formula, or in [55], which expands the trace directly. Alternatively, [56] provides proof based on the fermionic action. While we do not provide the full proof here, we present a brief intuitive argument supporting the correctness of the formula. By interpreting the determinant as the result of an integration over fermionic fields, we can restore the temporal dimension using a simple bock matrix of the form

$$D_\sigma = \begin{pmatrix} \mathrm{I} & 0 & \cdots & 0 & B_1 \\ -B_2 & \mathrm{I} & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathrm{I} & 0 \\ 0 & \cdots & 0 & -B_{N_t} & \mathrm{I} \end{pmatrix}, \tag{2.42}$$

where all entries are matrices of size $L^d \times L^d$, and we have dropped the spin indices. Taking the determinant over both spatial and temporal dimensions naturally restores the form given in (2.40).

In terms of discrete field operators, we may now write the action as

$$S = \sum_{\sigma} \psi_{\sigma}^{\dagger} D_{\sigma} \psi_{\sigma} \,. \tag{2.43}$$

In the limit $\Delta t \to 0$, the time evolution matrices can be expanded to first order in the time steps $B_t^{\sigma} \approx \mathrm{I} + \Delta t H_{\sigma,t}$, leading to the action

$$\begin{aligned} S =& \Delta t \sum_{\sigma,t} \psi_{\sigma,t}^{\dagger} \left(\psi_{\sigma,t} - \psi_{\sigma,t-1} \right) / \Delta t + \psi_{\sigma,t}^{\dagger} H_{\sigma,t} \psi_{\sigma,t-1} \\ =& \sum_{\sigma} \int dt\, d^d x\, \psi_{\sigma}^{\dagger} \left(\frac{\partial}{\partial_t} + H_{\sigma} \right) \psi_{\sigma} \,, \end{aligned} \tag{2.44}$$

where we have taken the continuum limit in the second line and suppress the time argument. This is just the expected form of the action given out Hamiltonian.

We are now ready to calculate the weight in (2.40) by constructing the matrix product $U_{\sigma} = \prod_{\sigma} B_t^{\sigma}$ directly and subsequently computing the determinant. The procedure can be technically demanding due to the computational cost of matrix multiplications and the stability challenges posed by large condition numbers at low temperatures. A detailed discussion of these aspects can be found in Chapter 4.

2.3.4 The sign problem for two fermion species

Whether or not a sign problem occurs is influenced by the specifics of the chosen Hubbard-Stratonovich transformation and the nature of the coupling. In the absence of interactions, sampling becomes unnecessary since the weight is independent of the auxiliary field, thus eliminating the sign problem.

To analyze the interacting system, we begin with a density-channel decomposition under attractive interactions, where the matrices $V_{\sigma}(\phi)$ remain real. However, the matrix product can have negative eigenvalues, leading to configurations with negative determinants. The weight remains positive only if the two spin species have equal chemical potentials, resulting in spin-balanced populations, expressed as

$$\mu_{\uparrow} = \mu_{\downarrow} \to \det\left[\mathrm{I} + e^{\beta \mu_{\uparrow}} U_{\uparrow} \right] = \det\left[\mathrm{I} + e^{\beta \mu_{\downarrow}} U_{\downarrow} \right], \tag{2.45}$$

$$P(\phi) = \det\left[\mathrm{I} + e^{\beta \mu_{\uparrow}} U_{\uparrow} \right]^2 \in \mathbb{R}^+. \tag{2.46}$$

Here, the weight is real and non-negative, and we can apply the full arsenal of stochastic sampling methods. This construction can be trivially extended to larger numbers of fermion flavors with $SU(N)$ symmetry as long as N is even. However, if the populations are imbalanced, the weight, while remaining real, is not guaranteed to be positive, and a sign problem may arise in simulations. In the presence of repulsive interactions, the potential matrix becomes imaginary, resulting in an inherent sign problem. The Hubbard model at half-filling is a notable exception to this, where the positivity of the weight is maintained through a particle-hole transformation [57].

In spin channel, the situation is similar, albeit slightly more involved. For attractive interactions, the potential matrices are purely imaginary, and the matrix products are related to each other via complex conjugation,

$$U_{\uparrow} = U_{\downarrow}^{*} \,. \tag{2.47}$$

Channel	$\lambda < 0$ $\mu_\uparrow = \mu_\downarrow$	$\lambda < 0$ $\mu_\uparrow \neq \mu_\downarrow$	$\lambda > 0$ $\mu_\uparrow = \mu_\downarrow$	$\lambda > 0$ $\mu_\uparrow \neq \mu_\downarrow$
Density	$\mathbb{R}^+$	$\mathbb{R}$	$\mathbb{C}$, HF: $\mathbb{R}^+$	$\mathbb{C}$
Spin	$\mathbb{R}^+$	$\mathbb{C}$	$\mathbb{R}$, HF: $\mathbb{R}^+$	$\mathbb{R}$
Coupling	$\mathbb{C}$	$\mathbb{C}$	$\mathbb{C}$	$\mathbb{C}$

Table 2.1: Summary of the sign problem for the different transformations and physical scenarios. The entries denote the domain of the weight. HF indicates half filling.

By utilizing the transposition invariance of the determinant, we can replace the complex conjugation with the Hermitian adjoint, allowing the weight to be written as

$$P(\phi) = \det\left[\mathrm{I} + e^{\beta\mu_\uparrow} U_\uparrow\right] \det\left[\mathrm{I} + e^{\beta\mu_\downarrow} U_\uparrow^\dagger\right] = \left|\det\left[\mathrm{I} + e^{\beta\mu_\uparrow} U_\uparrow\right]\right|^2 \in \mathbb{R}^+, \tag{2.48}$$

where we have assumed $\mu_\uparrow = \mu_\downarrow$ in the last line. If the spin species are imbalanced, the weight becomes complex instead, incurring a sign problem. Although all matrices are real in systems with repulsive interactions, negative weights may appear, as in the spin-imbalanced density channel case. Nevertheless, half-filling remains free of sign problems. A summary of the sign problem for the different transformations and physical scenarios can be found in Table 2.1.

2.3.5 Observables for lattice fermions

Equipped with a suitable weight, we now proceed to compute physical observables from the sampled lattice configurations. Given some bilinear operator $\hat{O} = \hat{\psi}^\dagger_{\sigma,i} O_{ij} \hat{\psi}_{\sigma,j}$, expectation values are given by

$$\langle O(t)\rangle^\phi = \frac{\mathrm{Tr}\left[e^{-t\hat{H}_\phi} \hat{O} e^{-(\beta-t)\hat{H}_\phi}\right]}{\mathrm{Tr}\left[e^{-\beta\hat{H}_\phi}\right]} \tag{2.49}$$

where we have inserted the operator at time t of the imaginary time evolution, and $\hat{H}_\phi$ indicates the Hamiltonian in the background of a specific auxiliary field. Note that the superscript ϕ indicates that we evaluate the observable on a specific field configuration. The average over all samples gives the actual expectation value of the physical system. For simplicity, we will omit the spin index and chemical potential from here on, as the latter can be absorbed into the Hamiltonian.

Next, we insert an exponential and rewrite the expectation value in terms of a derivative with respect to an external source J as

$$\begin{aligned}\langle O(t)\rangle^\phi &= \frac{\partial}{\partial J}\left(\log \mathrm{Tr}\left[e^{-t\hat{H}_\phi} e^{J\hat{O}} e^{-(\beta-t)\hat{H}_\phi}\right]\right)\Big|_{J=0} \\ &= \frac{\partial}{\partial J}\left(\log\det\left[\mathrm{I} + B_1 \dots B_t e^{JO} B_{t+1} \dots B_{N_t}\right]\right)|_{J=0}, \end{aligned} \tag{2.50}$$

where the second step introduced the matrix representation, as before. The logarithm of the deter-

minant can be replaced by a trace logarithm, allowing for a direct derivative calculation, yielding

$$\begin{aligned}\langle O(t)\rangle^{\phi} &= \frac{\partial}{\partial J}\left(\mathrm{Tr}\log\left[\mathrm{I}+B_1\dots B_t e^{JO}B_{t+1}\dots B_{N_t}\right]\right)|_{J=0}\,,\\ &= \mathrm{Tr}\left[\left(\mathrm{I}+B_1\dots B_t e^{JO}B_{t+1}\dots B_{N_t}\right)^{-1}B_1\dots B_t O e^{JO}B_{t+1}\dots B_{N_t}\right]|_{J=0}\,,\\ &= \mathrm{Tr}\left[B_{t+1}\dots B_{N_t}\left(\mathrm{I}+B_1\dots B_{N_t}\right)^{-1}B_1\dots B_t O\right]\,, \end{aligned} \tag{2.51}$$

where in the final line, we used the cyclic property of the trace and evaluated the expression at zero external field.

Finally, we use the Woodbury matrix identity to obtain

$$\langle O(t)\rangle^{\phi} = \mathrm{Tr}\left[\left(\mathrm{I}-\left(\mathrm{I}+B_{t+1}\dots B_{N_t}B_1\dots B_t\right)^{-1}\right)O\right]\,. \tag{2.52}$$

To identify the elements of the reduced density matrix, we set the operator matrix to contain only a single element, yielding

$$\langle\hat{\psi}_i^{\dagger}(t)\hat{\psi}_j(t)\rangle^{\phi} = n_{ij}^{\phi}(t) = \delta_{ij}+\left[\left(\mathrm{I}+B_{t+1}\dots B_{N_t}B_1\dots B_t\right)^{-1}\right]_{ji}\,. \tag{2.53}$$

The second term corresponds to the equal-time single-particle Green's function, given by the inverse of the Fermi matrix as

$$G_{ij}^{\phi}(t) = \left[\left(\mathrm{I}+B_{t+1}\dots B_{N_t}B_1\dots B_t\right)^{-1}\right]_{ij}\,. \tag{2.54}$$

The single-particle operators considered here do not depend on the specific time slice when averaged, a property that arises from the time-translation invariance of the bosonic fields. This invariance is reflected in the cyclic symmetry of the weight, as can be observed in (2.46). We can use this property when computing observables by averaging over the time slices, enhancing the statistical accuracy depending on the observable in question.

Many-body observables

So far, our focus has been on bilinear operators. However, higher-order observables, such as density-density correlators, are also of significant physical interest. Fortunately, the Hubbard-Stratonovich allows us to understand the systems as non-interacting fermions in the presence of an external field. This reinterpretation enables us to apply Wicks theorem, which decouples the higher-order operators into products of bilinear ones.

Conveniently, correlations between the different spin species are also fully decoupled and can be written as products like

$$\begin{aligned}\langle\psi_{i,\uparrow}^{\dagger}\psi_{j,\uparrow}\psi_{k,\downarrow}^{\dagger}\psi_{l,\downarrow}\rangle^{\phi} &= \langle\psi_{i,\uparrow}^{\dagger}\psi_{j,\uparrow}\rangle^{\phi}\langle\psi_{k,\downarrow}^{\dagger}\psi_{l,\downarrow}\rangle^{\phi}\\ &= n_{ij,\uparrow}^{\phi}n_{kl,\downarrow}^{\phi}\,,\end{aligned} \tag{2.55}$$

where we have assumed equal times. Generally, if operators involving different spin species are considered, the expectation values factorize, allowing each spin sector to be computed independently.

When considering higher-order contractions involving the same spin species, additional terms arise. For example, in the case of four-fermion interactions within the same spin species, the full density correlation is given by

$$\langle \psi^\dagger_{i,\uparrow}\psi_{j,\uparrow}\psi^\dagger_{k,\uparrow}\psi_{l,\uparrow}\rangle^\phi = n^\phi_{ij,\uparrow}n^\phi_{kl,\uparrow} - n^\phi_{il,\uparrow}n^\phi_{kj,\uparrow}\,. \tag{2.56}$$

Note that permutations of the field operators are easily accounted for by simply using the commutation rules.

For a generic number of fermionic operators, the contractions can be performed using the determinant formula

$$\langle \psi^\dagger_{i_0,\uparrow}\psi_{j_0,\uparrow}\dots\psi^\dagger_{i_n,\uparrow}\psi_{j_n,\uparrow}\rangle^\phi = \det\left[G^\phi_{i_\alpha j_\beta}\right], \tag{2.57}$$

where $k,l = 1,\dots,n$ and $G^\phi_{i_\alpha j_\beta}$ is a matrix with $n\times n$ elements constructed from the full equal-time greens function. A detailed proof of this relation can be found in the Appendix of [58].

2.4 Canonical computations

In the previous sections, we discussed the computation of observables in the grand canonical ensemble, where the chemical potential is fixed, but the particle number is allowed to fluctuate. However, in many practical scenarios the canonical ensemble, with a set particle number, is more relevant. To achieve this, we can use projective techniques, which we will discuss in the following sections, for both ground state and finite temperature calculations.

2.4.1 Ground state

To study systems at zero temperature, the general lattice approach described above must be refined to eliminate contributions from excited states. This is achieved by realizing that the ground state wave function can be obtained from the infinite time limit of the time evolution operator as

$$|\Psi_0\rangle \propto \lim_{N_t\to\infty}\prod_{N_t} e^{-\Delta t\hat{H}}|\Psi_T\rangle\,. \tag{2.58}$$

Here, $|\Psi_0\rangle$ is the ground state, and $|\Psi_T\rangle$ is some trial state that has a non-zero overlap with the ground state. In practice, the limit cannot be taken exactly, and one needs to monitor the convergence of the ground state energy with respect to the number of time slices. Within this framework, the ground state partition function can be expressed as

$$\begin{aligned} Z &= \langle\Psi_0|\Psi_0\rangle \\ &= \lim_{N_t\to\infty}\langle\Psi_T|\prod_{N_t} e^{-\Delta t\hat{H}}|\Psi_T\rangle\,, \end{aligned} \tag{2.59}$$

and observables can be computed analogously via

$$\langle O\rangle = \frac{\langle\Psi_T|\prod_{N_t/2} e^{-\Delta t\hat{H}}\hat{O}\prod_{N_t/2} e^{-\Delta t\hat{H}}|\Psi_T\rangle}{\langle\Psi_T|\prod_{N_t} e^{-\Delta t\hat{H}}|\Psi_T\rangle}\,, \tag{2.60}$$

where we use a symmetric splitting of the time evolution operator. Following the procedure discussed in Section 2.3.2, various auxiliary field transformations can now be applied to decouple the interaction.

One of the most attractive features of the ground state approach is the reduced size of the matrices involved in practical computations. The trail states can be written as Slater determinants composed of single-particle basis states. In matrix representation, they take the form of $N_s \times N$ matrices, where N is the number of particles, and $N_s = V$ is the total number of single-particle states permitted by the lattice, which is set by the number of sites. This means that the matrix-matrix operations required to compute the weight become $\mathcal{O}(N_s^2 N)$ compared to the original $\mathcal{O}(N_s^3)$, reducing computational cost significantly. Similar to the finite temperature case, the weight can be computed via a determinant formula and is given by the overlap of Slater determinants as

$$\langle \Psi | \Psi' \rangle = \det \left[S_\psi S'_\psi \right], \tag{2.61}$$

where S_ψ and $S_{\psi'}$ are the Slater determinant matrices.

The efficiency of the ground state approach strongly depends on the choice of trial states. A good initial overlap between the actual ground state and the trial state is desirable, as it reduces the number of time slices required to reach convergence. On a rectangular lattice, plane waves provide a natural and commonly used choice for the single-particle basis. However, in systems with strong correlations or inhomogeneities, other trial states have been used, such as projected BCS wavefunctions [59] or harmonic oscillator orbitals [60, 61].

2.4.2 Finite temperature

Computing the canonical partition function at finite temperatures is more challenging, but various methods to do so exist. For example, we may expand the partition function on a field configuration in the fugacity $z = e^{\beta\mu}$ as

$$\begin{aligned} Z^\phi &= \det\left[\mathrm{I} + zU\right] = \sum_{n=0}^{N_s} z^n Z_n^\phi \\ &= \exp\left[\mathrm{Tr}\log\left(\mathrm{I} + zU\right)\right] \\ &= \exp\left[\sum_{n=1}^{N_s} \frac{(-1)^{n+1}}{n} z^n \,\mathrm{Tr}\, U^n\right], \end{aligned} \tag{2.62}$$

where Z_N^ϕ is the canonical trace for a fixed particle number N. This technique, known as the activity expansion, has been applied in contexts such as the nuclear shell model [62]. However, there are several drawbacks to the approach, including the need to compute high powers of the matrix U, which can be numerically expensive. Moreover, substantial cancellations and large condition numbers can lead to problems with numerical precision. Although diagonalizing the matrix can mitigate these problems, we instead turn to the projective approach to canonical computations first introduced in [63].

Starting from the grand canonical partition function, we introduce a phase factor $\varphi_n = \frac{2\pi n}{N_s}$ to the weight, resulting in

$$\det\left[\mathrm{I} + e^{i\varphi_n} e^{\beta\mu} U\right] = \sum_{m=0}^{N_s} e^{im\varphi_n} e^{m\beta\mu} Z_m^\phi. \tag{2.63}$$

To isolate the canonical component for a specific particle number N, we leverage the orthogonality of the phase factors and project out the undesired terms, yielding

$$Z_N^\phi = \frac{1}{N_s} \sum_{n=1}^{N_s} e^{-iN\varphi_n} e^{-N\beta\mu} \det\left[\mathrm{I} + e^{i\varphi_n} e^{\beta\mu} U\right], \tag{2.64}$$

where the number of quadrature points is fixed to match the total number of single-particle states N_s. This expression requires the evaluation of N_s determinants, which can be prohibitively expensive but may be avoided via diagonalization [64, 65]. It is important to note that the chemical potential here does not have a physical meaning, but serves as a stabilizing parameter. It is tuned to maximize the overlap between the canonical and grand canonical partition functions, either once for the whole run or on a per-configuration basis.

In a two-component Fermi gas, separate projections must be performed for each spin species, which is easily possible due to the decoupling of the respective contributions to the full weight. If we want to compute observables, they will also require a projection. Working with a generic bilinear operator $\hat{O} = \hat{\psi}_i^\dagger O_{ij} \hat{\psi}_j$, the expectation value on a single field configuration becomes

$$\langle O \rangle_N^\phi = \frac{e^{-N\beta\mu}}{Z_N^\phi N_s} \sum_{n=1}^{N_s} e^{-iN\varphi_n} \det\left[\mathrm{I} + e^{i\varphi_n} e^{\beta\mu} U\right] \mathrm{Tr}\left[\left(\mathrm{I} - \left(\mathrm{I} + e^{i\varphi_n} e^{\beta\mu} U\right)^{-1}\right) O\right]. \tag{2.65}$$

Higher-order observables may be computed using Wick's theorem in analogy to the grand canonical case.

There are different strategies for sampling the canonical weight. One approach involves making projections at each step of the Markov chain, while another samples the grand canonical ensemble at a fixed chemical potential and reweights the configurations to the canonical weight. The latter method requires a good guess for the chemical potential to be known a priori. Alternatively, promising results for particle projection have been achieved using the complex Langevin method in [66], where small three-dimensional lattices were considered. This approach leaves the Fourier sum continuous and samples it along with the auxiliary field in the complex plane.

Whether a sign problem is present in canonical computation depends on the specifics of the auxiliary field transformation, as well as the interaction and spin or mass balance, similar to the grand canonical case discussed Table 2.1.

While direct sampling of the canonical weight is appealing, this work will focus on the reweighting approach. It allows the use of the same configurations for grand canonical and canonical computations, facilitating a direct comparison between the two ensembles.

Chapter 3

Stochastic quantization

Stochastic quantization originated in the late 60s in the context of quantum mechanics [67] and underwent significant further development in the 80s [68, 69]. It is an approach in which Euclidean quantum field theory and quantum mechanics emerge from a stochastic process, which can either be treated perturbatively or non-perturbatively via lattice computations.

This chapter introduces the Langevin method for real-valued actions, a sampling algorithm based on stochastic quantization, in Section 3.1, after which we discuss the extension to complex weights in Section 3.2. Complex Langevin is a powerful tool for systems with sign problems and can allow for the computation of observables that are Section 3.2.2. Lastly, Section 3.3 presents an investigation of Lee-Yang zeros in scalar field theories based on complex Langevin.

3.1 Langevin sampling

The central idea behind stochastic quantization is to replace the conventional path integral with a stochastic evolution of the field's degrees of freedom, governed by the Langevin equation

$$\frac{\partial \phi}{\partial \tau} = -\frac{\delta S}{\delta \phi} + \eta(\tau), \tag{3.1}$$

where $\eta(\tau)$ is Gaussian white noise with zero mean and variance $\langle \eta(\tau)\eta(\tau')\rangle = 2\delta(\tau - \tau')$ and τ is a non-physical stochastic time. To grasp the intuitive picture, we can draw an analogy to classical Brownian motion. The field can be envisioned as a particle moving in a potential $V(\phi) = -S(\phi)$, with the random noise acting as thermal fluctuations. In the absence of noise, the field would follow the path of steepest descent, ultimately leading to the classical equations of motion for the theory in question. Consequently, the noise introduces quantum fluctuations into the system.

Observables of interest are computed as averages over stochastic time. They agree with the path integral formulation in the infinite time limit, as long as the stochastic process is ergodic, so we have

$$\langle \mathcal{O} \rangle = \lim_{T\to\infty} \frac{1}{T} \int_0^T d\tau \mathcal{O}[\phi(\tau)] = \frac{1}{Z} \int \mathcal{D}\phi \mathcal{O}[\phi] e^{-S[\phi]}. \tag{3.2}$$

This equivalence is guaranteed by the correspondence between the Langevin evolution and a Fokker-Planck equation, which possesses a unique equilibrium distribution given by $P(\phi) = e^{-S(\phi)}$.

3.1.1 Formal justification

We now briefly sketch the correspondence between the Langevin and Fokker-Planck equations. To this end, we consider the time evolution of observables $O(x)$ in both formalisms. Writing Equation (3.1) for a lattice system as a stochastic differential equation of Ito-type, we have

$$d\phi_i = -\frac{\partial S}{\partial \phi_i} d\tau + dw_i \,. \tag{3.3}$$

Here, dw_i is a Wiener measure, for which the expectation values are as $\langle dw_i \rangle = 0$ and $\langle dw_i dw_j \rangle = 2\delta_{ij} dt$ at equal time. These expectation values vanish for unequal times, and the Wiener process scales as $\sqrt{d\tau}$.

To explore the time evolution of a field configuration $\phi(\tau)$, we expand the observable $O(\phi)$ to first order, yielding

$$\begin{aligned} O\big(\phi(\tau+d\tau)\big) =& O\big(\phi(\tau)\big) + \sum_i \frac{\partial O}{\partial \phi_i} d\phi_i + \frac{1}{2}\sum_{i,j} \frac{\partial^2 O}{\partial \phi_i \partial \phi_j} d\phi_i d\phi_j + \mathcal{O}(d\tau^{3/2}) \,. \\ =& O\big(\phi(\tau)\big) + \sum_i \frac{\partial O}{\partial \phi_i} \left(-\frac{\partial S}{\partial \phi_i} d\tau + dw_i \right) \\ &+ \frac{1}{2}\sum_{i,j} \frac{\partial^2 O}{\partial \phi_i \partial \phi_j} dw_i dw_j + \mathcal{O}(d\tau^{3/2}) \,. \end{aligned} \tag{3.4}$$

Terms of order dt^2 are dropped in the second line. This expression allows us to determine the time evolution of the expectation value, given an initial distribution $P(\phi)$ as

$$\begin{aligned} \frac{d}{dt}\big\langle O\big(\phi(\tau)\big)\big\rangle &= \frac{1}{dt}\big\langle O\big(\phi(\tau+d\tau)\big) - O\big(\phi(\tau)\big)\big\rangle \\ &= \langle \sum_i -\frac{\partial O}{\partial \phi_i}\frac{\partial S}{\partial \phi_i} + \frac{\partial^2 O}{\partial \phi_i^2} \rangle \\ &= \int \mathcal{D}\phi \left(\sum_i -\frac{\partial O}{\partial \phi_i}\frac{\partial S}{\partial \phi_i} + \frac{\partial^2 O}{\partial \phi_i^2} \right) P(\phi) \\ &= \int \mathcal{D}\phi\, O(\phi) \left(\sum_i \frac{\partial}{\partial \phi_i}\frac{\partial S}{\partial \phi_i} + \frac{\partial^2}{\partial \phi_i^2} \right) P(\phi) \,. \end{aligned} \tag{3.5}$$

In the final step, we used integration by parts, which requires sufficiently fast decay of the distribution towards infinity. We are now left with a Fokker-Plank equation for the time evolution of the distribution itself instead of the fields

$$\partial_\tau P(\phi) = \left(\sum_i \frac{\partial}{\partial \phi_i}\frac{\partial S}{\partial \phi_i} + \frac{\partial^2}{\partial \phi_i^2} \right) P(\phi) \,. \tag{3.6}$$

The equilibrium solution to this equation is $P(\phi) = e^{-S(\phi)}$, demonstrating that the Langevin equation naturally leads to the correct equilibrium distribution and ensuring that the expectation values computed in the infinite-time limit are accurate. For a more detailed proof, refer to [70].

3.1.2 Discretization & Technical aspects

The discussion thus far has centered on a continuous stochastic time t, an idealization that is not practical in numerical simulations. Instead, we need to discretize the time evolution of the fields, which can be done in several ways. As with ordinary differential equations, a variety of integration schemes are available, each varying in terms of accuracy and stability.

The most straightforward method is the Euler-Maruyama scheme, which can be seen as the equivalent of the Euler method and is given by

$$\phi_{i,t+1} = \phi_{i,t} - \frac{\partial S}{\partial \phi_i}\epsilon_L + \sqrt{2\epsilon_L}\eta_i\,, \tag{3.7}$$

where ϵ_L is the discretized time step and η_i is a Gaussian random number with zero mean and unit variance. While this scheme is only first-order accurate, it is simple to implement and computationally inexpensive. The discretization maintains stability if the time step ϵ_L is sufficiently small and the action is bounded. However, observables computed in this scheme will carry an error that has to be removed by an extrapolation $\epsilon_L \to 0$ to continuous stochastic time.

Alternative integration techniques have been developed and employed in lattice computations. These include second-order Runge-Kutta schemes [71] and implicit methods [72], which offer improved accuracy and stability under certain conditions.

3.2 Complex Langevin

The Langevin approach described above is well suited for cases where the action is real and bounded from below, avoiding sign problems and ensuring a well-defined equilibrium distribution. However, many physical systems do not allow for a description via a real action, and complex parts appear. In such cases, the Langevin approach can be extended via complexification so that sampling remains possible. This was already noticed in the 80s [73, 74], but interest waned due to significant numerical and convergence challenges encountered at the time [75, 76].

Recently, there has been a resurgence of interest in the method, with various applications in the context of QCD [77–80] and beyond, in particular in the realm of ultracold many-body quantum systems [32, 50, 81–85]. These contemporary investigations also introduced partial solutions to previously identified issues with the method [86–89].

3.2.1 Complexification

To extend the Langevin formalism to cases involving complex-valued actions, we allow the field variables to become complex $\phi_i = \phi_i^R + i\phi_i^I$. Here, ϕ_i^R and ϕ_i^I are the real and imaginary components. This leads the Langevin equation to change domains from $\mathbb{R}^n$ to $\mathbb{C}^n \equiv \mathbb{R}^{2N}$. The resulting system of coupled equations is given by

$$\frac{\partial \phi_i^R}{\partial \tau} = -\,\mathrm{Re}\left(\frac{\partial S}{\partial \phi_i}\right) + \eta_i(\tau)\,, \tag{3.8}$$

$$\frac{\partial \phi_i^I}{\partial \tau} = -\,\mathrm{Im}\left(\frac{\partial S}{\partial \phi_i}\right), \tag{3.9}$$

where the derivative of the action with respect to the complex field appears. For the equations to be well-defined, the action must be holomorphic (or at least meromorphic) to avoid ambiguities in the differentiation process. In Equation (3.9), the noise is chosen to be purely real. While noise with a complex part is generally possible under a suitable normalization condition, it has been found to hinder numerical stability [90, 91] and is therefore typically avoided.

The discretization of the complex Langevin equation follows the same procedure as in the real-valued case, using the Euler-Maruyama method. However, numerical stability is often an issue when using the naive discretization scheme with a fixed step size. Instabilities usually arise from repulsive directions in the complexified space, where large drift values can cause the stochastic process to diverge. An adaptive stepping scheme was introduced in [86] to avoid these instabilities. This method adjusts the step size dynamically based on the magnitude of the drift term, ensuring that the product of the step size and drift magnitude remains constant. The approach allows for smaller steps in regions with large drift and larger steps where the drift is small, thus improving numerical stability.

Sampling in the complex Langevin framework proceeds similarly to the real case, but now occurs in the complexified field space. Unlike conventional sampling methods, the exact form of the probability distribution we sample is not generally known. However, under certain conditions, the complex Langevin process has been shown to converge to the desired distribution, as will be discussed in the next section. We then have a direct correspondence of observables

$$\int \prod_i d\phi_i \, O(\phi) \, e^{-S(\phi)} \leftrightarrow \int \prod_i d\phi_i^R d\phi_i^I \, O\left(\phi^R + i\phi^I\right) P\left(\phi^R, \phi^I\right), \tag{3.10}$$

allowing for computations even in the presence of otherwise prohibitive sign problems.

3.2.2 Criteria for correctness

Complex Langevin has been applied in various contexts, perhaps most notably in QCD, for which it was initially proposed, as well as in systems involving non-relativistic fermions [32, 50, 81]. Unfortunately, the method is not guaranteed to converge to the correct result. Indeed, there are now two separate time evolutions: one for the real-valued weight with complex domain $P(\phi^R, \phi^I)$ and one for the complex weight with real domain $\rho(\phi)$. The former is given by

$$\partial_\tau P(\phi^R, \phi^I) = L^T P(\phi^R, \phi^I) \tag{3.11}$$

$$= \sum_i \frac{\partial}{\partial \phi_i^R} \left(\frac{\partial}{\partial \phi_i^R} + \mathrm{Re}\left(\frac{\partial S}{\partial \phi_i} \right) \right) + \frac{\partial}{\partial \phi^I} \mathrm{Im}\left(\frac{\partial S}{\partial \phi_i} \right), \tag{3.12}$$

where L^T is the real Fokker-Planck operator acting on the complexified field space. The second evolution equation is given by

$$\partial_\tau \rho\phi = L_c^T \rho(\phi) = \sum_i \frac{\partial}{\partial \phi_i} \left(\frac{\partial}{\partial \phi_i} + \frac{\partial S}{\partial \phi_i} \right) \rho(\phi), \tag{3.13}$$

which resembles the standard Fokker-Planck operator but is extended to a complex action. Here, L_c^T is the complex Fokker-Plank operator, extended to act on a complex-valued function of real variables.

Extensive research has been conducted on the convergence properties of the complex Langevin method. Perhaps most notably, there is proof of convergence that relates the time evolution of observables under both of the operators mentioned above [87]. The proof relies on several partial integrations, which require a sufficiently fast decay of the equilibrium distribution towards infinity. However, this condition is not always met, and in some cases, the complex Langevin process may converge to incorrect results.

Various criteria have been proposed to identify situations where the method might fail. An easily used approach is the so-called "drift criterion", for which one monitors the magnitude of the drift term and its decay towards large values [88]. In practice, one computes

$$\max_i \left(\frac{\partial S}{\partial \phi_i} \right), \tag{3.14}$$

for each field configuration and bins the values into a histogram. If the distribution decays at least exponentially, the complex Langevin process is expected to converge correctly. However, slower than exponential decay does not necessarily imply a failure of the complex Langevin process, as it may still converge to the correct result.

Additionally, methods for correcting wrong convergence have been proposed, involving the computations of so-called "boundary-terms" [92]. These are related to the difference between the complex and real Fokker-Plank operators and can be computed order-by-order under some assumptions to an interpolating function between both evolutions. In terms of observables that can be computed from the stationary complex Langevin process, they are given as [93]

$$B_n = \lim_{Y \to \infty} \int_{|\phi| < Y} d\phi \, P(\phi^R, \phi^I) \, L_c^n \, \mathcal{O}(\phi) \, . \tag{3.15}$$

Here, a cutoff in field space is introduced to ensure the convergence of the integral. It must be carefully removed by searching for a plateau in the observable as a function of the cutoff. Note that a boundary term computed like this is not a generic feature of the complex Langevin process but depends on the observable under study. There may be observables that converge to the exact result, while others do not, even when evaluating the same Langevin trajectory.

Nonetheless, caution is warranted, as the complex Fokker-Planck operator is not guaranteed to have a non-degenerate zero mode or a non-positive spectrum. Recent studies have explored such cases [94], indicating that these issues can significantly affect the results.

A similar problem may arise in the presence of singularities in the complexified action. They are often encountered in fermionic theories, where zeros in the fermionic determinant can lead to diverging actions and singularities in the drift. The impact of singularities depends on their location relative to the sampled distribution. If a pole lies entirely outside the sampled region, it typically does not cause problems. In contrast, if the pole lies inside the sampled area, correctness depends on the decay of the distribution towards the pole, which is a similar requirement to the one for large field values. Unfortunately, the correction procedures described earlier are ineffective in cases where singularities lead to incorrect convergence.

3.3 Evaluating complex Langevin for the search of Lee-Yang zeros

This section is based on [2].

In this section, we will discuss the potential of complex Langevin in the search of Lee-Yang zeros based on simple toy models and full-dimensional lattice computations. A secondary aim is to provide a simple, practical introduction to complex Langevin based on these toy models.

The Yang-Lee theorem [95, 96] states that O(N) models have branch cuts for purely imaginary values of h_0, and the cuts end at second-order critical points, known as Lee-Yang edge singularities. These singularities provide an upper bound on the radius of convergence of a Taylor expansion in h_0 around the origin.

O(N) models are of great interest. Indeed, around the chiral phase transition the matter dynamics of QCD can be well-approximated by a theory of massless pions and scalar quark condensate. In this low-energy regime, nucleons and heavier hadrons play no role, and the model is described by an $O(4)$ scalar field with quartic self-interactions. In this model, an external source h_0 coupled to the quark condensate plays the role of chemical potential. The breaking of chiral symmetry, signaled by a vacuum expectation value of the quark condensate, is controlled by the sign and strength of the mass parameter.

Thus, the expansion in h_0, would correspond to an expansion in the chemical potential μ around the chiral phase transition in QCD. Knowing the location of the edge singularity in the O(N) case provides information on how far in the chemical potential a Taylor expansion in QCD can be trusted. An earlier investigation on this topic, using random matrix models, can be found in [97]. Recent works on this topic include Taylor expansions [98], lattice QCD results for non-universal parameters [99], reweighting [100], Functional Renormalisation Group methods [101], Padé resummations [102], as well as studies in QCD via imaginary chemical potential [103].

3.3.1 One-site toy model

We start by considering a zero-dimensional, single-site toy model with the action defined as

$$S(\phi) = \frac{m^2}{2}\phi^2 + \frac{\lambda}{4!}\phi^4 + i h_I \phi, \tag{3.16}$$

where m and λ are real and positive. The external field h_I is also real, so the coupling to ϕ is purely imaginary. The partition function is a simple integral given by

$$Z(h) = \int d\phi\, e^{-S(\phi)} = \int d\phi\, e^{-m^2/2\phi^2 - \lambda/4!\phi^4 - i h_I \phi}, \tag{3.17}$$

which takes the form of a Fourier transform. In the following, we set $m = 0$ and $\lambda = 1$, noting that the parameters do not significantly influence qualitative features as long as both are non-negative. The Langevin evolution is then given by

$$\partial_\tau \phi^R = -\frac{1}{3!}\,\mathrm{Re}\left(\phi^3\right) + \eta(t), \tag{3.18}$$

$$\partial_\tau \phi^I = -\frac{1}{3!}\,\mathrm{Im}\left(\phi^3\right) - h_I, \tag{3.19}$$

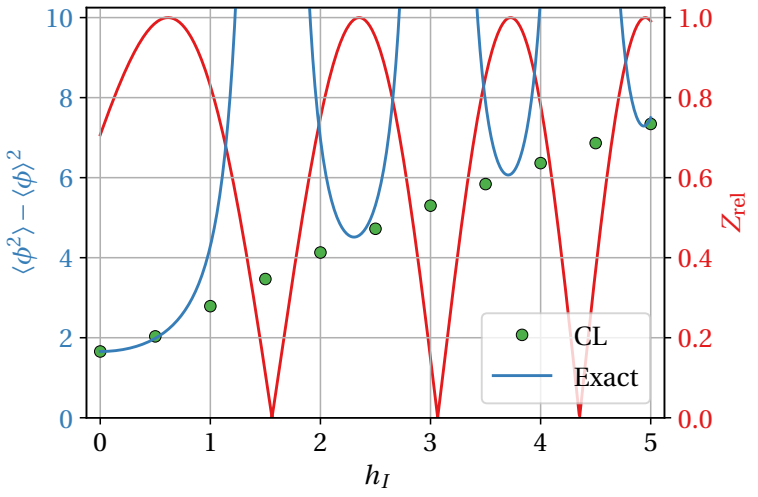

Figure 3.1: Susceptibility for the one-site toy model from complex Langevin (green points) compared to the exact result (blue lines) as a function of the external field. The statistical errors are smaller than the symbol size in all cases. The red line shows the relative weight between the two equally contributing thimbles.

which we discretize in the Euler-Maruyama scheme. We employ an adaptive step size with a reference step size of $\epsilon_L = 0.05$ and compute the susceptibility for various external field values. Since a simple integral solves the model, we can compare the results of the complex Langevin process to the exact solution. The susceptibility is expected to diverge at the positions of the Lee-Yang zeros, where the partition function vanishes.

We present results for the imaginary part of the susceptibility in Figure 3.1. While complex Langevin agrees well with the exact value for external field $h \leq 0.5$, it begins to deviate heavily as the external field increases. Notably, the complex Langevin approach fails to capture the divergence at the Lee-Yang zeros. Indeed, it is unclear how the complex Langevin process could capture such a feature in the first place, as a large imaginary magnetization would require the Langevin distribution to be localized far from the origin in the complex plane.

This model was previously examined by Salcedo [104], who pointed out an interesting pathology: Since the imaginary part of the drift term is always negative for any real-valued field, the stochastic process cannot cross the real axis from below, making the sampling non-ergodic. This is a clear violation of the conditions for the convergence of complex Langevin and can be seen as a reason for the failure of the method. Although this argument is compelling, it can be countered by introducing a small imaginary component to the noise, allowing for random jumps in the imaginary direction. However, this modification does not resolve the issue of incorrect convergence.

Finally, we note that there appears to be a point where the exact and complex Langevin results coincide, which will be discussed further in the following section.

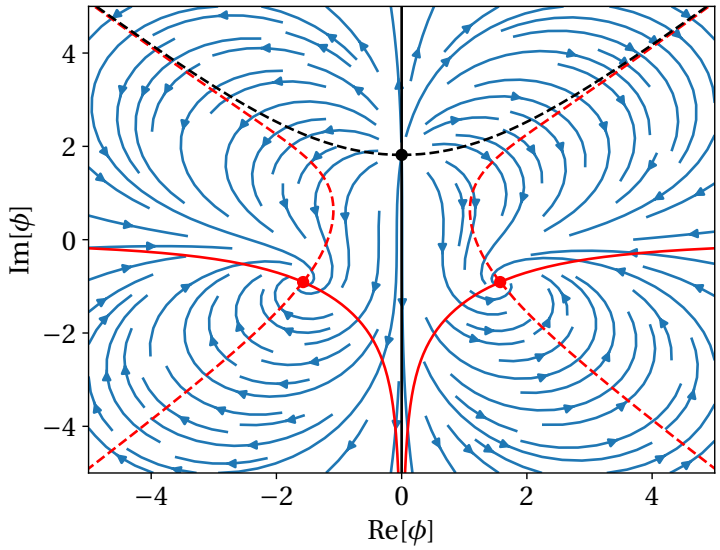

Figure 3.2: Drift and thimble structure of the quartic one-site model (3.16). The Langevin drift is shown in blue. There are two attractive fixed point, which are associated with two contributing thimbles (solid red lines). The fixed point of the single non-contributing thimble is repulsive (black). The dashed lines are the corresponding anti-thinbles.

3.3.2 Lefschetz Thimbles - a short excursion

To better understand the failure of the complex Langevin approach in the simple one-site model, this section briefly introduces Lefschetz Thimbles [105, 106] and their use for the model in question. A Lefschetz thimble is a manifold in complex field space, defined as the set of points where the imaginary part of the action remains constant, starting from a fixed point in the drift. Formally, the thimble is described by the differential equation

$$\frac{\partial \phi}{\partial \theta} = \overline{\frac{\partial S}{\partial \phi}}, \tag{3.20}$$

where the path flows into a fixed point z_σ of the action in the infinite time limit. Analogously, we define the anti-thimble as the set of lines with

$$\frac{\partial \phi}{\partial \theta} = -\overline{\frac{\partial S}{\partial \phi}}, \tag{3.21}$$

ending in the same fixed point. There exists a mapping of the original integral to an integral over the various available thimbles, given by

$$Z = \sum_\sigma n_\sigma \int_{D_\sigma} d\phi\, e^{-S(\phi)} = \sum_\sigma n_\sigma e^{-\mathrm{Im}[S(z_\sigma)]} \int_{D_\sigma} d\phi\, e^{-\mathrm{Re}[S(\phi)]}. \tag{3.22}$$

Here, D_σ represents the domain of the thimble, and n_σ is the intersection number of the anti-thimble with the original integration domain. The latter governs whether a specific thimble contributes to the original integral. The second equality follows directly from (3.20), as the imaginary

part of the action remains constant along the thimble. It is important to note that while the thimble decomposition addresses part of the sign problem, two sources of a sign problem persist.

First, the Jacobian induced by the change of variables necessary to sample the individual thimbles is not guaranteed to have a real and positive determinant. This leads to a residual sign problem, even when sampling a single thimble. We call this the local residual sign problem. It is quite mild in most practical applications and does not typically pose a significant problem.

Secondly, a sign problem can arise if multiple thimbles are contributing with a similar magnitude but different phases $\exp(-\mathrm{Im}[S(z_\sigma)])$, leading to cancellations of contributions between thimbles. We call this the global residual sign problem.

This global sign problem is observed in the model (3.16), where the thimble structure is illustrated in Figure 3.2. The symmetry $\mathrm{Re}\,\phi \to -\mathrm{Re}\,\phi$ of the action is reflected in the drift and thimble structure. Two contributing thimbles are associated with attractive fixed points in the drift, while the sole repulsive fixed point does not contribute. These contributing thimbles are related via symmetry and must carry the same absolute magnetization. In order to find the vanishing weight required for Lee-Yang zeros, the contributing thimbles must have opposite phases. As the absolute value of the partition functions is the same, the relative weight is just given by

$$Z_{\mathrm{rel}} = \frac{|Z_1 + Z_2|}{|Z_1| + |Z_2|}, \tag{3.23}$$

where Z_i are the parts of the partition function associated with the individual thimbles. We show results for the relative weight in Figure 3.1. The points where Z_{rel} goes to zero also mark the zeros in the full partition function and the divergence of the magnetization. Interestingly, the points where $Z_{\mathrm{rel}} = 1$ coincide perfectly with the values for which the exact and complex Langevin results agree. This suggests that the complex Langevin method fails to capture the relative phases between thimbles correctly, effectively sampling a system where the global sign between thimbles is quenched.

Finally, it is worth noting that this same model was recently studied in the context of boundary terms in [94], where the authors identified positive real eigenvalues in the spectrum of the complex Fokker-Planck operator, which hinders convergence.

3.3.3 O(2) site model

Since the lattice model we are interested in is a $O(4)$ model, we modify our toy to include an additional field, making it $O(2)$ symmetric. The action is then given by

$$S(\phi_1, \phi_2) = \frac{m^2}{2}\left(\phi_1^2 + \phi_2^2\right) + \frac{\lambda}{4}\left(\phi_1^2 + \phi_2^2\right)^2 + i h_I \phi_1, \tag{3.24}$$

where the h_I is coupled to the first field only. The partition function can be cast into the form of an integral over the radial coordinate as

$$Z(h) = 2\pi \int_0^\infty dr\, r\, I_0(-ihr) \exp\left[-\frac{m^2 r^2}{2} - \frac{\lambda r^4}{4}\right], \tag{3.25}$$

where $I_n(x)$ denotes the modified Bessel functions of the first kind. The magnetization can then be obtained via a derivative of the partition function with respect to the external field,

$$M = \langle x \rangle = -\frac{\partial Z}{\partial i h} = \frac{2\pi}{Z} \int_0^\infty dr\, r^2\, I_1(-ihr) \exp\left[-\frac{m^2 r^2}{2} - \frac{\lambda r^4}{4}\right]. \tag{3.26}$$

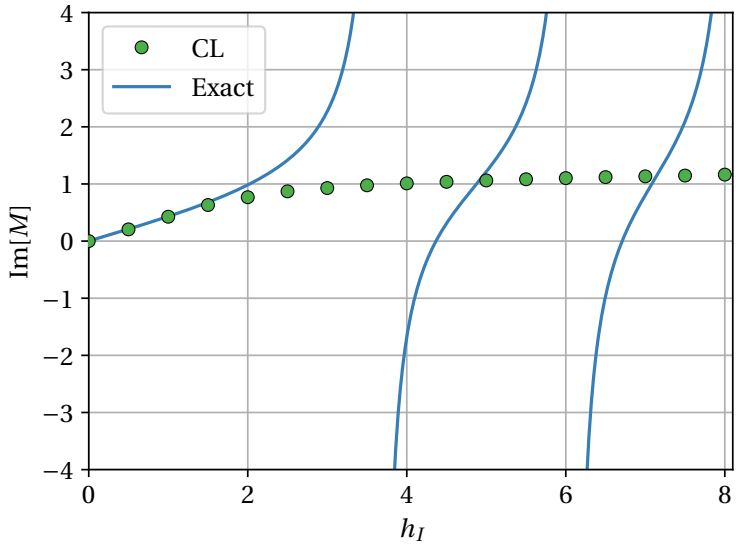

Figure 3.3: Imaginary part of the magnetization of the single site O(2) model as a function of h. The continuous lines represent the exact solutions, while the points resulted from complex Langevin simulations. The divergences correspond to zeros of the partition function, which the complex Langevin fails to capture.

To perform simulations, we again discretize the Langevin equation in the Euler-Maruyama scheme, and use an adaptive step size. As can be seen in Figure 3.3, the complex Langevin process fails to capture the divergence of the magnetization at the Lee-Yang singularities, analogous to the one-site model. Such behavior appears to be somewhat common with complex Langevin, as similar features have been observed in studies of Random Matrix theory [107] and models with singular Langevin drift [108].

3.3.4 Three-dimensional field theory

Moving on from the simple models with only the local potentials, we now consider the full three-dimensional $O(N)$ model with an imaginary external field. The action is given by

$$S = \int d^3x \left[\frac{1}{2}\partial_\mu \varphi_i \partial_\mu \varphi_i + \frac{m_0^2}{2}\varphi_i\varphi_i + \frac{\lambda_0}{4!}(\varphi_i\varphi_i)^2 + h_0\varphi_1 \right], \tag{3.27}$$

where $1 \leq i \leq N$ and the external field h_0 only couples to the first component. According to the Lee-Yang theorem, in the symmetric phase, the magnetic equation of state exhibits branch cuts terminating in edge singularities. Consequently, our focus lies on computing the average magnetization, given by

$$M = \frac{1}{V}\frac{\partial}{\partial h_0}\ln Z = \frac{1}{V}\langle \Phi_1 \rangle , \quad \Phi_1 = \int d^3x\, \varphi_1(x) \tag{3.28}$$

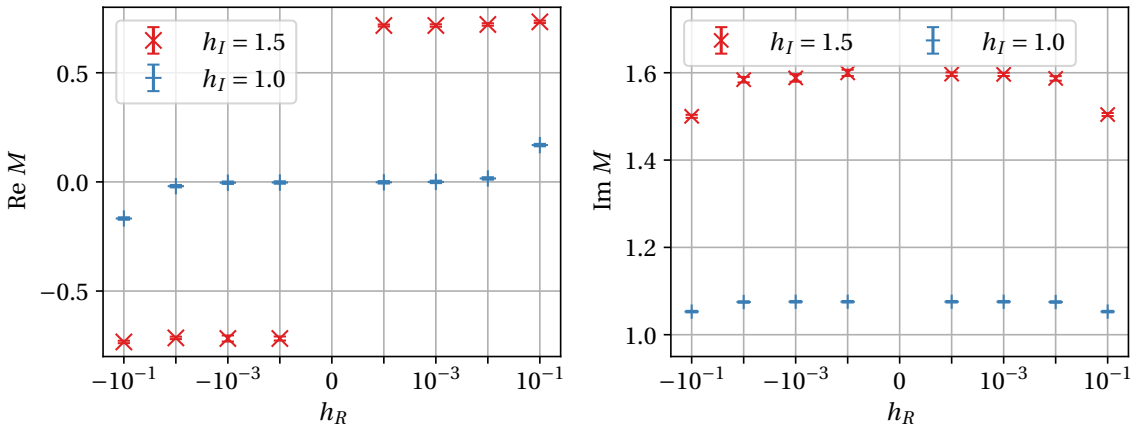

Figure 3.4: Real (left) and Imaginary (right) part of the magnetization of the three-dimensional $O(2)$-Model on a $V = 8^3$ lattice. Note that the horizontal axis is in units of powers of 10, except at 0. The real part shows the expected behavior: for $h_I < h_c$ the magnetization is continuous across $h_R = 0$, while for $h_I > h_c$ a discontinuity appears.

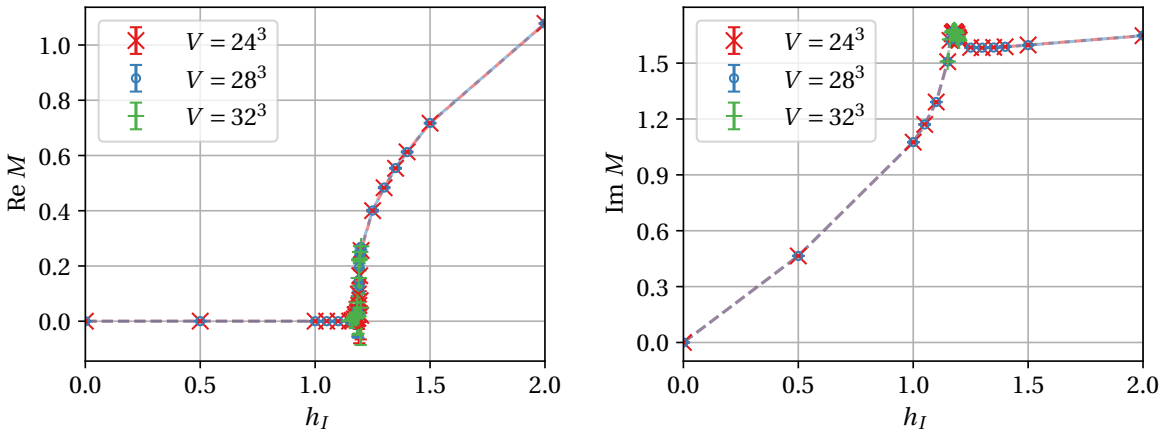

Figure 3.5: Real (right) and Imaginary (left) part of the average magnetization as a function of the imaginary part of the external magnetic field. Dashed lines connecting the data points are drawn to guide the eye. The results for the three different volumes collapse onto a single curve in the given scales. Notably, there is a kink in the real part of the magnetization at $h_I \approx 1.185$, while the imaginary part peaks around the same value.

where V is the total volume of the system. The per-site magnetic susceptibility can be computed from the partition function via a second derivative with respect to the external field,

$$\chi = \frac{1}{V}\frac{\partial^2}{\partial h_0^2}\ln Z = \frac{1}{V}\left[\langle \Phi_1^2\rangle - \langle \Phi_1\rangle^2\right]. \tag{3.29}$$

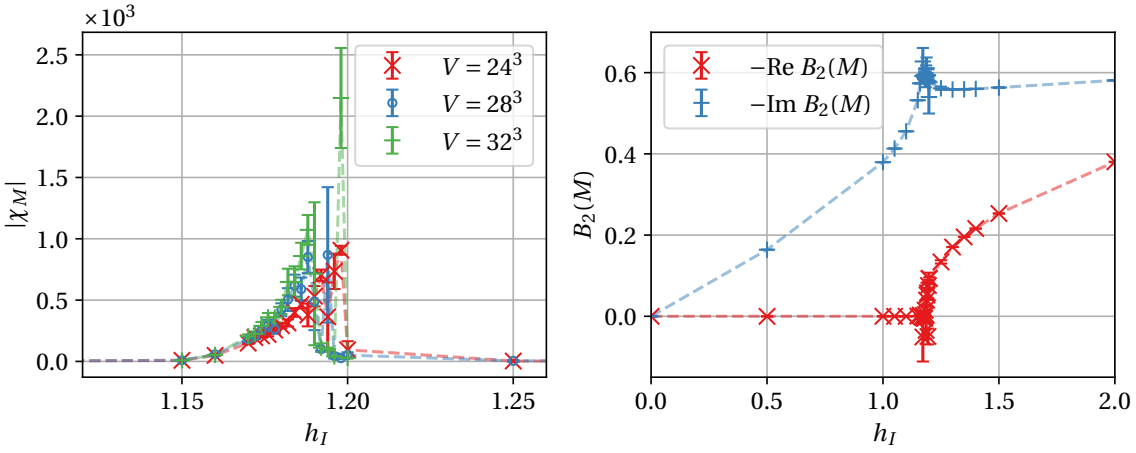

Figure 3.6: Left: Absolute value of the (complex) magnetic susceptibility as a function of the imaginary external field h_I for the three volumes considered. The peaks exhibit the typical behaviour of finite volume phase transitions of growing with the volume. Right: Negative real and imaginary part of the second boundary term associated with the magnetization, computed for $V = 28^3$. Different volumes show similar behavior. Note the similarity to the magnetization itself. The dashed lines connecting data points are drawn to guide the eye.

Splitting up the external field into a real and imaginary part, $h_0 = h_R + i h_I$, we anticipate a discontinuity in the magnetization when crossing the branch cut at constant h_R, whereas the susceptibility should peak at a critical value $|h_I| \simeq h_c$

Our simulations utilize the lattice parameters $m_0 = 1$, $\lambda_0 = 1$, and $N = 2$ for various external field values and total volume values. To verify the aforementioned (dis)continuous behavior, we conduct preliminary investigations at small volumes $V = 8^3$ and at two values of h_I, for which the results are shown in Figure 3.4. We find the expected continuous behavior in the $h_I < h_c$ regime, while a discontinuity appears for values $h_I > h_c$. For all computations, we introduce a slight offset $h_R = 10^{-4}$, similar to techniques employed in Ising model studies.

Results of a large-scale scan with Volumes $V = 24^3, 28^3$, and 32^3 can be found in Figure 3.5. The behavior resembles a second-order phase transition, as can be found, e.g., in the Ising model, where the magnetization is continuous but displays a kink and onset to a non-zero value at the critical point. The imaginary part of the magnetization is finite everywhere but at the origin and shows a cusp at $h_I \approx 1.185$, which is the exact location as the peak of the real part. This behavior is further investigated on the left plot of Figure 3.6, where we show the absolute value of the magnetic susceptibility. As is also typical for second-order phase transitions, the susceptibility shows a peak at the critical value, which grows with the system volume.

Additionally, we computed the first two boundary terms of the magnetization to validate the complex Langevin results. While we find the first order term $B_1(M)$ to vanish within statistical errors, the second term $B_2(M)$ takes on a finite value even at small external fields and is shown on the right side of Figure 3.6. This behavior is relatively unusual and makes the correction procedure proposed in [93] unusable, as it relies on the ratio B_1/B_2. The observed hierarchy of boundary

terms casts some doubt on the reliability of complex Langevin simulations for the given model.

3.3.5 Conclusion

We have studied the complex Langevin method in the search for Lee-Yang zeros for simple toy models and a full three-dimensional lattice model. Overall, the findings paint a mixed picture. The method failed to capture the divergence in both toy models, producing smooth curves instead. In the single-field toy model, we identified the source of this problem as a failure to capture the relative phases between thimbles correctly. Further investigation of the three-dimensional $O(2)$ model with complex Langevin has revealed somewhat different behavior, including a discontinuity and onset of the magnetization with an increasing imaginary external field reminiscent of a second-order phase transition. This is very promising for the search of Lee-Yang zeros. However, boundary terms are present, and thus, CL results must be interpreted with care.

In recent work [109–112], several technical advancements in the complex Langevin method have been made via the introduction of a kernel aided by Machine Learning. This has enabled the use of the method in real-time equilibrium simulations at far greater imaginary times than was previously possible and may also prove useful for the search for Lee-Yang zeros in the future.

Chapter 4

DQMC & Technical aspects

Since the introduction of Determinant Quantum Monte Carlo (DQMC) calculations for non-relativistic fermions, the Blankenbecler, Scalapino, and Sugar (BSS) algorithm [56] has been a workhorse for the field, and has seen a large range of developments and enhancements. In this chapter, we will discuss various technical aspects of the sampling algorithms used in this work, which are largely derived from BSS, and the specifics of their implementation. We start by detailing methods for constructing the Green's function in a stable and efficient way via Fourier acceleration, matrix decompositions, and wrapping. Subsequently, updating schemes are described, followed by a discussion of recent developments in the field of DQMC, which enable simulations of significantly larger systems through systematic model space truncation.

4.1 Constructing the Green's function

As already discussed in Section 2.3, the Green's function is a central quantity in the DQMC method and is directly related to observables and the weight of configurations. Therefore, handling and constructing the Green's function stably and efficiently is one of the most important aspects of the DQMC approach. In the following, we will work with the equal time Green's function $G^{\phi}(t) = \langle \psi(t)\psi^{\dagger}(t)\rangle$ given by

$$G^{\phi}(t) = \left(\mathrm{I} + B_{t+1}\dots B_{N_t}B_1\dots B_t\right)^{-1}. \tag{4.1}$$

The weight contribution can be computed by taking the determinant and does not depend on the time slice t due to cyclicity. Since the matrices are of size $N_s \times N_s$, the naïve computational cost of constructing the Green's function for a single time slice and field configuration is $\mathcal{O}(N_s^3 N_t)$, assuming standard matrix multiplication. This significant bottleneck in simulations calls for specialized methods to accelerate computations.

4.1.1 Fourier acceleration

The first improvement we can make to constructing the Green's function is using sparse matrix operations. Remember that the matrices B_n are a product of kinetic and interaction contributions,

$$B_n = e^{-\Delta t T} V(\phi_n), \tag{4.2}$$

where the potential part is diagonal in position space for the decompositions discussed in Section 2.3.2. The form of the kinetic part depends on the desired lattice dispersion. Suppose we use the Hubbard dispersion, which allows for hopping only between nearest neighbors. In that case, the kinetic part is sparse in position space, and one may use checkerboard breakup [113] to decrease the cost to effectively $\mathcal{O}(N_s^2 N_t)$. However, the kinetic part is not sparse in position space for more general dispersion relations. In particular, the continuum dispersion of free particles $p^2/2m$ yields a dense matrix, rendering the checkerboard breakup ineffective. Instead, we make use of so-called Fourier acceleration in all our computations by noticing that the kinetic part is diagonal in momentum space. The matrix product is constructed sequentially, performing a Fourier transform on the partial product of matrices at each step. The full product can be constructed from right to left, starting in momentum space, where we denote it as $\tilde{U}_{\text{partial}}$. Each step proceeds as

$$B_n \tilde{U}_{\text{partial}} = \overbrace{V(\phi_n) \underbrace{\left(e^{-\Delta t T} \tilde{U}_{\text{partial}} \right)}_{\text{iFFT Col}}}^{\text{FFT Col}}, \tag{4.3}$$

where we act with an (inverse) Fourier transformation on the matrix between the multiplications, indicated by the curly brackets. In this case, the transformations act on each matrix column but leave the rows untouched. If we instead construct the matrix from right to left, we begin with the matrix in position space, yielding

$$U_{\text{partial}} B_n = \overbrace{\underbrace{\left(U_{\text{partial}} V(\phi_n) \right)}_{\text{FFT Row}} e^{-\Delta t T}}^{\text{iFFT Row}}, \tag{4.4}$$

and each row is taken to momentum space and back. In practice, we use both directions extensively in a modified form. Since all of our matrices are stored in column-major format, we only create FFT plans for the columns of the matrices. This is because they are contiguous in memory, making the FFTs easily parallelizable. We found plans created on the rows to degrade performance significantly. When constructing the matrix from right to left, we transpose the entire matrix and apply the FFTs column-wise. The transpose is a simple operation and does not significantly impact the performance for the matrix sizes we are dealing with.

The overall procedure can lead to a significant reduction in runtime. Matrix operations are reduced to $\mathcal{O}(N_s^2 \log N_s)$, as they only involve diagonal matrices. The FFTs acting on each column of the matrix are $\mathcal{O}(N_s \log N_s)$, which is the dominant cost. Overall, the scaling using Fourier acceleration becomes $\mathcal{O}(N_s^2 N_t \log N_s)$, which is a significant improvement over the naïve $\mathcal{O}(N_s^3 N_t)$. In Section 4.3, we will see that this cost can be further reduced by using matrix decomposition and truncation techniques.

4.1.2 Stable matrix algebra

Quite early on in the development DQMC, it was realized that computing the product of the matrices B_n is numerically unstable at low temperatures [114, 115]. To illustrate the issue, we will consider the simple example of non-interacting fermions in a single spatial dimension. We use a lattice with $L = 8$ sites and Hubbard dispersion, an anisotropy of $\Delta t = 0.05$, and tune the chemical

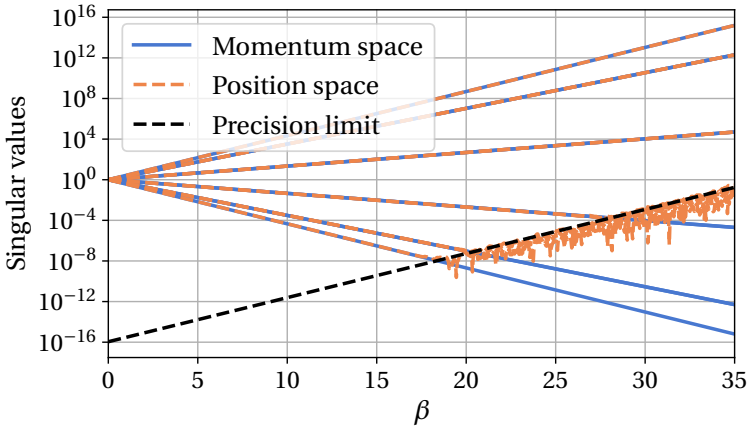

Figure 4.1: Singular values of the matrix U for a system of $L = 8$ sites at half-filling with Hubbard dispersion and an anisotropy of $\Delta t = 0.05$. The solid lines indicate the singular values computed directly from the momentum space matrix, while the dashed orange lines are computed via a singular value decomposition (SVD) after computing the product in position space. The dashed black line indicates the maximum relative precision expected for a double-precision floating-point number.

potential to half-filling. In momentum space, constructing the matrix U is unproblematic, as the B_n are all diagonal, and no mixing of energy states occurs. With decreasing temperature and fixed anisotropy, the number of component matrices increases, leading to an exploding condition number in their product[1]. An issue arises if we try to work in position space instead. Each component B_n and the partial U matrices are now dense, and matrix elements mix at different scales.

Once the condition number becomes larger than the maximum relative precision of a double precision floating point number, roughly 1.11×10^{-16}, errors in the matrix product will accumulate. This can lead to numerically unstable simulations and wrong results. The singular values for the toy model as a function of the temperature are shown in Figure 4.1[2]. The dashed black line indicates the theoretical precision bound relative to the biggest singular value. All singular values computed directly from the diagonal matrix in momentum space are well-behaved and untouched by precision problems. However, mixing in position space leads to a loss of precision, starting with the smallest singular values and progressively destroying any information below the precision limit.

Initially, this may not appear to be a major issue. After all, the states with tiny singular values correspond to energy states far above the Fermi surface and should not contribute to the physics of the system. However, once the precision limit becomes comparable to the scale of the Fermi surface, given by 10^0, numerical errors will inevitably appear in physical observables. To illustrate the problem we consider the total particle number in a given system, which we may write in terms of the eigenvalues d_i of U as

$$N = \sum_i \frac{d_i}{d_i + 1}. \tag{4.5}$$

[1] The condition number is directly related to the ratio of the maximum and minimum eigenvalues of the matrix, but equal only if the matrix is normal.

[2] For the given example, the singular values and eigenvalues are identical.

Clearly, as long as any single eigenvalue stays well below unity, its actual value is irrelevant to the overall sum. If, however, the small eigenvalues are lifted up, as can be seen in Figure 4.1, the total particle number will be overestimated. As the temperature is lowered, the system appears to move away from half-filling, and the computed particle number will wrongly indicate a fully filled lattice in the $\beta \to \infty$ limit.

We have seen that the non-interacting system allows a simple solution to the stability problem by computing in momentum space, where the Hamiltonian is diagonal. Unfortunately, the generic case with finite interaction between spin species and potentially complex trapping geometry does not allow such a straightforward resolution. Instead, matrix decomposition techniques are typically used to keep the diverging scales separated throughout the simulation [114, 116]. In the following, we will introduce the general idea behind the aforementioned decompositions, enabling us to construct the weight, Green's function, and any observables in a numerically stable manner.

Constructing matrix products

We can construct the matrix products using various matrix decomposition methods, including SVD, (pivoted) QR, and eigenvalue decompositions. While all approaches have an asymptotic computational cost of $\mathcal{O}(N_s^3)$, prefactors can vary significantly. An exhaustive study regarding the properties of the different decompositions and their performance in the context of DQMC can be found in [116]. We will mostly use column-pivoted QR decompositions, since they offer relatively fast runtimes and near-perfect numerical stability when paired with the truncation approaches we will discuss in Section 4.3. Note that the general framework for computations is the same for all methods. We start by writing the partial product of matrices as

$$U_n = B_n U_{n-1} = B_n \dots B_1 \,, \tag{4.6}$$

and denote the decomposition at each step as

$$U_n = L_n D_n R_n \,. \tag{4.7}$$

Here, D_n is a diagonal matrix, while the nature of L_n and R_n depends on the specific method. For QR decomposition, L_n is unitary while R_n is an upper triangular with a unit determinant (up to a phase). Starting with this form, we multiply with a single matrix B_{n+1} to obtain the next partial product. We can keep the scales separated throughout this whole procedure by organizing the operations appropriately, namely

$$U_{n+1} = \left(B_{n+1} L_n \begin{bmatrix} d & & & \\ & d & & \\ & & d & \\ & & & d \end{bmatrix} \right) R_n = \begin{bmatrix} X & X & X & x \\ X & X & X & x \\ X & X & X & x \\ X & X & X & x \end{bmatrix} R_n \,. \tag{4.8}$$

The letters of different sizes indicate the scales, which are initially separated and ordered. Importantly, the scales in the matrix B_n need to be under numerical control so that the multiplication with the diagonal matrix does not lead to information loss. In the first step, we perform the matrix multiplications of all matrices on the left, leading to a column-stratified result and avoiding mixing.

Note that if we were to include the final multiplication with R_n, the scales would be mixed, and information about the small values would be lost. To avoid this, we perform another decomposition on the column matrix,

$$\underbrace{\begin{bmatrix} X & X & X & x \\ X & X & X & x \\ X & X & X & x \\ X & X & X & x \end{bmatrix}}_{L_{n+1}D_{n+1}\tilde{R}_{n+1}} R_n = L_{n+1}D_{n+1}R_{n+1}\,, \tag{4.9}$$

and only multiply the resulting unit scale matrices. For QR decomposition, L_{n+1} stays unitary, but R_{n+1} will not generally be triangular. We can now construct the full matrix product without losing relevant information on the smaller scales. In practice, several multiplications with B_n can be made from the left before performing a new decomposition, saving significant computational time. The matrix can also be constructed from left to right, leading to an analogous procedure for row-stratified matrices.

In some situations, we may want to multiply two decomposed matrices in a numerically stable way, which can be achieved via

$$\begin{aligned} B_1B_2 &= L_1D_1R_1L_2D_2R_2 \\ &= L_1\underbrace{(D_1\,(R_1L_2)\,D_2)}_{\tilde{L}D\tilde{R}}R_2 \\ &= LDR\,, \end{aligned} \tag{4.10}$$

where the different scales behave as follows:

$$\begin{aligned} &\begin{bmatrix} d_1 & & & \\ & d_1 & & \\ & & d_1 & \\ & & & {\scriptstyle d_1} \end{bmatrix} \underbrace{\begin{bmatrix} X & X & X & X \\ X & X & X & X \\ X & X & X & X \\ X & X & X & X \end{bmatrix}}_{\text{unit scale}} \begin{bmatrix} d_2 & & & \\ & d_2 & & \\ & & d_2 & \\ & & & {\scriptstyle d_2} \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} XX & XX & X_X & X_x \\ XX & XX & X_X & X_x \\ XX & XX & XX & X_X \\ XX & XX & XX & XX \end{bmatrix}}_{\tilde{L}D\tilde{R}}\,. \end{aligned} \tag{4.11}$$

The matrix in the center of (4.10) is the product of two unit scale matrices and easy to compute. After multiplication with the diagonal matrices, the scales remain well separated, decreasing with both the column and row index. Moreover, there is no numerical issue when multiplying values of different scales. The decomposition of the product can be performed as before and is usually stable. However, it can exhibit instability in extreme situations compared to the construction with purely column-stratified matrices.

Computing the Green's function

For a successful DQMC simulation, we next need viable way to compute the weight and Green's function. Both can be hard to obtain, as the calculations involve sums and inverses of the ill-conditioned matrices. The weight is given by the determinant of the inverse Green's function $\mathrm{I}+U$, which we compute as

$$\begin{aligned} G^{-1} &= \mathrm{I} + U \\ &= \mathrm{I} + LDR \\ &= L\underbrace{\left(L^{-1}R^{-1} + D\right)}_{\tilde{L}\tilde{D}\tilde{R}} R, \end{aligned} \tag{4.12}$$

where the inversions of L and R are unproblematic. In the last step, we perform an additional decomposition to remain in LDR form. From this, the weight can be computed as the product of all diagonal values. Inevitably, adding the identity matrix to the product destroys the small scales beyond 10^{-15}, making the procedure not invertible in practice. The weight itself is not affected by this, but it can become an issue in downstream computations of observables using the Green's function, which must be kept in mind. Moreover, the procedure described above can be insufficient in extreme scenarios. An extended scheme for these cases is provided in [117], where the diagonal matrix is factorized into scales smaller and larger than the Fermi surface $D = D_l D_s$ with the $N_s \times N_s$ matrices

$$D_l = \max(D, 1), \qquad D_s = \min(D, 1)\,. \tag{4.13}$$

The inverse Green's function can be computed safely with two intermediate decomposition as

$$\begin{aligned} G^{-1} &= 1 + LDR \\ &= 1 + LD_l D_s R \\ &= \underbrace{\Big[\underbrace{\left(R^{-1}D_l^{-1} + LD_s\right)}_{\tilde{L}\tilde{D}\tilde{R}} D_l\Big]}_{L'D'R'} R. \end{aligned} \tag{4.14}$$

In most of the computations in this work, we have used the first method, as we found the more elaborate scheme unnecessary in the presence of the model space truncation, discussed in Section 4.3.

Stabilization for the Canonical Ensemble

Thus far, we have discussed methods to compute and construct the Green's function for grand canonical computations in a stable manner. In canonical computations, as discussed in Section 2.4, the projections naïvely require an independent computation of the weight for every point in the discretized frequency domain, which is prohibitively expensive. However, this issue may be avoided by using the diagonalized form of the Green's function, which can be computed from the decomposed matrix product [65]. To do so, we consider a cyclic permutation of the decomposition,

which leaves the eigenvalues invariant, yielding

$$\begin{aligned} LDR &= L\,[DRL]\,L^{-1} \\ &= L[\tilde{P}D_\lambda \tilde{P}^{-1}]L^{-1} \\ &= PD_\lambda P^{-1}, \end{aligned} \tag{4.15}$$

where D_λ is the diagonal matrix of eigenvalues and P is the matrix of eigenvectors. Conveniently, the permuted matrix is in row stratified format, avoiding the mixing of scales,

$$DRL = \begin{bmatrix} X & X & X & X \\ X & X & X & X \\ X & X & X & X \\ x & x & x & x \end{bmatrix}. \tag{4.16}$$

The eigenvalues can be computed safely with solvers using balancing, as provided by LAPACKS `zgeevx` routine. The computation of the weight via the Fourier sum in Equation (2.64) reduces from complexity $\mathcal{O}(N_s^4)$ to $\mathcal{O}(N_s^3)$, assuming the number of frequencies corresponds to the number of states in the system, and is given by

$$Z_N^\phi = \frac{1}{N_s}\sum_{n=1}^{N_s} e^{-iN\varphi_n} e^{-N\beta\mu} \prod_j \left[1 + e^{i\varphi_n} e^{\beta\mu} d_j\right]. \tag{4.17}$$

Here, d_j are the eigenvalues of U, and a product of eigenvalues replaces the computation of the determinant. Note that in practical computations, the factor $e^{\beta\mu}$ is already included in the diagonal values. To ensure the precision of the eigen decomposition, we compare the $\varphi_n = 0$ term to the grand canonical weight computed via the QR decompositions in every step.

A similar procedure can be used to reduce the complexity of computing observables. For a general bilinear operator, as used in Section 2.4.2, we find

$$\begin{aligned} \langle O\rangle_N^\phi &= \frac{e^{-N\beta\mu}}{Z_N^\phi N_s}\sum_{n=1}^{N_s} e^{-iN\varphi_n} \det\left[\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}U\right] \mathrm{Tr}\left[\left(\mathrm{I} - \left(\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}U\right)^{-1}\right) O\right] \\ &= \frac{e^{-N\beta\mu}}{Z_N^\phi N_s}\sum_{n=1}^{N_s} e^{-iN\varphi_n} \prod_j \left[\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}d_j\right] \mathrm{Tr}\left[\left(1 - \left(\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}PD_\lambda P^{-1}\right)^{-1}\right) O\right] \\ &= \frac{e^{-N\beta\mu}}{Z_N^\phi N_s}\sum_{n=1}^{N_s} e^{-iN\varphi_n} \prod_j \left[\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}d_j\right] \mathrm{Tr}\left[\left(\mathrm{I} - \left(\mathrm{I} + e^{i\varphi_n}e^{\beta\mu}D_\lambda\right)^{-1}\right) P^{-1}OP\right], \end{aligned} \tag{4.18}$$

where the last step uses the cyclic property of the trace. The matrix product $P^{-1}OP$ needs to be computed only once for each auxiliary field configuration, and the preceding term reduces to a delta. Higher order correlations may be obtained similarly using Wicks theorem, and a second projection must be included to fix the particle number for both species.

4.2 Updating

In the next section, we discuss some general procedures for efficient updating with DQMC for continuous fields using global updating schemes and discrete fields with time-local updates. Both

update schemes rely on an efficient way to compute the equal time Green's function (2.54) and weight. The different methods each offer their advantages and disadvantages in terms of computational cost, the autocorrelation of the Markov chain, and variance of the measured observables [118–120]. Notably, time-local updates typically lead to lower variances.

4.2.1 Wrapping

Naively computing the matrix product, addition, and inversion for each time slice independently would incur a high computational cost. However, we can avoid this by recognizing that Green's functions evaluated at neighboring time slices are related by

$$\begin{aligned} G^{\phi}(t+1) &= \left(\mathrm{I} + B_{t+2} \dots B_{N_t} B_1 \dots B_{t+1}\right)^{-1} \\ &= B_{t+1} \left(\mathrm{I} + B_{t+1} \dots B_{N_t} B_1 \dots B_t\right)^{-1} B_{t+1}^{-1} \\ &= B_{t+1} G^{\phi}(t) B_{t+1}^{-1} . \end{aligned} \tag{4.19}$$

Using this relationship, we efficiently move the Green's function between different time slices by multiplying by B_{t+1} and its inverse, for which we use Fourier acceleration. This procedure is known as wrapping and is a crucial part of the algorithm. We take the same approach for the Matrix product and the inverse of the Green's function. Unfortunately, wrapping for an arbitrary number of time slices is not possible due to the large condition number at low temperatures. Instead, we typically set a fixed timescale t_w for wrapping, after which the Green's function is recomputed from scratch. The number of complete computations of the matrix product thus scales with N_t / t_w instead of N_t. Our simulations typically use between $t_w = 10$ and $t_w = 20$, which offers a significant speedup while keeping errors below the numerically relevant threshold.

4.2.2 Time local updates

Various updating schemes for DQMC calculations exist. The most straightforward approach changes only a single spin in each accept-reject step, while updating the Green's function and weights accordingly [56]. This corresponds to space-time local updating, which is efficient for small systems but becomes impractical as system size increases. Instead, we employ a time local scheme, inspired by the one used in [41] for ground state computations, to perform updates on all fields simultaneously on a single time slice.

Consider the fully computed matrix product $U(\phi, t)$, where t as an argument indicates the permutation. We may change the auxiliary field values on a single time slice by multiplying with the appropriate potential factors from the right side, like

$$U(\phi, t) = B_{t+1}(\phi_{t+1}) \dots B_{N_t}(\phi_{N_t}) B_1(\phi_1) \dots B_t(\phi_t) , \tag{4.20}$$

$$U(\phi', t) = U(\phi, t) \frac{V(\phi'_t)}{V(\phi_t)} , \tag{4.21}$$

where ϕ' differs from the original field only on the time slice t. Unlike in space-time local approaches where the Green's function is used to compute the relative weight between different configurations [114], we typically do not track it during computations. Instead, we only keep $U(\phi, t)$ and compute (4.12) for every proposed field configuration to obtain the weight. This comes with an

$\mathcal{O}(N_s^3)$ cost for the matrix decomposition, which is the same order one would obtain from updating all fields on a single time slice separately with space-time local updates.

When updating multiple fields at once, the choice of proposal distribution becomes crucial to maintain a high acceptance rate in the Markov chain. We use the finite temperature version of the approach employed in [41], where proposed updates are conditioned on the current configuration of all other time slices. This works particularly well for discrete and finite auxiliary fields, as the normalization of the proposal distribution is easy to compute. For the discrete decomposition in density channel, the proposal distribution becomes

$$P(\phi'_t|\phi) = \frac{1}{\mathcal{N}} \prod_x \left(1 + A\phi'_{t,x}\tilde{n}_{\uparrow,x}\right)\left(1 + A\phi'_{t,x}\tilde{n}_{\downarrow,x}\right), \tag{4.22}$$

and the partial densities $\tilde{n}_x$ are computed like the density but dropping the time slice under consideration, yielding

$$\tilde{U}(\phi, t) = U(\phi, t)\frac{1}{V(\phi_t)}, \tag{4.23}$$

$$\tilde{n}_x = \mathrm{diag}\left[\frac{\tilde{U}(\phi, t)}{1 + \tilde{U}(\phi, t)}\right]. \tag{4.24}$$

The proposals reproduce the exact distribution to order Δt, improving as the time step decreases in the continuum limit. In our computations, acceptance rates are typically well above 90% with $\Delta t = 0.05$ and sufficient to ensure fast decorrelation of the Markov chain.

4.2.3 Global updates

In order to perform global updates via Hybrid Monte Carlo or Langevin dynamics, we need to compute the gradient of the action,

$$S[\phi] = \mathrm{Tr}\log\left[1 + U_\uparrow\right] + \mathrm{Tr}\log\left[1 + U_\downarrow\right], \tag{4.25}$$

with respect to the auxiliary fields. The gradient is computed straightforwardly, yielding

$$
\begin{aligned}
K_{t,x} &= \partial_{\phi_{t,x}} \operatorname{Tr} \log 1 + U \\
&= \operatorname{Tr}\left[\left(\mathrm{I} + B_1 \dots B_{N_t}\right)^{-1} B_1 \dots B_{t-1} \frac{\partial B_t}{\partial \phi_{t,x}} B_{t+1} \dots B_{N_t} \right] \\
&= \operatorname{Tr}\left[\left(1 + B_1 \dots B_{N_t}\right)^{-1} \left(B_{N_t}^{-1} \dots B_{t+1}^{-1} \left(\frac{\partial B_t}{\partial \phi_{t,x}} \right)^{-1} B_{t-1}^{-1} \dots B_1^{-1} \right)^{-1} \right] \\
&= \operatorname{Tr}\Bigg[\Bigg(B_{N_t}^{-1} \dots B_{t+1}^{-1} \left(\frac{\partial B_t}{\partial \phi_{t,x}} \right)^{-1} B_{t-1}^{-1} \dots B_1^{-1} \; + \\
&\qquad B_{N_t} \dots B_{t+1} \left(\frac{\partial B_t}{\partial \phi_{t,x}} \right)^{-1} B_n^{-1} \dots B_{N_t}^{-1} \Bigg)^{-1} \Bigg] \\
&= \operatorname{Tr}\left[\left(\left(\frac{\partial B_t}{\partial \phi_{t,x}} \right)^{-1} B_{t-1}^{-1} \dots B_1^{-1} B_{N_t}^{-1} \dots B_{t+1}^{-1} + \left(\frac{\partial B_t}{\partial \phi_{t,x}} \right)^{-1} B_t \right)^{-1} \right] \\
&= \operatorname{Tr}\left[B_t^{-1} \underbrace{\left(B_{t-1}^{-1} \dots B_1^{-1} B_{N_t}^{-1} \dots B_{t+1}^{-1} B_t^{-1} + \mathrm{I} \right)^{-1}}_{\tilde{G}(t)^{-1}} \frac{\partial B_t}{\partial \phi_{t,x}} \right]
\end{aligned}
\tag{4.26}
$$

In this expression, $\tilde{G}(t)$ is directly related to the equal-time Green's function via

$$
\begin{aligned}
\tilde{G}^{-1}(t-1) &= \left(B_{t-1}^{-1} \dots B_1^{-1} B_{N_t}^{-1} \dots B_t^{-1} + \mathrm{I} \right)^{-1} \\
&= B_t \dots B_{N_t} B_1 \dots B_{t-1} \left(B_t \dots B_{N_t} B_1 \dots B_{t-1} + \mathrm{I} \right)^{-1} \\
&= \mathrm{I} - \left(B_t \dots B_{N_t} B_1 \dots B_{t-1} + \mathrm{I} \right)^{-1} \\
&= \mathrm{I} - G(t-1)^{\phi},
\end{aligned}
\tag{4.27}
$$

and may be propagated by wrapping as before. The matrix product can again be computed using the decomposition techniques discussed in the last section. Importantly, to compute the gradient, we use the final expression from (4.26), as it avoids extra multiplications that could reduce numerical accuracy.

To obtain the full gradients on all fields, we still need to wrap through all time slices, which comes with the same computational cost as a full sweep of single-site updates. However, the global update does not require a weight comparison at every time slice and may be more efficient regarding autocorrelation times.

4.3 Model space truncation

The regular BSS or HMC algorithms provide an exact approach for computations of systems like the Hubbard model at half-filling and ultracold fermions at finite temperatures. However, they are not as efficient as the respective ground state variants of the algorithms (see Section 2.4). This inefficiency is particularly problematic for simulations of ultracold fermions, which require a continuum extrapolation where the number of lattice points increases, but the particle content stays

fixed. The finite temperature algorithm scales with a complexity of $O(N_t N_s^3)$, whereas the ground state algorithm scales more efficiently with $O(N_t N N_s^2)$, where N is the number of particles. Thus, there is a scaling difference of N_s/N that increases as the continuum is approached. Here, we have assumed that the matrix multiplications are the primary computational cost source. The scaling difference can become even more crucial when using Fourier acceleration, as the ground state algorithm does not have other N_s^2 terms.

Recent studies have addressed this discrepancy for both the canonical [121] and grand canonical ensembles [122] using truncation schemes on the decomposed matrix product. These schemes dynamically exclude low-lying, unoccupied modes from computations. As a result, matrices become rectangular instead of square, leading to computational costs comparable to the ground state method as the temperature decreases. In the following section, we will give an overview of the truncations used in this work and discuss their effectiveness and implementational details.

4.3.1 Truncations in the grand canonical ensemble

Since we perform all sampling in the grad canonical ensemble, we begin by detailing the truncation procedure for this case. The main quantity of interest both during sampling and when evaluating observables is the (inverse) Green's function. Assuming the matrix product U is computed as an eigendecomposition $U = PDP^{-1}$, the Green's function can be written as

$$G^{-1} = \mathrm{I} + U = \mathrm{I} + PDP^{-1} = P\,(\mathrm{I} + D)\,P^{-1}\,, \tag{4.28}$$

and the contribution to the probability weight becomes $\prod_i 1 + d_i$. Thus, the small eigenvalues of U do not contribute to the weight, nor do they significantly affect the Green's function from which observables are computed. This insight can be utilized when constructing U by dropping the smallest modes during intermediate matrix decompositions. While we typically do not use eigendecomposition, opting instead for SVD or column-pivoted QR decomposition, the general principle of dropping the smallest values in the diagonal matrix remains valid as long as the scales can be kept separated.

We start with a matrix decomposition $U = LDR$, where the diagonal values are ordered in descending magnitude, and we have already applied a truncation to include only the first m modes. Clearly, we must have $1 \leq m \leq N_s$, and the truncation effectively sets all values beyond the m-th entry in D to zero. The decomposed matrices are of shape $N_s \times m$, $m \times m$, and $m \times N_s$, respectively, and the general procedure of multiplying additional time slices goes as follows:

1. Multiply a new batch of matrices $B_{j+i} \dots B_i$ from the left. This has a cost of $\mathcal{O}(mN_s \log Ns)$ with Fourier acceleration, since the right side has a reduced number of columns.

2. Multiply the left and diagonal matrices with a cost of $\mathcal{O}(m^2 N_s)$.

3. Perform a new decomposition on the rectangular matrix, which also comes at a cost of $\mathcal{O}(m^2 N_s)$. This yield three more matrices in LDR format, the first is of size $N_s \times m$, while the others are $m \times m$.

4. Cut all diagonal values below a certain threshold ϵ_T to obtain a new truncation rank m'. We use $\epsilon_T = 10^{-5}$ in all computations where the truncation is applied.

5. Multiply the right matrix in the new truncation, which is $m' \times m$, with the original $m \times N_s$ matrix, at a cost of $\mathcal{O}(m' m N_s)$.

It must be noted that the described procedure still contains terms that scale as $\mathcal{O}(N_s^3)$ since the initial QR decompositions are performed on the full rank matrix. However, all subsequent operations can benefit from the truncation, avoiding the cost of the conventional method, which carries an additional factor N_t proportional to the number of time slices.

The computation of the weight can also be made more efficient when using the truncated decomposition since

$$\det\left(\mathrm{I}_{N_s} + LDR\right) = \det\left(\mathrm{I}_m + DRL\right), \tag{4.29}$$

where DRL is an $m \times m$ matrix and needs $\mathcal{O}(m^2 N_s)$ operations to compute. We obtain the determinant via an additional decomposition on the $m \times m$ matrix, which is an $\mathcal{O}(m^3)$ operation and a significant improvement over the $\mathcal{O}(N_s^3)$ cost of the full rank matrix. One may use a procedure similar to the one in (4.14) to ensure numerical stability as discussed, e.g., in [123]. However, since the smaller scales are already excluded from the computation, we have not found it necessary in our simulations.

The equal-time Green's function can be obtained from the truncated decomposition via the Woodbury matrix identity, yielding

$$\begin{aligned} G &= \left(\mathrm{I}_{N_s} + LDR\right)^{-1} \\ &= \mathrm{I}_{N_s} - L\left(D^{-1} + RL\right)^{-1} R. \end{aligned} \tag{4.30}$$

The final line can be evaluated in $\mathcal{O}(m N_s^2)$ operations but is only relevant when we compute observables and not during the generation of samples.

To demonstrate the efficiency of the truncation scheme, we perform a test on a one-dimensional model with attractive contact interactions. We measure the efficiency by comparing the number of active diagonal values m with the number of particles in the system N as

$$\eta_T = \frac{N_s - m}{N_s - N} \tag{4.31}$$

where $\eta_T = 1$ indicates that the truncation is optimal and $\eta_T = 0$ that no truncation is performed [122]. Figure 4.2 shows the test results for systems with and without a harmonic trapping potential, using $10 + 10$ particles on a lattice with $N_s = 120$ sites, $N_t = 1200$ time slices, $\Delta t = 0.05$, and a coupling of $g = -1$, which we tune in the tapped system. We measure the efficiency for several thousand steps during sampling to yield a smooth average curve. All computations are grand canonical and achieve a fixed particle content by tuning the chemical potential.

Both cases show a rapid decimation of states as the matrix product is assembled, which can be interpreted as decreasing the temperature. In the trapped case, $\eta_T \geq 0.9$ once the temperature becomes larger than the single particle gap of the harmonic oscillator. In contrast, the untrapped case shows similar behavior when the temperature becomes comparable to the Fermi temperature. Notably, density channel computations show faster state decimation compared to the spin channel due to the different ways interactions enter the matrix product. In the density channel, the real-valued interactions increase the condition number, while interactions of pure phases in spin channel lead to slower mode decimation. This behavior is specific to attractive contact interactions and is expected to reverse for repulsive systems.

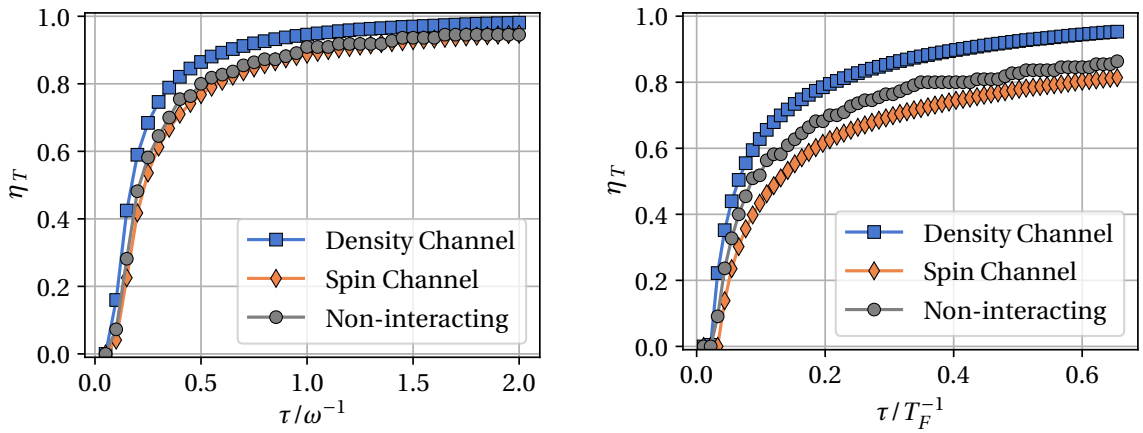

Figure 4.2: Efficiency of the truncation as the matrix product in build up for the trapped (left) and uptrapped (right) systems. All computations are done in 1D, and we use $g = -1$. The density channel computations (blue squared) result in a faster decimation of states compared to the spin channel (orange diamond) and non-interacting (gray circle) cases.

4.3.2 Truncations in the canonical ensemble

The truncation scheme in the canonical ensemble works similarly to the grand canonical case, with a few additional steps. In particular, the reconstruction of the eigenvectors, needed for efficient computations from the truncated matrices, must be handled with care.

Overall, the procedure stays much the same, but the eigendecomposition takes the form

$$U = P_m D_m P_m^{-1}\,, \tag{4.32}$$

where P_m and P_m^{-1} are the first m columns/rows of the original eigenvector matrix and its inverse. From this, we compute bilinear observables as

$$\langle \psi_i^\dagger \psi_j \rangle_N^\phi = \frac{e^{-\beta\mu}}{Z_N N_s} \sum_{n=1}^{N_s} e^{-iN\varphi_n} \prod_j^m \left[1 + e^{i\varphi_n} e^{\beta\mu} \lambda_j\right] \sum_k^m P_{m,jk} \frac{\lambda_k}{1+\lambda_k} P^{-1}_{m,ki}\,, \tag{4.33}$$

where the λ_j are the truncated eigenvalues of U and the $P_{m,ni}$ are the elements of the truncated eigenvector matrix. In (4.33), the Fourier sum runs over a full set of states, but may be reduced to $N_f < N_s$ quadrature points to reduce computational costs. We typically use values comparable or larger than the number of particles per spin species.

To obtain the truncated eigendecomposition from the decomposed matrix product, we first compute the eigenvectors of the $m \times m$ matrix, yielding

$$DRL = \tilde{P}\tilde{D}\tilde{P}^{-1}\,, \tag{4.34}$$

where $\tilde{P}$ and $\tilde{D}$ are also $m \times m$ matrices. The eigenvalues correspond to the first m eigenvalues of U, as the (full rank) matrices are related via a similarity transform. To compute the truncated

eigenvector matrix and its inverse, we may use the relations

$$P_m = L\tilde{P} \tag{4.35}$$

$$P_m^{-1} = X^{-1}R, \tag{4.36}$$

where L and R are the truncated left and right matrices from the decomposition of U. X is the $N_s \times N_s$ matrix

$$X = RL\tilde{P}. \tag{4.37}$$

We refer to the appendix of [121] for a detailed derivation of these relations. Our approach differs from the one presented therein, as we truncate the decomposition during the construction of the matrix products and only have part of the left and right unit scale matrices. This is not an issue, since none of the relations rely on the full rank matrix, which saves computational effort during construction.

Overall, the truncation scheme in the canonical ensemble reduces computational costs similarly to the grand canonical case. The leading contribution to the eigendecomposition comes from (4.37) with $\mathcal{O}(N_s^2 m)$. Furthermore, the full observable in (4.33) is computed in $\mathcal{O}(N_s^2 N_f m)$ operations.

4.4 Implementation

We now discuss some implementational details regarding the truncation schemes, which help us balance the number of repeated operations with memory usage. So far, we have assumed that the full matrix is recomputed after applying a certain number of wrapping steps t_w. However, this approach is costly because the initial decomposition when reconstructing the matrix is of full rank each time. Instead, we split the matrix product into two parts, which are computed only once and partially cached to avoid excessive repeated operations. At the start of each sweep over the lattice, we compute the full product once moving from left to right. We store the decomposed partial product of matrices to memory as P_i after each decomposition, such that

$$U = \overbrace{\underbrace{B_1 \dots B_{t_w}}_{P_1} \dots B_{2t_w}}^{P_2} \dots, \tag{4.38}$$

where $i \leq N_t / t_w$. Since the involved matrices are complex and double precision, storing all partial products requires $32 \times N_s^2 \times N_t / t_w$ bytes, disregarding the diagonal matrix. This memory requirement can become hard to handle for systems in higher dimensions or at low temperatures. Therefore, we only store intermediate products when truncation reaches a certain value m_T. The cached matrices are $N_s \times m_T$ or $m_T \times N_s$, resulting in a decreased storage cost of $32 \times N_s m_T N_t / t_w$ bytes.

A sweep begins with the fully constructed matrix product, which we wrap t_w times as discussed in Section 4.2. After each block, we sequentially construct the product of matrices with the updated fields,

$$P_{L,n} = B_{nt_w} \dots B_{N_t}, \tag{4.39}$$

built from right to left. After applying a set of wrapping steps, we recompute the matrix product via

$$U(nt_w) = P_{L,n} P_n \tag{4.40}$$

as described in (4.10). Since both decompositions are already truncated, the operations involved are inexpensive, scaling as $\mathcal{O}(m^2 m')$, where $m' > m$. If the decomposition of P_n is not in storage due to insufficient truncation $m > m_T$, we compute the full matrix by applying the standard products and decompositions on $P_{L,n}$ from the right. Although suboptimal, this does not pose a major problem, as P_n is usually already quite small in cases where we do not have the remaining matrices in memory and vice versa.

Unfortunately, the product for the updated field, computed on the fly, cannot be used in the next sweep, as it is constructed from right to left. As a result, we must compute the matrix product twice per sweep, which doubles the number of full-rank matrix operations.

Another critical aspect of the implementation involves determining when to compute observables and thus when to take measurements[3]. Indeed, DQMC simulation can suffer from a slow decay of sampled distributions, which may lead to exploding variance when computing standard errors for some higher-order observables or slow convergence of the observable itself. The issue is particularly relevant for the two-dimensional canonical computations performed in Chapter 6. Several strategies to mitigate the problem have been proposed. We considered the reweighting approach discussed in [119, 120], as well as the measurement synchronization scheme discussed in [124]. Since we find the latter to be sufficient for our purposes, and the computational effort required is smaller, we use it for all grand canonical computations in two dimensions. For completeness, a brief introduction to the reweighting method is provided in Appendix A.

Since our updating scheme is time-local, we cannot use the space-time synchronized measurements proposed in [124], but use time synchronized ones instead. For equal-time observables, (imaginary-)time translation allows for a simple average over all times, yielding

$$\langle O \rangle = \frac{1}{N_t} \sum_{t=1}^{N_t} \langle O(t) \rangle \,, \tag{4.41}$$

which boosts statistics. Typically, we would apply a full sweep to update, followed by evaluating the observables for all times. However, in the time-synchronized scheme, we measure the observable at a specific time immediately after updating the corresponding time slices. This approach reduces variance and decreases computational effort, as we avoid building the relevant matrix products twice when computing observables simultaneously with updates.

[3] Computing an observable and storing its result to memory is often referred to as 'measuring' in the context of lattice Monte-Carlo simulations. This is not to be confused with actually measuring observables in experiments.

Chapter 5

Trapped fermions in one dimension

In early cold atoms experiments, (approximately) harmonic potentials were used to trap the atoms. While arbitrary potentials can be realized today [125–128], the harmonic trap remains a popular choice due to its simplicity and the precise control over, e.g., the particle number that can be achieved [129–133].

The majority of lattice simulations of ultracold quantum gases have instead been conducted without an external potential and with the goal of reaching a thermodynamic limit. They encompass studies of the ground state and finite temperature properties in one to three spatial dimensions [41, 59, 134], and the thermodynamics of the unitary gas [52, 135] in particular. More recently, the BKT-transition temperature [42] and pseudogap effects [136] were computed in the BEC-BCS crossover regime of the 2D gas.

In this section, our focus is on harmonically trapped fermions in one spatial dimension and methods to make them amenable to lattice computations. Previous theoretical investigations on the ground state of the trapped system have employed diverse methodologies. Exact diagonalization approaches, typically confined to small particle numbers, have been utilized [137–139]. Non-uniform lattice Monte Carlo methods [60, 61] have been used to study systems of up to 20 particles, while coupled cluster [140, 141] and diffusion Monte Carlo [142] approaches have allowed for computations with higher particle numbers. At finite temperature, the canonical system has been studied with exact diagonalization [143], while lattice-based methods include grand canonical auxiliary field and canonical stochastic Green's function [48] approaches.

While quantum Monte Carlo methods have been successfully applied to study trapped systems, as stated above, they often rely on the one-dimensional nature of the problem to avoid a sign problem. The lattice approach used here does not suffer from a sign problem for spin-balanced systems in any dimension. This chapter thus lays the groundwork for future studies of trapped systems via lattice approaches, particularly in two spatial dimensions.

We structure our discussion into three main components. In Section 5.1, we discuss how the harmonic potential is added to the lattice model. Section 5.2 is dedicated to the application of the complex Langevin method to trapped fermions, encompassing a discussion of the method's general properties and a specific focus on repulsive contact interactions. Finally, Section 5.3 presents results from canonical Monte Carlo simulations of trapped fermions from finite temperature down to the ground state.

5.1 Harmonic potentials on the lattice

The continuum Hamiltonian of a harmonically trapped gas of fermions with two components and contact interactions is given by

$$\hat{H} = \int \mathrm{d}x\, \hat{\psi}^\dagger_\sigma(x) \left(\frac{-\nabla^2}{2m} + \frac{1}{2} m\omega^2 x^2 \right) \hat{\psi}_\sigma(x) + g \int \mathrm{d}x\, \hat{n}_\uparrow(x) \hat{n}_\downarrow(x) \,. \tag{5.1}$$

Here, $\hat{\psi}^\dagger_\sigma$ and $\hat{\psi}_\sigma$ are creation and annihilation operators of fermions in spin states $\sigma \in \{\uparrow, \downarrow\}$ respectively. The particle number density operators are given by $\hat{n}_\sigma = \hat{\psi}^\dagger_\sigma \hat{\psi}_\sigma$. In what follows, we set $m = 1$ in lattice units, leaving the trap frequency ω and lattice coupling λ as the key parameters for the lattice system. Note, that the coupling is affected by cutoff effects and must be tuned to match the continuum theory. The given Hamiltonian is put on a rectangular spatial grid of size L and spacing a with $N_x = L/a$ sites and periodic boundaries. In units of the lattice spacing, the discretized Hamiltonian then reads

$$\hat{H} = \sum_{p,\sigma} \epsilon_p \hat{\psi}^\dagger_{p\sigma} \hat{\psi}_{p\sigma} + \sum_{x,\sigma} \frac{1}{2} \omega^2 x^2 \hat{n}_{x\sigma} + \sum_x \lambda \hat{n}_{x\uparrow} \hat{n}_{x\downarrow} \,, \tag{5.2}$$

where $\hat{\psi}^\dagger_{p\sigma}$ and $\hat{\psi}_{p\sigma}$ are creation and annihilation operators for fermions with momentum p. The dispersion relation used here is $\epsilon_p = p^2/2$, corresponding to that of free particles. For a brief discussion on the choice of dispersion relation, see Section 5.3.1. By inclusion of the frequency ω, we introduce an additional length scale into the system, known as the harmonic oscillator length given by

$$L_T = \frac{1}{\sqrt{\omega}} \,, \tag{5.3}$$

To mitigate finite size and volume effects, the oscillator length must be larger than the inter-site spacing while remaining significantly smaller than the overall extent of the lattice, i.e., $1 \ll L_T/a \ll N_x$. In practice, the lattice size required for good results also depends on the system's temperature and number of particles. Larger lattices are needed to resolve the wave functions of higher energy levels, which have a larger spatial extent.

Most aspects of the numerical procedure such as the Hubbard-Stratonovich transformation, weight construction, stabilization, and sampling methods are largely consistent with what is discussed in Chapter 2. The primary distinction lies in the construction of the matrix product for U, a crucial element in the simulation process. Each component of this matrix product becomes

$$B_n = e^{-\Delta t T} V_I e^{-\Delta t V_T} \,, \tag{5.4}$$

where V_I indicates the field dependent interaction potential, while V_T is the trapping potential. They are applied simultaneously, as both terms are diagonal in position space, and Fourier acceleration is used as before.

In order to tune the lattice to the continuum coupling, we use the ground state energy of the two-particle system as a reference. An exact solution for the two-body problem is available [144,

145], providing a relation between the continuum coupling g and the two-body binding energy E, namely

$$g(E) = \frac{\sqrt{2}(E-1)\Gamma\left(1-\frac{E}{2}\right)}{\Gamma\left(\frac{3}{2}-\frac{E}{2}\right)}. \tag{5.5}$$

To determine the appropriate bare lattice coupling, we adjust it to match the exact ground state energy of the two-particle system, as given by the relation above. We do not rely on Monte Carlo simulations to fine-tune the interaction strength. In the case of the harmonic trap in one dimension, lattices are generally small enough to allow for the exact diagonalization of the two-body lattice Hamiltonian. When dealing with larger lattices (> 100 spatial sites), we use a central frame, which introduces a small error. Consequently, we compute the two-body ground state energy using Monte Carlo simulations to ensure that any errors remain within statistical precision.

5.2 Complex Langevin for trapped fermions

This section is based on [3].

Applications of the complex Langevin method to systems of non-relativistic field theories, both bosonic and fermionic, have garnered significant interest in recent years. For the former, the literature offers investigations of quantum vortices [83], spin-orbit coupling [82], as well as the Berezinskii–Kosterlitz–Thouless transition [84, 85]. In the context of fermions, a sizeable body of work has been developed, including studies of the thermal properties of the spin-imbalanced unitary Fermi gas [32, 33], the ground state and thermal properties of spin and mass imbalanced fermions in a single spatial dimension [50, 81], as well as approaches to computations in the pairing decoupling channel and direct computation of particle number projections [53, 66].

In this section, we use the complex Langevin approach to study trapped fermions, focusing on the one-dimensional case. Additionally, we provide a general discussion of the method's applicability to spin-imbalanced systems. In the following we will primarily use bare lattice parameters for simplicity. Subsequent work will offer further investigations using physical parameters and direct comparisons to experimental results.

5.2.1 Regularization

Since the Langevin process requires a continuous field variable, we use a Hubbard-Stratonovich transformation that yields a continuous scalar field in density channel (2.35). As noted in [81], the sine terms produce unstable directions upon complexification. A Gaussian regulator term may be included in the drift force to counteract these instabilities, arranging for trajectories that do not wander too far out in field space. Consequently, the action becomes $S \to S + \xi\phi^2$, which results in the complexified Langevin equations

$$\partial_t \mathrm{Re}\{\phi\} = \mathrm{Re}\{\partial_\phi S[\phi] + \xi\phi\} + \eta, \tag{5.6}$$

$$\partial_t \mathrm{Im}\{\phi\} = \mathrm{Im}\{\partial_\phi S[\phi] + \xi\phi\}. \tag{5.7}$$

For all simulations, including those without a sign problem, we use a regulator with a strength of $\xi = 0.01$. This value is sufficiently small to have a negligible effect on the result, as confirmed by extrapolations for various cases.

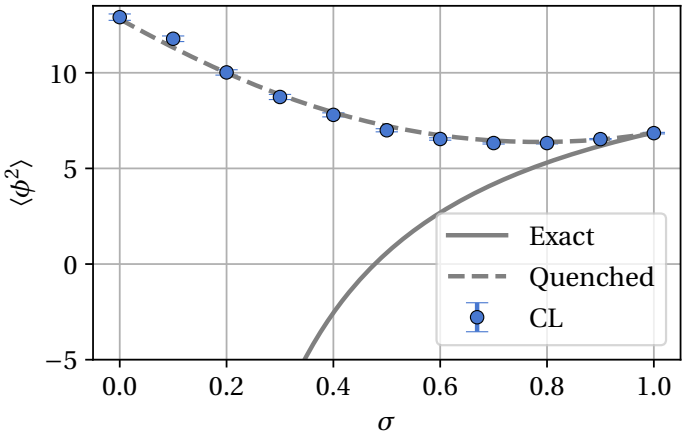

Figure 5.1: Second moment of the field for the sine toy model. The solid line indicates the exact result, while the dashed line is the phase-quenched expectation value. The Langevin simulations are shown as blue points and agree with the phase-quenched result up to artifacts from the finite step size.

5.2.2 Reliability considerations for complex Langevin

In Chapter 3, we have already discussed various pathologies of the complex Langevin method. Here, we will focus on a specific issue arising when the weight is real-valued but not necessarily positive. This scenario is encountered in the study of spin-imbalanced Fermi gases in the density channel, as widely used in the literature [32–34]. In a different context, it has been observed that the complex Langevin approach can produce phase-quenched results when dealing with a weight that is not complex, but still faces a sign problem [79]. To illustrate the problem, let us consider a simplified model with a real-valued weight that depends only on a single field $\phi \in [-\pi, \pi)$, given by

$$P(\phi) = e^{-S(\phi)} = \sin\phi + \sigma\,, \tag{5.8}$$

where σ is a real-valued offset parameter. For $\sigma = 0$, the weight has equal positive and negative contribution, and the partition function vanishes. Increasing σ will increase the positive contributions until the weight becomes entirely positive at $\sigma = 1$, so that no sign problem is present anymore. Evaluating the Langevin drift force, we find

$$\partial_\phi S = \frac{\cos\phi}{\sin\phi + \sigma}\,, \tag{5.9}$$

which has poles at the zeros of the weight $\sin\phi = -\sigma$. Poles in the drift immediately raise questions of ergodicity, and can spoil results in the limit of zero step size [146]. This is not an issue since we keep the step size large enough for the stochastic process to jump over the singularities.

After running the toy model simulation with a fixed step size of $\epsilon_L = 1.0 \times 10^{-3}$ and evolving to $t_{\max} = 2000$ in Langevin time (after thermalization), we calculated the second moment of the field as displayed in Figure 5.1. We opt for a small step size because larger ones tend to underestimate the expected value for intermediate σ. Hints of this behavior are still visible in our computation.

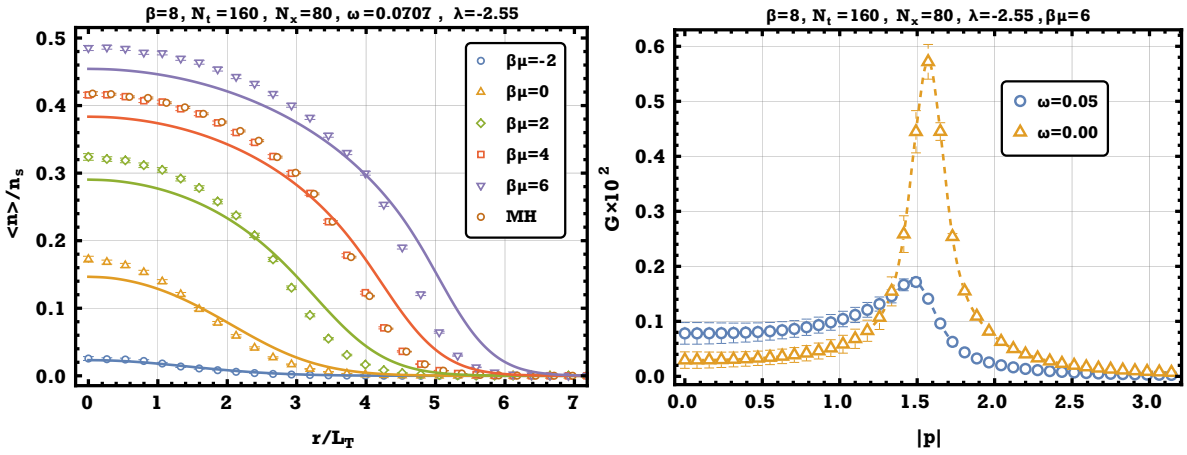

Figure 5.2: LHS: Density distribution for various chemical potentials and a comparison to Metropolis results, which are shifted for visibility. The position is given in length scales of the trap, while the density is normalized by its saturation value. Solid lines indicate the free case with matched average particle number. RHS: Connected density-density correlator of the trapped compared to the untrapped system at constant chemical potential with momenta in the Brillouin zone.

The results indicate that we are sampling from the phase quenched weight. This is unsurprising, as the Langevin process does not differentiate between positive and negative contributions. Indeed, for any positive weight $\exp(-S)$, we can find a corresponding negative weight $\exp S + i\pi$, which will evaluate to the same drift force. Moreover, the points at which phase changes occur represent lower-dimensional sub-manifolds in field space and are thus not visited by the Langevin process in practice. Consequently, the phase plays no role, and the sampled result must be the phase-quenched one. One might argue that this could be circumvented by initializing the stochastic process in the complex plane. However, our computations show the same behavior for such initializations.

Although these considerations limit the usefulness of Langevin simulations for spin-imbalanced fermions, they do not necessarily invalidate previously obtained results. As we will see in Chapter 6, the phase-quenched expectation values can be very close to the exact ones, especially when considering collective observables such as density.

5.2.3 Attractive systems and pairing

We first test the standard Langevin approach for an attractive, trapped system without a sign problem. The following takes $m = 1$ in lattice units for all cases. We use an inverse temperature of $\beta = 8$, with $N_t = 160$ time slices, leading to a $\Delta t = 0.05$ Trotter time discretization. The average Langevin step size is $\epsilon_L = 10^{-2}$ for all cases, and we employed an adaptive step size to prevent runaways [86]. This is sufficient for our results to converge within the statistical error. Our spatial lattice consists of $N_x = 80$ points, which sits comfortably above the length scale of the trap ($L_t = 3.76$) and the thermal wavelength ($\lambda_T = 7.1$), both in lattice units.

On the left-hand panel of Figure 5.2, we compare the density profiles in the trap with their free[1] counterparts for different chemical potentials. To achieve the same average particle content as in the interacting system, we adjusted the chemical potential for the non-interacting one accordingly. Our computations reveal an enhancement of the peaks at the center due to the attractive interaction. Because the average particle content is low, the effect is barely visible for $\beta\mu = -2$. Moreover, we compare the Langevin results for $\beta\mu = 4$ to runs performed via Metropolis sampling with global proposals and find that the density profiles agree within error. Note that the Langevin result was sampled with a finite regulator and finite step size, without performing an extrapolation. Hence, some deviations are expected, given the extensive statistics. On the other hand, we can rule out ergodicity issues since the Metropolis never sees a sign change in the determinants, indicating the absence of boundaries in the relevant regions in configuration space.

In addition to the density profile, we compute the connected density-density correlator, also called shot-noise in the experimental context,

$$G_{\uparrow,\downarrow}(k) = \langle \hat{n}_\uparrow(k)\hat{n}_\downarrow(-k)\rangle - \langle \hat{n}_\uparrow(k)\rangle\langle \hat{n}_\downarrow(-k)\rangle\,. \tag{5.10}$$

The plot on the right side of Figure 5.2 shows the shot-noise correlator for a system at fixed chemical potential with and without a trapping potential. In the untrapped case, a distinct peak at the Fermi momentum of the system is visible, indicating the appearance of BCS-type pairing around the Fermi surface. When a harmonic trapping potential is present, the peak appears less distinct but does not fall off as steeply towards smaller momenta. Since we do not work in physical units here, and the temperatures and particle numbers are not directly comparable, we cannot make definitive statements about pairing differences between trapped and untrapped systems.

For a discussion of the systems in the presence of a finite population imbalance, we refer to Section 5.3.5.

5.2.4 Density profiles for repulsive systems

In repulsive systems, there is a sign problem, specifically a complex weight problem. This is the regime where standard algorithms fail, and we resort to complex Langevin simulations for possible solutions. The density profiles for different interaction strengths, from small ($\lambda = 0.85$) to moderate ($\lambda = 1.7$), are shown in Figure 5.3. We use a trotterization step size of $\Delta t = 0.05$, an inverse temperature of $\beta = 8$, a spatial lattice size of $N_x = 80$, and a trap length of $L_t = 3.76$. In contrast to the attractive system, we observe a flattening and outward displacement of the distribution compared to its free, particle-content matched counterpart. This effect becomes more pronounced with stronger repulsion.

Since the stochastic process now explores the complex plane, we need to monitor the occurrence of potential boundary terms that may spoil convergence toward the correct result. Indeed, as was found in previous studies, repulsive interactions induce some degree of slow decay in the field and drift distributions at any coupling strength. However, while slow decay can indicate possible wrong convergence, no one-to-one correspondence exists. In systems at zero temperature with moderate interaction strengths, complex Langevin results have aligned with those from a formulation using hard-wall bosons, despite slow decay [147, 148]. Further investigations are necessary to

[1] We use the terms non-interacting and free interchangeably. "Free" does not indicate the absence of a trapping potential.

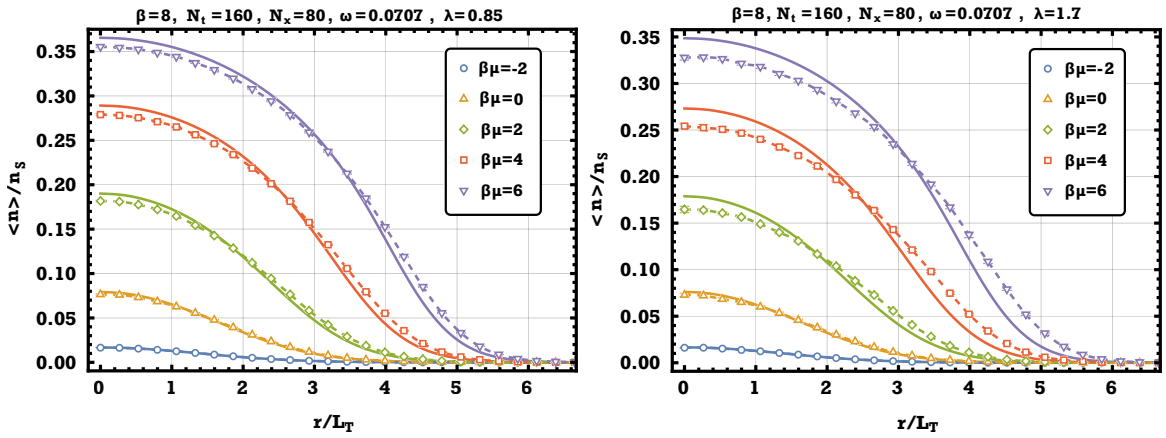

Figure 5.3: Density profiles for small (left) and moderate (right) repulsive interactions at various chemical potentials in the presence of a harmonic trapping potential. The dashed lines are spline fits to the data, while solid lines indicate the free case with matched average particle number

understand the full impact of boundary terms in these systems and whether correction terms are needed.

Finally, the repulsive case can easily be extended to imbalanced mixtures of spins, since there is no qualitative change in the nature of the phase problem under investigation. We rewrite the chemical potentials using

$$\mu = \frac{\mu_\uparrow + \mu_\downarrow}{2}, \qquad h = \frac{\mu_\uparrow - \mu_\downarrow}{2}. \tag{5.11}$$

In Figure 5.4, we show results for a system with repulsive interaction $\lambda = 1.7$, an imbalance of $\beta h = 2$ and $\beta = 32$. The spatial lattice size is $N_x = 40$, and the trap frequency is $\omega = 0.0707$. The left panel displays the position space density profiles for various chemical potentials, while the right panel shows the local polarization. We observe a maximum at a finite radius, which moves further out as the filling increases. It is worth noting that the bump in the majority profile is not solely due to the repulsive interaction. Instead, it partly originates from the oscillations appearing in density profiles at lower temperatures. Some further caution is required here since low temperatures can lead to a significant separation of scales and consequently to a loss of precision when calculating the drift force and observables.

5.2.5 Summary & Outlook

In this section we have discussed the application of complex Langevin simulations to trapped non-relativistic fermions in situations with attractive and repulsive couplings. Our investigation focused on one-dimensional systems, for which we have computed density profiles and correlations. We found that systems with repulsive interactions present a significant sign problem, and have demonstrated that complex Langevin provides a way to tackle this issue. Our observations indicate that density measurements are well-behaved within specific parameter ranges. However, incorrect convergence cannot be ruled out due to the sub-exponential decay in the distribution of the drift and observables. Notably, the case of spin and mass imbalance with attractive interactions

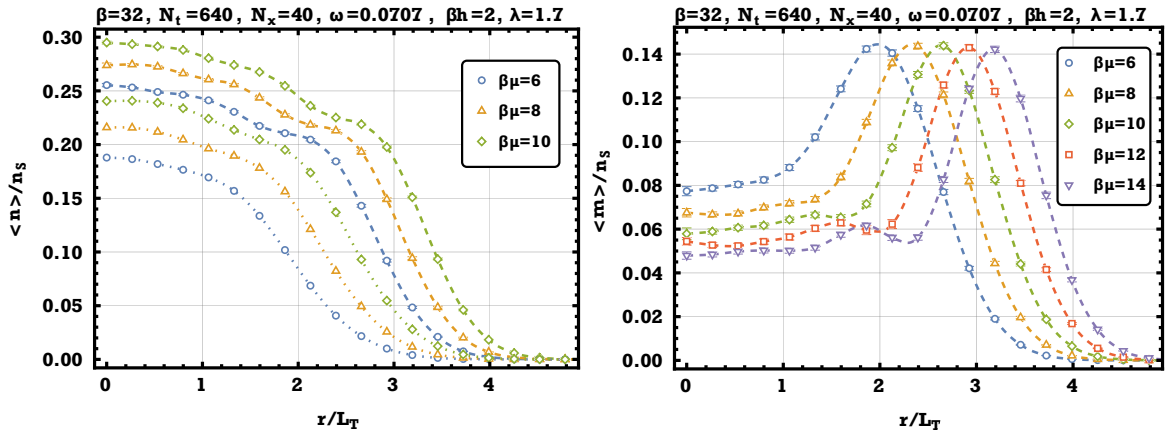

Figure 5.4: LHS: Density profile of spin imbalanced systems at various particle contents. The dotted lines are the minority, the dashed lines the majority species. RHS: Local polarizations. The lines are spline fits to the data.

requires further investigation, as it represents a scenario where complex Langevin simulations are known to face challenges. As an extension, we will apply lattice approaches to ground state and fixed particle number systems in the next section with more conventional Markov chain methods.

5.3 Monte Carlo computations from the ground state to finite temperature

This section is based on [1].

This section focuses on Metropolis-like Monte Carlo simulations of trapped fermions with attractive contact interactions in one dimension. We use the stabilization and truncation methods introduced in Chapter 4, which enable computations for both canonical and grand canonical systems at any temperature, down to the ground state. Our analysis includes density profiles, correlations, and separation energies. We also explore the effects of spin species population imbalance. While the latter can introduce a sign problem, we do not encounter one in the parameter ranges studied here.

Our results are structured as follows. In Section 5.3.1, we provide a brief discussion on the lattice size and parameters. Section 5.3.2 presents results for density profiles and correlations in systems with balanced population. In Section 5.3.4 we provide a comparison to exact diagonalization results for separation energies and compute the temperature dependence of the pairing gap for various particle numbers. Finally, in Section 5.3.5, we discuss the sign problem and provide a tomographic picture of density-density correlations in the imbalanced system.

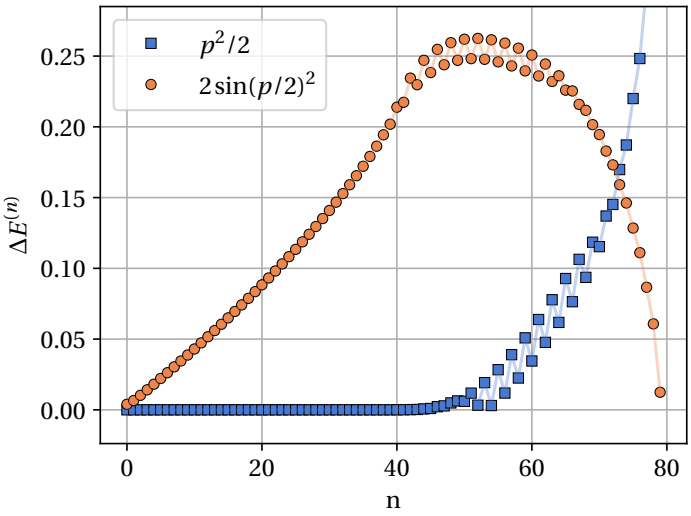

Figure 5.5: Relative deviation of the energy levels with p^2 dispersion (orange discs) and Hubbard dispersion (blue squares) in the non-interacting trapped lattice system compared to the continuum theory.

5.3.1 Technical setup

In the following, we use a discrete Hubbard-Stratonovich transformation in density channel, specifically

$$e^{-\Delta t\lambda \hat{n}_{i\uparrow}\hat{n}_{i\downarrow}} = \frac{1}{2}\sum_{\phi_i=\pm 1} e^{(\gamma\phi_i - \Delta t\lambda/2)(\hat{n}_{i\uparrow}+\hat{n}_{i\downarrow}-1)}, \tag{5.12}$$

where $\gamma = \cosh(\gamma) = e^{\Delta t\lambda/2}$. The discrete auxiliary field does not allow sampling via drift-based methods, and we opt for the time local update discussed in Section 4.2.2 instead. We have found that this method reduces autocorrelation times compared to global updates with a continuous field.

An essential aspect of lattice simulations is the choice of parameters to ensure rapid convergence towards the continuum limit. In our case, we need to consider several variables. One of these is the temporal lattice spacing a_t, which dictates the number of time slices used and thus influences the error of the Trotter decomposition. We find $a_t/a = 0.05$ to be sufficiently small to ensure a negligible Trotter error.

The ratio L_t/a governs the finite size and finite distance errors, and we set it to $L_t/a = 4$. Simultaneously, the total number of spatial sites is $N_x = 80$, corresponding to a box size of $N_x/L_T = 20$ in units of the harmonic oscillator. For the system with up to $N = 20$ particles, corresponding to a total filling of $N/N_x = 0.25$, we use these values unless stated otherwise. However, it is worth noting that a higher number of particles necessitates a larger spatial extent and smaller spacing to avoid finite size and filling effects. In particular, one must ensure that the trap's center is sufficiently below the saturation density, which is the case in all our computations. In Section 5.3.3, we consider systems with up to $N = 80$ particles, for which we use an $N_x = 200$ lattice and $L_t/a = 8$, leading to $N_x/L_t = 25$ and a filling of $N/N_x = 0.4$. The proximity to the continuum theory is evidenced by

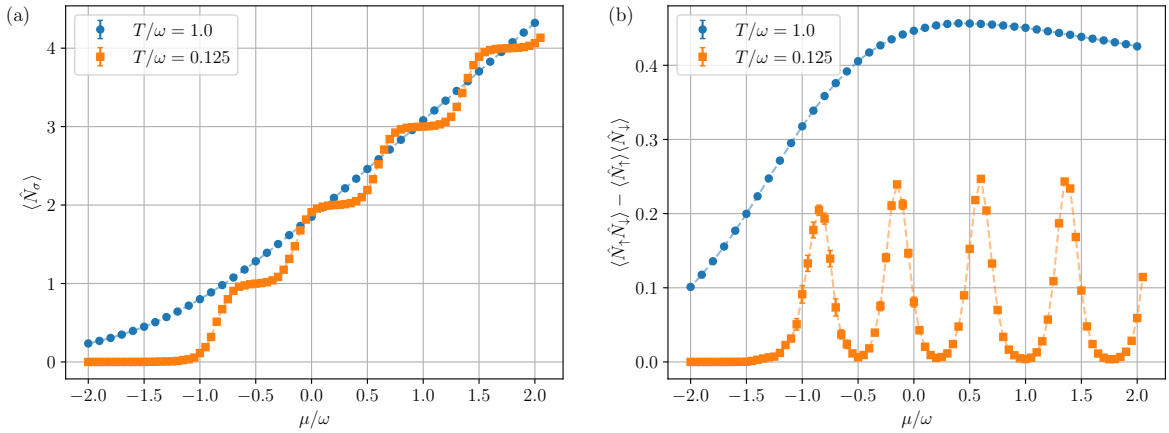

Figure 5.6: (a) Expectation value of the particle number operator as function of the chemical potential $\mu = \mu_\uparrow = \mu_\downarrow$ for different temperatures at $g/\sqrt{\omega} = -3$. (b) Correlation between the particle number operators for each spin species as function of the chemical potential $\mu = \mu_\uparrow = \mu_\downarrow$ for different temperatures at $g/\sqrt{\omega} = -3$. Dashed lines are drawn to guide the eye.

the excellent agreement with previous results, e.g., in Figure 5.11 and Figure 5.10. Additionally, we have compared the density profiles for $N = 20$ fermions for various lattice parameters and found them to fall on universal curves, see Figure 5.10.

A note on the particle dispersion

Besides the discretization and bare parameters, the choice of the particle dispersion relation is crucial for the accuracy of results. In this work, we use

$$\epsilon_p = \frac{p^2}{2}, \tag{5.13}$$

for the kinetic energy of the lattice system. Another form often used is given by the finite difference approximation of the derivatives, yielding

$$\epsilon_p = 2\sin^2\left(\frac{p}{2}\right). \tag{5.14}$$

This is the standard dispersion of the Hubbard model and was previously employed for the trapped system in, for example, [48]. To perform a simple check, we diagonalize the non-interacting lattice Hamiltonians with both dispersions and compare the resulting energies to the continuum theory. We use the smaller lattice with $N_x = 80$, $\omega = 0.0625$. Figure 5.5, depicts the relative deviation of the energy levels for both cases given by

$$\Delta E^{(n)} = \frac{|E^{(n)} - E^{(n)}_{\text{cont}}|}{E^{(n)}_{\text{cont}}}. \tag{5.15}$$

The comparison shows that the quadratic dispersion agrees significantly better with the continuum result, with only small deviations up to half-filling. In contrast, the Hubbard dispersion deviates even for the lowest-lying states. Both cases exhibit an even-odd effect in higher shells. Note

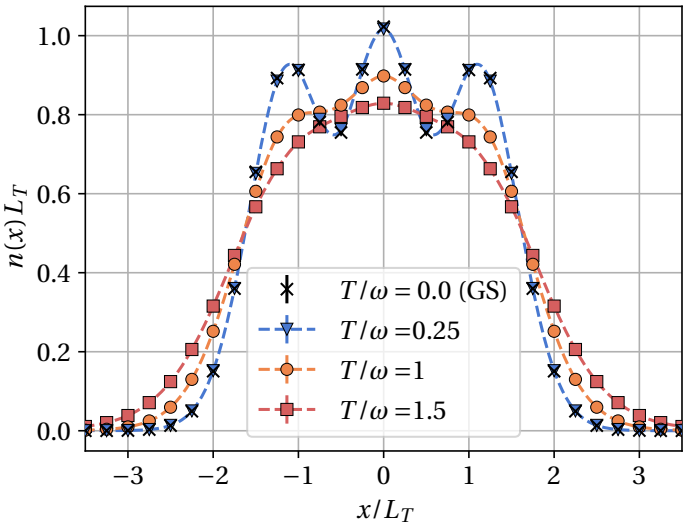

Figure 5.7: Density profile per spin species $n(x) = n_\uparrow(x) = n_\downarrow(x)$ for $3+3$ particles and $g/\sqrt{\omega} = -3$ at three different temperatures. Dashed lines are splines through the data points and crosses indicate data from projection to the ground state. The peaks characteristic for the ground state gradually disappear as the temperature increases beyond the level spacing of the trapping potential.

that this issue does not affect our simulations, as they are conducted in more dilute regimes. Moreover, these results suggest that using the quadratic dispersion is necessary to obtain accurate results for observables such as the total energy or separation energies. It is worth mentioning that methods which keep the energy levels exact are available, but they require a non-uniform spatial lattice [60, 61].

5.3.2 Balanced spin species

In the case of spin-balanced systems, the determinants in (2.46) are real and equal since $\mu = \mu_\uparrow = \mu_\downarrow$ and thus no sign problem is present. This is true for the density channel Hubbard-Stratonovich transformation used in this section and the spin-z channel not applied here, where the determinants are complex conjugates of each other.

Before making use of the particle number projection methods described in Section 2.4, we compute the expectation value of the particle number operator $\hat{N}_\sigma = \int dx\, \hat{n}_\sigma(x)$ as a function of the chemical potential for different temperatures, along with the correlation between up and down spin number operators, $\langle \hat{N}_\uparrow \hat{N}_\downarrow \rangle - \langle \hat{N}_\uparrow \rangle \langle \hat{N}_\downarrow \rangle$.

Figure 5.6 (a) depicts clear steps at nearly integer particle numbers for the lowest temperature of $T/\omega = 0.125$, indicating the thresholds in the chemical potential where each energy level is filled. This effect is smoothened by thermal fluctuations, as can also be seen. Moreover, the attractive interaction between up and down spins causes the departure from the empty system at negative chemical potentials. Conversely, the non-interacting theory would have this threshold around $\mu/\omega = 1$ for small but finite temperatures.

Figure 5.6 (b) shows the (connected) correlation between $\hat{N}_\uparrow$ and $\hat{N}_\downarrow$. Similar to what we saw for

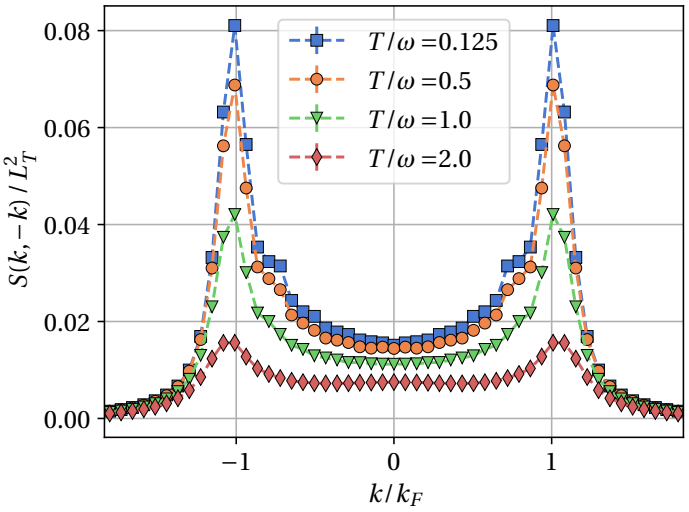

Figure 5.8: Density-density correlation function for opposite momenta and $10+10$ particles at $g/\sqrt{\omega} = -3$ and four different temperatures. The BSC peaks at the Fermi are visible in all cases, but become less pronounced with increasing temperature.

the average particle number, this correlation function for $T/\omega = 0.125$ exhibits repeating patterns following the filling of the different energy levels. The behavior is absent at high temperatures, and the correlator changes more smoothly with the chemical potential.

Turning to the canonical ensemble, we investigate the density profile of a fixed number of particles in the trapping potential. The ground state exhibits distinct particle peaks that stem directly from the wave functions of the harmonic oscillator states. These oscillations are already present in the ground state of the non-interacting system $g = 0$, where the density is determined by $n^{(0)}(x) = \sum_{n=0}^{N_\sigma - 1} |\psi_n(x)|^2$, using the well-known harmonic oscillator wavefunctions $\psi_n(x)$. The density peaks are thus not necessarily a sign of pairing but rather a direct consequence of the harmonic potential.

In Figure 5.7, the density profiles $n(x) = n_\uparrow(x) = n_\downarrow(x)$ for $g/\sqrt{\omega} = -3$ with $3+3$ particles are shown for various temperatures. In the ground state, the density distribution displays three peaks equal to the number of filled oscillator shells. We find the same behavior at temperatures significantly smaller than the spacing between energy levels in the trapping potential, since the ground state dominates in these cases. As the temperature is raised to the point where the thermal energy becomes comparable to the energy gap, the peaks vanish, and the density profile gradually becomes smoother, highlighting the significant impact of temperature on the density profile. With larger particle numbers, the amplitude of the particle peaks is expected to decrease, and they become more frequent. Eventually, the density converges towards a smooth profile in the thermodynamic limit.

As an indicator for the existence of pairing in the system, we compute the connected density-density correlation function in momentum space, given by

$$S(k,k') = \langle n_\uparrow(k) n_\downarrow(k') \rangle - \langle n_\uparrow(k) \rangle \langle n_\downarrow(k') \rangle \,. \tag{5.16}$$

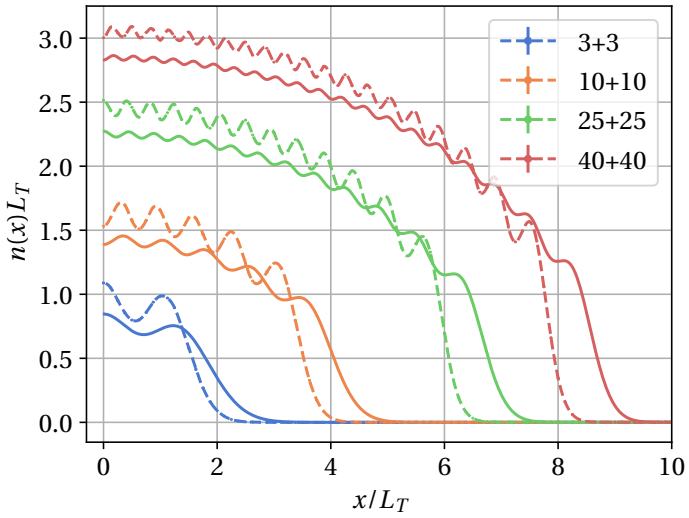

Figure 5.9: Density profiles per spin species for $N = 6$, 20, 50 and 80 particles at $g/\sqrt{\omega} = -5$ and $T/\omega = 0.25$. Splines connecting the data points (dashed lines) are drawn to guide the eye.

Pairing around the Fermi surface will manifest as positive correlation peaks at $S(\pm k_F, \mp k_F)$. In Figure 5.8 we present results for the system at $g/\sqrt{\omega} = -3$ and $S(k,-k)$ for various temperatures. Close to the ground state, the peaks around the Fermi surface are more pronounced. The correlations decrease rapidly towards larger momenta, but they do not disappear at vanishing opposite momenta due to the finite particle number and the resulting finite size of the systems. We observe a weakening in correlations, similar to the breakup of ground state features in the density profile when temperatures become comparable to the level spacing of the trapping potential. Moreover, the overall picture appears more flat. However, in contrast to the oscillations in the density profile, the density-density correlations at lower temperatures remain clearly visible.

5.3.3 Few to many

To show the applicability of our approach to systems of larger particle numbers, we compute the density profiles and energies for systems of up to $N = 80$ particles. To this end, we adjust the lattice parameters to $N_x = 200$ and $\omega = 0.015625$, corresponding to $L_T/a = 8$. In Figure 5.9, we plot the density profiles $N = 6$, 20, 50 and 80 particles at $g/\sqrt{\omega} = -5$ and $T/\omega = 0.25$, as well as the ground state densities of the corresponding non-interacting systems. The larger number of lattice sites and smaller lattice spacing allow for a satisfactory resolution of the density oscillations close to the ground state. We compare the energies obtained at different particle numbers to verify the computations further. In [140], the authors found the energy, normalized by the non-interacting energy, to only show very weak dependence on the particle number when plotted against the rescaled coupling

$$\gamma = \frac{\pi g}{\sqrt{\omega N}}. \tag{5.17}$$

This behavior is reproduced by our data, as seen in Figure 5.10, where we compare the energies against those of the two-particle system in the ground state. The lattice results generally slightly

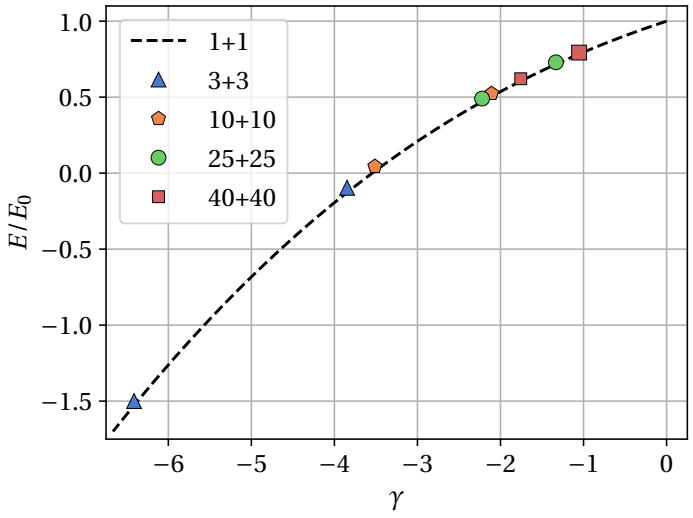

Figure 5.10: Energy of the system, normalized by the non-interacting energy, as function of the rescaled coupling γ. The energies of the two-particle system in the ground state are plotted as a dashed line. All points fall on an almost universal curve, with the lattice results slightly overshooting the two-particle system.

overshoot the two-particle system, which is expected at higher particle numbers but may also partially result from the finite temperature.

Our current approach faces limitations when studying systems with strong couplings due to ergodicity issues in the Markov process. These issues complicate the tuning of particle numbers, particularly at very low chemical potentials. A similar challenge arises in the Hubbard model at half-filling, where ergodicity problems can be alleviated by implementing global updates [149]. To address these limitations, alternative strategies could be employed. One possibility is to sample the canonical weight more directly or to perform computations in a different decoupling channel. Preliminary tests suggest that spin channel computations are less prone to being trapped in sectors with fixed particle numbers compared to density channel computations.

5.3.4 Energy observables

Having explored density observables and their correlations, we now focus on energy-related quantities. Specifically, separation energies have been experimentally measured [150] and theoretically analyzed through exact diagonalization [138, 141]. These measurements serve as a valuable benchmark to assess the accuracy of our computations near the ground state. The separation energy can be understood as the interaction energy cost of adding a single particle to the system and is defined in terms of the ground state chemical potential of the system,

$$\mu(N) = E(N) - E(N-1)\,, \tag{5.18}$$

$$\Delta_S(N) = \mu(N) - \mu^*(N)\,, \tag{5.19}$$

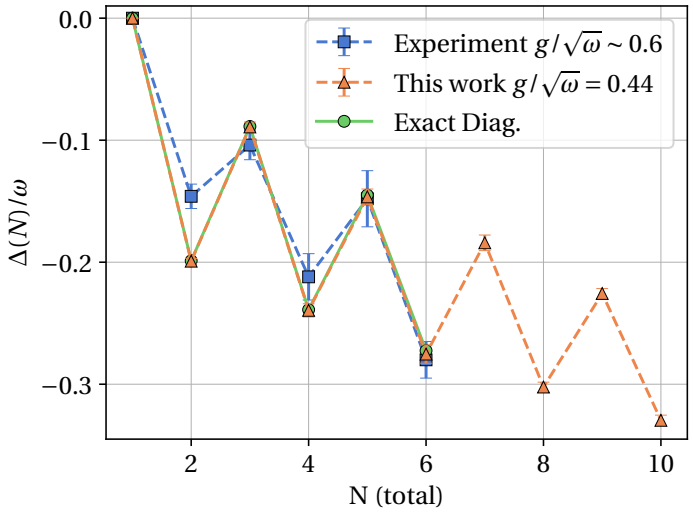

Figure 5.11: Separation energies in the ground state for up to $5+5$ particles. Experimental data [150] (blue squares) and exact diagonalization results [138] (green circles) are compared to results from finite temperature lattice simulations at $T/\omega = 0.125$ (orange triangles). We achieve excellent agreement with exact diagonalization and are able to extend the comparison to larger particle numbers.

where $\mu^*(N)$ is the chemical potential of the free system. Note that this is not the tuning parameter seen in the grand-canonical case and must be computed instead.

At low temperatures, direct simulations of the imbalanced systems are necessary because the overlap between the balanced canonical simulation and the imbalanced sector becomes small. In Figure 5.11, we see strong agreement between our results labeled as "Lattice", calculated at a finite temperature of $T/\omega = 0.125$, and results obtained using exact diagonalization from [138]. The step-like pattern, with lower energies for even particle numbers indicating shell closures, suggests the presence of pairing between particles of up and down spin. We also note the proximity to experimental values at a slightly different coupling. Some deviations are visible, likely due to anharmonicity of the trapping potential present in the experiment. While beyond the scope of this work, our method generally allows the study of an arbitrarily shaped external potential, which opens up avenues to investigate the effects of anharmonicity in a controlled manner. For the comparison to exact diagonalization results, we match both two-body energies to the same ground state values, resulting in a value for the coupling slightly smaller than what was given in [138]. While the results in Figure 5.11 are at weak coupling, similar agreement with exact diagonalization is found in more strongly coupled scenarios.

Next, we compute the energy staggering pairing gap, defined for even particle numbers as

$$\Delta_P(N) = \frac{1}{2}[2E(N/2-1,N/2) - E(N/2,N/2) - E(N/2-1,N/2-1)]\,. \tag{5.20}$$

This quantity serves as an indicator for pairing in the system and has been used to study pseudogap effects [136, 151] and pair correlations [152] in higher dimensions. Although higher-order estimators are available [153], we use a three-point estimator here for simplicity. At large enough

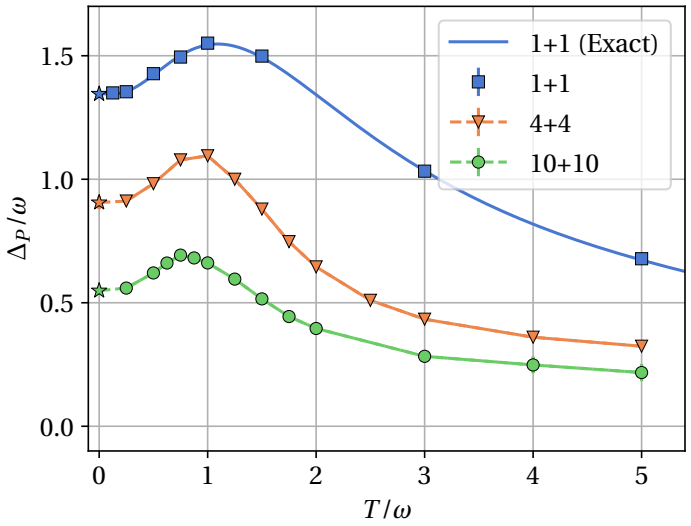

Figure 5.12: Pairing gap at $g/\sqrt{\omega} = -3$ for $1+1$, $4+4$ and $10+10$ particles at finite temperatures. The two particle system is compared to the exact result of the continuum theory (solid line). Stars indicate results from ground state projection.

temperatures, it is often enough to sample for a single chemical potential and reweight the different particle numbers from the same data. This approach breaks down around $T/\omega \sim 0.5$, where we perform independent simulations for the energies at different particle numbers.

Figure 5.12 depicts the pairing gap at $g/\sqrt{\omega} = -3$ for $1+1$, $4+4$, and $10+10$ particles at finite temperature. For two particles, we agree with the continuum theory's exact result up to $T/\omega = 5$. We see an increase in the gap when temperature increases before it decays again. Evidently, the energy of the single particle contribution increases more rapidly than that of the two particle one for small temperatures. The agreement between the two-particle systems and the exact result indicates the validity of the renormalization approach used to fix the coupling, where no excited states were considered. However, we expect the agreement to break down at higher temperatures as high-energy states beyond the limits of the lattice start to contribute significantly.

The systems with eight and twenty particles respectively follow a similar general trend to the two-particle one, but at lower total magnitude, consistent with available ground state results [140]. Interestingly, we still find peaks at finite temperatures, which move to lower temperatures when the particle number increases. It remains unclear, from our computations, whether the peak converges to a fixed temperature or vanishes in the thermodynamic limit, indicating a shell effect.

5.3.5 Imbalanced populations

An important open question in the study of ultracold fermionic gases is whether an exotic paring phase exists at finite spin imbalance. Such a phase, known as the FFLO phase, is characterized by pairing at non-degenerate Fermi surfaces, resulting in pairs with finite total momentum $k_{F\uparrow} - k_{F\downarrow}$.

However, the extent to which such systems can be studied using standard lattice methods is unclear. In the case of imbalanced spin species, the configuration weight is no longer necessarily

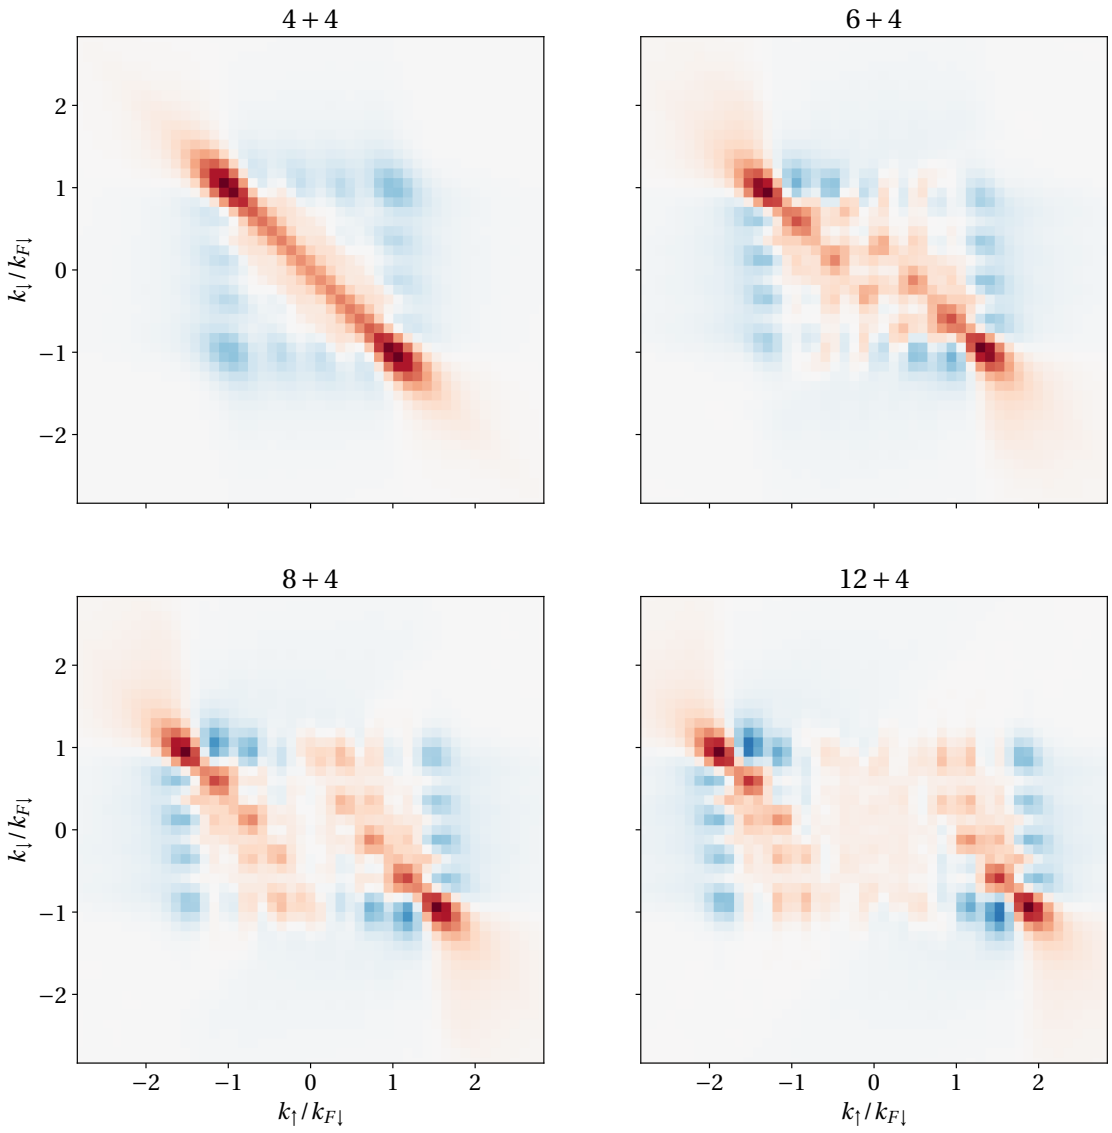

Figure 5.13: Density-density correlations of the system at $g/\sqrt{\omega} = -3$ and $T/\omega = 0.25$, with varying imbalance. We observe clear signals of unconventional pairing in the presence of an imbalance. To retain visibility, all color scales are normalized independently to span the full range of values in the corresponding data.

positive, and a sign problem can arise. This is the case in 2D and 3D simulations [35, 154]. In contrast, some previous studies of Fermi gases with contact interactions in 1D have not shown a sign problem in the considered parameter regions [3, 155, 156]. The trapped system studied here displays similar behavior. We find no negative weight configurations for any combination of temperature, chemical potential, and coupling. It is important to note that this is not an issue of ergodicity. Our code runs into the expected sign problems in 2D and 3D and reproduces known results in 1D, as has already been shown above in the case of separation energies, which require simulations at a slight imbalance.

We complement previous studies that consider trapped systems with imbalances [143, 156, 157], by computing pairing patterns at higher particle numbers than done before. In Figure 5.13, we present a visualization of the connected density-density correlations (5.16) for a system with $4+4$ to $4+12$ particles at $g/\sqrt{\omega} = -3$ and $T/\omega = 0.25$ in momentum space. The spin-balanced correlator is peaked at $k_\uparrow + k_\downarrow = 0$, while the imbalanced systems show a clear signal of pairing at finite momentum, which is consistent with the expected FFLO behavior. Note that Figure 5.13 does not show the relative magnitude of correlations at different particle numbers. Each plot is independently normalized to ensure visibility. Moreover, several pockets of positive and negative correlation are visible, forming an oscillatory pattern. Similar to the oscillatory pattern in the density profile, these oscillations are a feature of the harmonically trapped system and not found in the same way in the untrapped gas [50]. While a thorough analysis is left to future work, we generally find a decrease in correlations when going to higher temperatures and larger imbalance. This aligns with results from a recent exact diagonalization study of few-particle systems at finite temperature [143].

5.3.6 Summary & Outlook

We have conducted lattice simulations of trapped fermionic systems in one dimension with both balanced and imbalanced populations. The sampling was performed in the grand canonical ensemble with a reweighting process to obtain canonical expectation values. This method is more efficient than directly sampling the canonical weight but requires tuning the chemical potential. Our stabilization procedure enabled us to simulate the entire range of temperatures down to the ground state, which we then compared to a projective approach. We compared experimental and theoretical data for the separation energies of up to six particles to validate our results and found good agreement. Furthermore, we computed the energy staggering pairing gap for up to twenty particles and found agreement with exact results for two particles. The computation of separation energies, in particular, requires simulations for spin-imbalanced systems, which we find to be sign problem-free in the parameter ranges studied. This is also the case when computing density-density correlations in the presence of larger imbalances, where we find clear signals of unconventional pairing.

In future studies, it may be interesting to explore polaronic effects, as no sign problem appears to be present, allowing computations even at large imbalances. In contrast to, e.g., Path Integral Monte Carlo (PIMC) computations, the canonical lattice formulation used in this work does not incur a sign problem in higher dimensions as long as the population remains balanced. Given the good agreement of our lattice computation to previous results in, e.g., Figure 5.11, we expect the approach to generalize straightforwardly to higher dimensions. Indeed, using a very similar truncation approach as the one employed here, lattices for up to 75^2 were recently studied in the

two-dimensional untrapped gas [42], which is close to the linear extent we found sufficient for the 1D system.

5.4 Conclusion

In this chapter we have shown that a straightforward lattice approach is well suited for computations of trapped fermionic systems in one dimension. We were able to model both attractive and repulsive interactions. The former is made possible by complex Langevin, which allows for simulations even with a significant sign problem. However, it is important to note that the accuracy of our results is not guaranteed in all scenarios. We have observed slow decay in the sampling distribution towards infinity and around poles for all computations of the repulsive system. While this does not necessarily mean that complex Langevin has failed, it does require further investigation. Tools to address potential wrong convergence are available, but little success has been found when the errors originate from singularities instead of at infinity, as is the case here.

The Monte Carlo-based approach for attractive systems is free of the sign problem and facilitates exact computations. Even relatively small lattices of only 80 spatial points provide excellent agreement with exact results for separation energies. Moreover, we have demonstrated that the lattice dispersion chosen is crucial for the accuracy of results. In particular, a naïve discretization of the derivatives leads to significant deviations from the continuum theory even without interactions. In comparison, the quadratic dispersion we used gives accurate and stable results.

For the spin imbalanced gas, we have discussed the challenges faced by complex Langevin due to the nature of the sign problem. To the best of our knowledge, this had yet to be addressed in the literature. However, we found the sign problem to be wholly absent in the parameter ranges studied for the one-dimensional trapped gas, which allowed us to uncover signals of unconventional pairing. The absence of the sign problem is fascinating, as there is not a priori reason to expect it.

Many direct extensions of this chapter are worth pursuing. Chief among them is the application to higher dimensions, where other methods suffer from sign problems even for balanced populations. Applying complex Langevin to rotating systems in this context might also be possible. These are accessible experimentally [130], but challenging for Monte Carlo methods, due to the complex weight. Additionally, the complex Langevin simulations did not utilize the stabilization and truncation scheme and could be improved to handle lower temperatures by incorporating these methods. The severity of the sign problem in higher dimensions needs to be fully understood and warrants further review in light of the results presented here. This will be the focus of the next chapter, where we will discuss untrapped gases in two dimensions.

Chapter 6

Fermi gases in two dimensions

When the number of spatial dimensions is increased from one to two, a plethora of new phenomena and physical effects emerge. Chiefly among them is the appearance of a superfluid phase at low temperatures induced by a Berezinskii–Kosterlitz–Thouless transition. The two-dimensional case is exceptional in this regard, as it represents the marginal dimension beyond which superfluidity may appear. Although the Mermin-Wagner theorem rules out the possibility of a second-order phase transition and long-range order, as seen in three-dimensional systems, quasi-long-range order may exist. This is made possible by the formation of vortex-antivortex pairs, which generate algebraically decaying correlations. The two-dimensional system is of great interest, not least due to the proximity to condensed matter phenomena such as high-T_c [158] and topological superconductors [159].

For fermions with contact interactions in particular, the literature provides a rich body of work concerning, for example, the BKT transition in the BEC-BCS crossover [160–164], the density equation of state [134, 165–167], pairing properties [131], spin imbalances [30, 31] and many more. Lattice and Monte Carlo methods have been used extensively in the context of the Fermi-Hubbard model [168, 169], which describes a Fermi gas in the dilute limit. Moreover, direct applications to cold atoms systems, where the continuum limit is considered, have recently been performed [42, 136, 170].

This chapter focuses on lattice simulation for two-dimensional gases and explores their behavior under finite population imbalances. To start, in Section 6.1, we introduce the model, discuss the various scales, and describe how the lattice system can be tuned to match experimental parameters. In Section 6.2, we briefly review the balanced gas, which serves as a benchmark for our computations. Finally, Section 6.3 presents results for the imbalanced gas, focusing on the normal phase of the system. We also address the sign problem and its impact on the system and compare our results with those from the virial expansion.

6.1 Model parameters and scales

We first discuss the parameters and scales relevant to the grand canonical ensemble and the relation between physical and bare lattice parameters. For a two-dimensional system, the lattice

Hamiltonian, including the chemical potential, is given by

$$\hat{H} = \sum_{p,\sigma} \epsilon_p \hat{\psi}^\dagger_{p\sigma} \hat{\psi}_{p\sigma} + \sum_{x,\sigma} \mu_\sigma \hat{n}_{x\sigma} + \sum_x \lambda \hat{n}_{x\uparrow} \hat{n}_{x\downarrow} , \tag{6.1}$$

where μ_σ is the chemical potential for each spin species, λ is the interaction strength, and ϵ_p is the dispersion relation. The temporal extent β of the lattice controls the temperature. In cold atoms experiments, the interaction strength is typically expressed either as the two-body binding energy e_b or the s-wave scattering length a. For a two-dimensional system with $\hbar = 1$ and $m = 1$, the coupling parameters are related by [171]

$$e_b = \frac{4}{e^{2\gamma} a^2} , \tag{6.2}$$

where $\gamma = 0.5772\ldots$ is Euler's constant. While some literature omits this constant in defining the scattering length, we include it here to maintain consistency with previous work.

The only other relevant scale in the system is the Fermi energy given by

$$E_F = 2\pi n . \tag{6.3}$$

Here, n represents the total density per spin species when assuming balanced populations[1]. These parameters allow us to express any physical quantities in dimensionless form, facilitating direct comparison between experimental and lattice results.

6.1.1 Tuning the coupling

To match a lattice system to a physical one, a precise connection between the lattice coupling and the physical scattering length must be established. The literature offers two similar approaches, both of which yield slightly different results for coarse lattices.

The first approach used, for example, in [170] involves computing the two-body binding energy produced by the lattice Hamiltonian and then matching it to the desired two-body energy in the continuum. Subsequently, the scattering length can be computed using Equation (6.2). Assuming the lattice to be of infinite spatial extent, one finds for the binding energy

$$\frac{1}{\lambda} = \int_{-\pi}^{\pi} \frac{d^2 p}{(2\pi)^2} \frac{1}{p^2 + \epsilon_p} . \tag{6.4}$$

We have also tested matching to the binding energy on the finite lattice, which makes the momentum integral above discrete. However, for small lattices, the bound state can be significantly affected by periodic boundary conditions. We found that thermodynamic observables yield a cleaner infinite volume extrapolation when tuning the coupling with continuous integral above.

The second approach defines the scattering length directly from the lattice system after solving the two-body problem [172]. The resulting relation between the lattice coupling and the scattering length is given by

$$\frac{1}{\lambda} = -\frac{\log\left[a e^{\gamma - 2\beta(2)/\pi} \pi\right]}{2\pi} , \tag{6.5}$$

[1] In parts of the literature n is taken to be the sum of the densities of the spin species instead. The formulations differ by a factor of one half in the Fermi energy.

where $\beta(2) = 0.915\ldots$ is Catalan's constant.

Given the s-wave scattering length and the Fermi momentum, we may define the dimensionless coupling as

$$\eta = \ln k_F a\,, \tag{6.6}$$

where $\eta \gg 0$ corresponds to weak interactions and the BCS regime, $\eta \ll 0$ indicates strong interactions with BEC behavior, and $\eta = 0$ sits directly in the BEC-BCS crossover regime. Since the Fermi momentum is entirely determined by the density, the target lattice coupling is also dependent on this parameter. To reach the continuum limit, the density of lattice points is increased while keeping the total number of particles constant, thereby decreasing the density. Consequently, Equation (6.5) implies that as the scattering length increases, and the numerical value of the lattice coupling must decrease to maintain a constant dimensionless coupling.

We have run tests using both Equation (6.5) and Equation (6.4) to tune the coupling. While both converge to the same values for small binding energies and large scattering lengths, they differ at finite densities and with strong interactions. We used the more straightforward binding based approach for all our computations in the following. This is motivated by the excellent agreement with virial expansion results when comparing thermodynamic observables, as will be discussed in Section 6.3. The second approach, although accurate when approaching the continuum limit, shows some deviations from the virial expansion when the lattice couplings are too large. Our preference for the first method is thus justified by its more rapid convergence to the continuum.

Alternatively, in the grand canonical ensemble, a natural set of dimensionless parameters is defined via combinations of the inverse temperature and the binding energy or chemical potential. The quantities $\beta\mu$ and βe_b fully characterize the system. These parameters have been used in theoretical [32, 134, 165, 173] and experimental [166, 167] studies of two-dimensional Fermi gases, and can be tuned without effort in lattice computations. However, the density and scattering length parameters remain the more popular choice in experiments.

6.1.2 Tuning grand canonical parameters

When using dimensionless grand canonical parameters to characterize the system, we can immediately begin computations. However, to achieve a specific value for η in equation (6.6), we also need to set the density of the system. This can be challenging, as the density is not a direct input of the lattice Hamiltonian. Instead, it requires adjusting the chemical potential through multiple independent simulations to reach the desired value.

To lessen the burden of finding the correct value of μ for each set of parameters T/T_F and η, we have performed a sweep of parameters space for a fixed lattice size of $V = 29^2$. A grid scan of the binding energy and chemical potential allows us to compute η and the density for each point and to create an interpolated map back to the original parameters. Figure 6.1, shows a small region in the T/T_F-η space with the corresponding chemical potential, extrapolated from a finite set of points. From this plot, we can infer the dimensionless chemical potential. The density is determined by the temperature in lattice units and the target Fermi temperature, which are also provided. Finally, the coupling on the lattice is also entirely determined via either Equation (6.5) or Equation (6.4).

This approach, while effective, is only an approximation. It doesn't yield the exact point in coupling space due to the fixed lattice size and finite volume and density effects. However, it provides a strong starting point and minimizes the time-consuming tuning process.

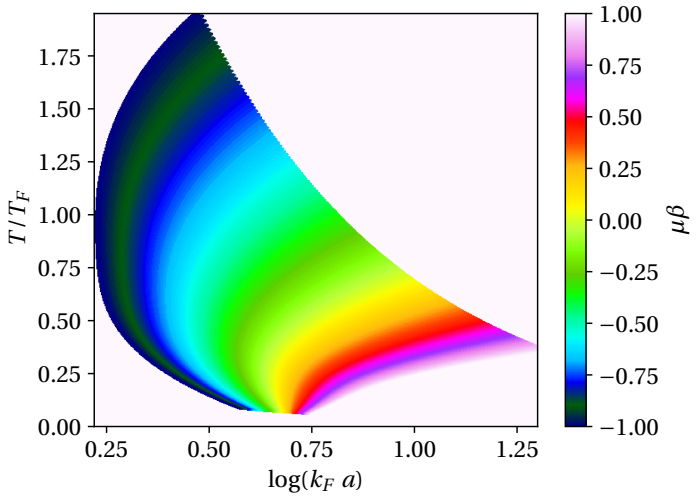

Figure 6.1: The dimensionless chemical potential as a function of the temperature in units of the Fermi temperature T_F and the common coupling $\log k_F a$. We use this plot to tune couplings in the given region. The color gradient represents an interpolation from a finite set of points obtained from a lattice of fixed size. The white spaces indicate the regions outside the boundaries of this specific example.

The next step is an iterative refinement based on observables measured directly from the lattice system. Given the initial guess, we choose a desired value for the density and dimensionless coupling, which we use to set the lattice coupling definitively. Now, λ and β are set, and what remains is to find μ so that a desired density is achieved. Since the chemical potential inferred from Figure 6.1 will not reproduce this density exactly, we compute the change in particle number with respect to the chemical potential as

$$\frac{\partial\langle N\rangle}{\partial\beta\mu} = \langle(N_\uparrow + N_\downarrow)^2\rangle - \langle N_\uparrow + N_\downarrow\rangle^2 = N_{\text{diff}}\,, \tag{6.7}$$

This result, which represents particle number fluctuations, is directly related to the compressibility of the gas. We now discretize the left-hand side of the equation to first order and obtain an expression for the change in chemical potential required to arrive at the desired particle number N_{target}, given by

$$\Delta\mu = \frac{N_{\text{target}} - \langle N\rangle}{N_{\text{diff}}}\,. \tag{6.8}$$

Updating the chemical potential like this over several steps leads to precise and efficient convergence toward the desired particle content. Our experience is that with suitable initial parameters, a single 'pre-run' is usually sufficient to tune the particle content to a precision within statistical error.

So far, we have assumed the spin species to be balanced, in which case the tuning is complete. However, the spin species' chemical potentials lose their degeneracy in the spin-imbalanced case. We use their difference

$$h = \frac{\mu_\uparrow - \mu_\downarrow}{2}\,, \tag{6.9}$$

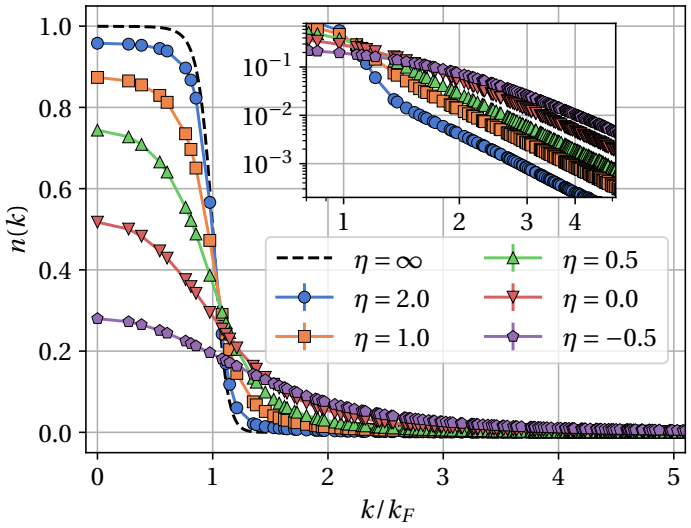

Figure 6.2: Momentum space density distribution of the 2D Fermi gas in the BEC-BCS crossover. The density is shown for different values of the coupling strength $\ln k_F a$, starting with the BCS regime at $\ln k_F a = 2.0$ indicated by the blue circles and ending in the BEC regime at $\ln k_F a = -0.5$ indicated by the purple pentagons. The dashed line corresponds to the exact solution for the non-interacting gas. The inset shows the density on double logarithmic scale, highlighting the polynomial decay for large momenta.

as an additional tuning parameter for simulations, while the central chemical potential is given by

$$\mu = \frac{\mu_\uparrow + \mu_\downarrow}{2}. \tag{6.10}$$

These parameters control the total magnetization M and particle number N, respectively, and we tune them in unison to land on the target values. The approach is much the same as before, but now involves three derivative relations, specifically

$$\frac{\partial\langle N\rangle}{\partial\beta\mu}, \quad \frac{\partial\langle N\rangle}{\partial\beta h}, \quad \frac{\partial\langle M\rangle}{\partial\beta h}. \tag{6.11}$$

The chemical potential derivative of the magnetization is equal to the second term and not computed separately.

Approaches to offload the effort of tuning the lattice parameters directly to a single simulation exist. For instance, the value of the chemical potential can be adjusted dynamically during the simulation or the thermalization period of the Markov chain [174]. However, we do not make use of these more sophisticated methods, as the overhead of conducting a single 'pre-run' for each set of parameters is quite manageable for our purposes.

6.2 The balanced gas in 2D

To verify the validity of our computations, we start by examining the balanced gas in two spatial dimensions. This system has been studied extensively in the literature, initially as the attractive Hub-

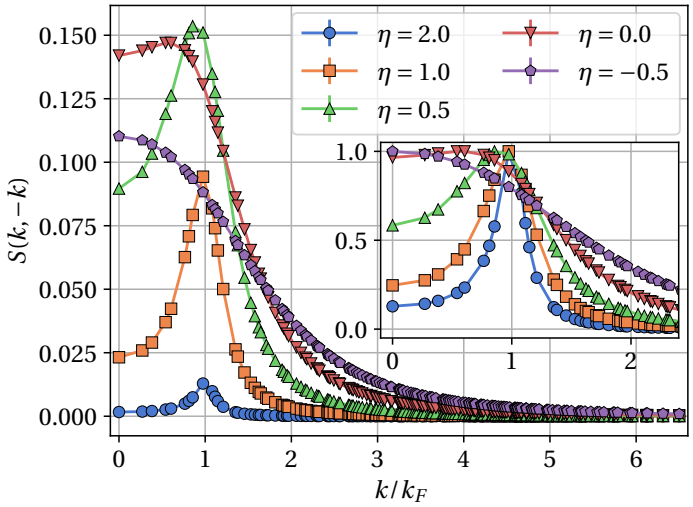

Figure 6.3: Shot-noise / density-density correlations of the 2D Fermi gas in the BEC-BCS crossover. The correlations are shown for different values of the coupling strength $\ln k_F a$, starting with the BCS regime at $\ln k_F a = 2.0$ indicated by the blue circles and ending in the BEC regime at $\ln k_F a = -0.5$ indicated by the purple pentagons. In the BCS regime, a clear peak at the Fermi surface is visible, which gradually vanishes as the system enters the BEC regime. The inset shows the correlations normalized by their maximum value, highlighting the general shape.

bard model for small fillings and later explicitly as a Fermi gas[41, 42, 114, 134, 136, 169, 170]. Fully featured extrapolations to the continuum limit were hard to obtain for a long time but have become possible due to recent breakthroughs with the truncation methods discussed in Section 4.3.

As a preliminary test, we compared our results with those found in the literature for the density equation of state, pressure, and compressibility [134]. Additionally, we calculated the double occupancy and condensate fraction for $N = 58$ particles and compared them to the values given in [42]. All results agree with the literature values within statistical precision.

The qualitative behavior of the 2D gas when varying the coupling strength parameter $\ln k_F a$ is well known. The gas is in the BCS regime for small coupling strengths and exhibits clear pairing peaks around the Fermi surface. As the coupling strength increases, the system smoothly transitions into a phase of tightly bound composite particles with bosonic quantum statistics.

6.2.1 Densities and correlations in the BEC-BCS crossover

To illustrate the effect of the crossover on the density distribution and correlations in the system, we perform a sweep of the coupling strength $\eta = \ln(k_F a)$ at a fixed temperature of $T/T_F = 0.125$ while keeping the total (average) particle number fixed at $\langle N \rangle = 86$. This is achieved by adjusting the chemical potential as described in Section 6.1.2. We use a lattice volume of $V = 35^2$ and a temporal extent of $N_t = 560$, which yields a total of 6.86×10^6 auxiliary field points. The four Fermi interaction is decoupled by a spin-channel decomposition, leading to faster convergence for this system, and we apply truncations with a threshold of $\epsilon_T = 10^{-5}$.

The momentum space density distribution is presented in Figure 6.2. It is important to note that this is not the Fourier transform of real space density but is represented by $\psi_p^\dagger \psi_p$. The non-interacting gas shows quick saturation of momentum states inside the Fermi surface and a steep slope in the distribution around k_F. As the coupling strength increases from $\eta = 2.0$ to $\eta = -0.5$, the slope becomes smoother, and the distribution extends to larger momenta while lower-lying states are depleted. The inset presents the same data in a double logarithmic scale, highlighting the polynomial decay of the density for larger momenta. This decay can be used to determine the contact parameter, as demonstrated in [42, 136].

In Figure 6.3, we show the connected density-density correlations, also called shot-noise, as defined in Equation (5.16), using the same parameters as before. The inset displays the correlations normalized by their maximum value. There is a clear peak around the Fermi surface for weak interactions, indicating BCS-type pairing. With increasing interaction strength, this peak broadens and eventually vanishes so that correlations are instead maximal at zero momentum. The plot also shows an increase in the magnitude of correlations in the system. The chosen temperature $T/T_F = 0.125$ is well above the BKT-transition temperature at $\eta = 2$ but lies in the superfluid phase at $\eta = -0.5$ when comparing with the results in [42]. The breakup of the Fermi surface as well as the phase transition may be reasons for the increased correlations in the BEC regime.

6.3 Imbalanced systems in 2D

We now explore the thermodynamic properties of the imbalanced Fermi gas in two dimensions using Monte Carlo sampling, while addressing the sign problem and its effects on the system. Although numerous mean-field studies have examined imbalanced gases in two dimensions [175–177], the low-temperature physics are far from settled. In particular, it is still unclear whether superfluidity is supported in the presence of finite polarization and, if so, what type of superfluid phase might emerge.

Concerning lattice-based Monte Carlo approaches, we are aware of only a single study [35], which considers the deep BCS regime exclusively. In contrast to the 1D and 3D cases, an attempt has yet to be made to apply the complex Langevin method to two-dimensional systems.

Experimentally, two major studies have investigated imbalanced Fermi gases in two dimensions. In the first work by Ong et al. [30], the authors used a harmonic trapping potential and found that excess fermions are pushed outwards so that a balanced center is formed. This behavior mirrors what was observed in three-dimensional systems, where balanced and imbalanced phases are separated by a first-order transition [28, 29]. In [31], the authors found the same behavior but were also able to measure a condensate fraction in the trap center, even with a slight imbalance. Whether there is a first-order transition between the balanced and imbalanced regions is still an open question that likely requires lower temperatures to resolve. More recently, an investigation using a rectangular confining potential found no signs of such a transition [178]. Additionally, the existence and nature of an FFLO-type or other unconventional phase remains unknown at the time of writing.

In the following, we focus on the normal phase of the system and the grand canonical parameters. This can be seen as a direct extension of previous works [134, 165–167], where the thermodynamics of the balanced gas were examined both theoretically and experimentally. We compute in the strongly interacting parameter regions of the BEC-BCS crossover regime, and discuss the

implications of population imbalances. We also discuss the sign problem in the system and its effects, particularly considering previous complex Langevin results for similar systems in three dimensions.

6.3.1 Model and technical setup

We use the same lattice Hamiltonian as in the balanced case (6.1), but allow for non-degenerate chemical potentials. Truncations are performed on the majority U-matrix, from which the minority matrix can be computed, see Section 4.3.

For the imbalanced case, we use the discrete density channel Hubbard-Stratonovic transformation given in Equation (2.34). The choice is motivated by the sign problem, which is typically less severe in density channel compared to spin channel.

6.3.2 Virial expansion

To facilitate a comparison with the lattice results, we compute select observables by virial expansion. For this, we use the coefficients computed in [179] via a time discretization and the method for Padé approximations therein. The virial expansion is an expansion in the fugacity $z = e^{\beta\mu}$, for which the grand potential becomes

$$\Omega - \Omega_0 = -\beta Z_1 \sum_{n,m=0}^{\infty} z_\uparrow^n z_\downarrow^m \Delta b_{n,m}\,, \tag{6.12}$$

where $\Delta b_{n,m} = b_{n,m} - b_{n,m}^0$ is the difference of virial coefficients between the interacting and non-interacting systems, Ω_0 denotes the non-interacting grand potential and Z_1 is the partition function for a single particle. Next, we compute the relevant observables via derivatives of the grand potential. For the density, we find

$$n - n_0 = \frac{Z_1}{V} \sum_{n,m} \Delta b_{n,m}(n+m) z^n z^m\,, \tag{6.13}$$

where $n_0 = \sum_\sigma \log[1 + z_\sigma]$ is the non-interacting density. In order to find the polarization, we compute the magnetization, which is given by

$$m - m_0 = \frac{Z_1}{V} \sum_{n,m} \Delta b_{n,m}(n-m) z^n z^m\,, \tag{6.14}$$

with the non-interacting magnetization $m_0 = \log\left[1 + z_\uparrow\right] - \log\left[1 + z_\downarrow\right]$. With these results, we compute [3/2] Padé approximants as a function of the parameters $\beta\mu$, while keeping βh fixed. A more in-depth introduction to the virial expansion and Padé approximation for the two-component Fermi gas is provided in Appendix C.

6.3.3 Thermodynamic properties

This section presents results for thermodynamic quantities at various external magnetic field values. To achieve this, we use a lattice of size $V = 27^2$, sufficient to provide a good description of the thermodynamic limit. The temporal spacing is set to $\Delta t = 0.05$. Additionally, we choose the temporal extent β such that the thermal wavelength satisfies $1 \ll \lambda_T = \sqrt{2\pi\beta} \ll L$. For the computations

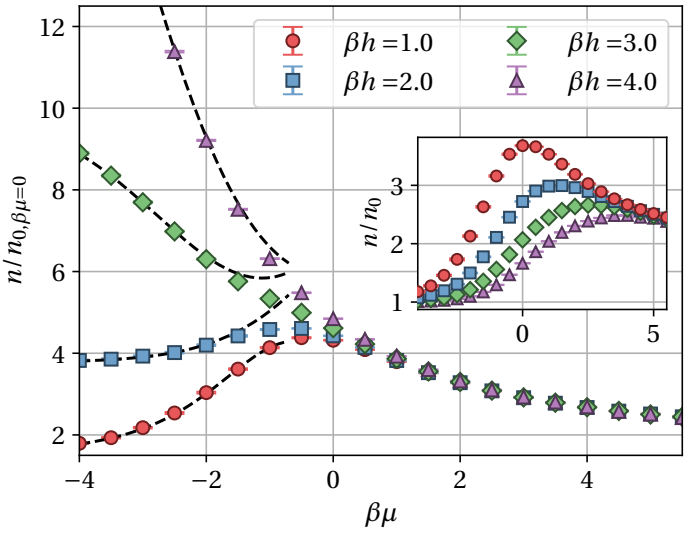

Figure 6.4: Density equation of state for $\beta\epsilon_b = 2$ and various values of the external field βh. The density is normalized by that of the non-interacting system at $\beta h = 0$. The dashed black line is a Padé approximation based on a fifth order virial expansion. The inset shows the same data, but normalized by the non-interacting density at the respective value of βh.

in this section, β is set to 16, corresponding to $\lambda_T \approx 10$. It is important to note that β controls the filling when the physical parameters are fixed and needs to be increased to reach the continuum limit. For the largest coupling considered, systematic errors in the density are negligible in the dilute regime. However, we observe deviations of 1% in the crossover regime at $\beta\mu = 0$ and 2.5% at $\beta\mu = 4$. A detailed discussion regarding the errors introduced by the finite lattice and filling can be found in Section 6.3.4.

In Figure 6.4, we show the density equation of state for $\beta\epsilon_b = 2$, and various external magnetic field values βh. We normalize the density by that of the non-interacting system at $\beta h = 0$ to keep the convention established in previous work [32]. This allows for a detailed comparison with results obtained via virial expansion, including a subsequent Padé approximation, to which we find excellent agreement up to $\beta\mu = -1$. At larger values of $\beta\mu$, the virial expansion breaks down. Interestingly, the densities for different external fields fall on a single curve quite rapidly as the chemical potential increases, even for the largest imbalance considered.

The inset shows the same quantities normalized with the non-interacting density at the same external field. There, we find a delayed onset of the enhancing behavior with increasing imbalance, owing to the smaller number of possible pairs, which is constrained by the number of minority species fermions.

The impact of the sign problem on the density is minimal, even in scenarios where the average sign is small, provided it remains distinguishable from zero. This minimal impact is due to a high correlation between the observables $\langle n \times \text{sign} \rangle$ and $\langle \text{sign} \rangle$. The correlation can be explained by

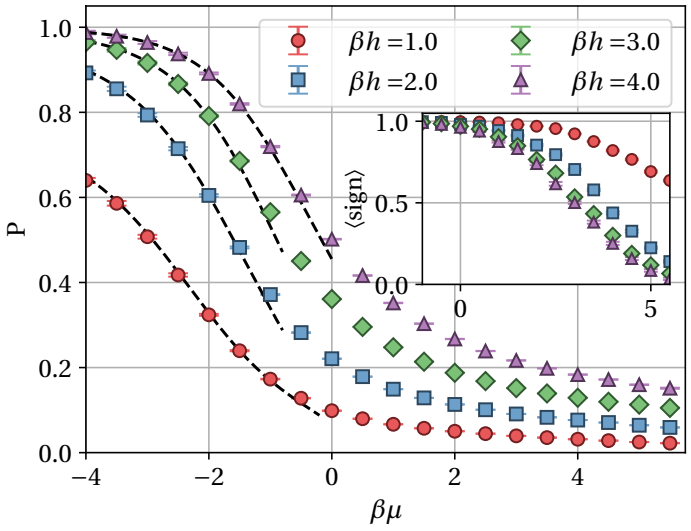

Figure 6.5: Polarization equation of state for $\beta\epsilon_b = 2$ and various values of the external field βh. The dashed black line is a Padé approximation based on a fifth order virial expansion. The inset shows the average sign of the lattice simulations for the same parameters.

inspecting the total density of a field configuration determined by

$$\begin{aligned} n &= \frac{1}{V}\,\mathrm{Tr}\left[\frac{U_\uparrow}{1+U_\uparrow}+\frac{U_\downarrow}{1+U_\downarrow}\right], \\ &= \frac{1}{V}\sum_i \frac{d_{i,\uparrow}}{1+d_{i,\uparrow}}+\frac{d_{i,\uparrow}}{e^{2\beta h}+d_{i,\uparrow}}, \end{aligned} \tag{6.15}$$

where $d_{i,\sigma}$ are the eigenvalues of the matrix U_σ. The relative shift in chemical potential does not affect the contribution of particles deep in the fermi-sea, owing to their large (absolute) eigenvalues, leading them to contribute as unity irrespective of the sign. In general, we find the density per field configuration to be positive, which means that $\langle n \times \mathrm{sign}\rangle$ only becomes negative if the sign itself is negative.

In Figure 6.5, we show the polarization, given by

$$P = \frac{n_\uparrow - n_\downarrow}{n_\uparrow + n_\downarrow}. \tag{6.16}$$

It generally increases with the external field and is suppressed when going to larger chemical potentials. Once again, we find good agreement with the virial expansion result, which we compute via the ratio of the Padé approximants for the density (6.13) and magnetization (6.14). The average sign for the corresponding parameters is shown in the inset. It decreases rapidly when increasing $\beta\mu$. Interestingly, when considering the average sign as a function of the external field, changes are most rapid for small βh, particularly between $\beta h = 0$ and $\beta h = 1$. We believe the smaller number of available pairs also causes this as βh becomes large, decreasing the effect of interactions on the overall system.

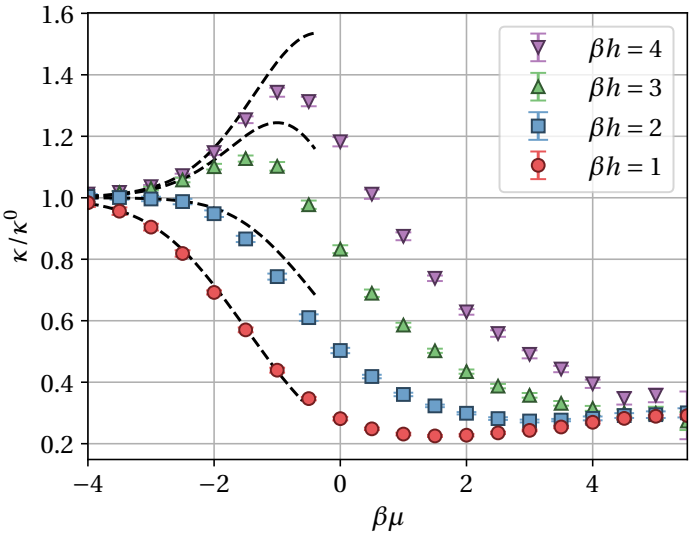

Figure 6.6: Compressibility in units of the non-interacting compressibility for $\beta\epsilon_b = 2$ for $\beta h = 1,2,3,4$ from bottom to top (circles, squares, triangles, diamonds). The dashed black line is a Padé approximation based on a fifth-order virial expansion.

The pressure of the imbalanced gas can be directly derived from the density equation of state through integration, yielding

$$P(\beta\mu) = \frac{1}{\lambda_T^2}\int n\, d(\beta\mu)'\,. \tag{6.17}$$

We also use this equation to calculate the non-interacting pressure for normalization purposes. We present results for the pressure in Figure 6.7. Given the smooth nature of the density data, we apply a simple spline extrapolation between data points, which we then integrate. To account for the contribution from the tail of $n(x)$ for $\beta\mu \to -\infty$, we integrate the virial expansion result up to the first lattice data point, indicated by the dashed lines. Since the statistical errors in the density are minimal, they are not explicitly shown for the pressure. In analogy to the density equation of state, the pressure enhancement grows weaker as the external field increases. This effect is especially evident in the crossover regime, where the enhancement is strongest for the balanced gas [134, 165, 166].

Finally, we compute the isothermal compressibility of the imbalanced gas, which is obtained by taking the derivative of the density with respect to the chemical potential, yielding

$$\kappa = \frac{1}{n^2}\frac{dn}{d(\beta\mu)}\,. \tag{6.18}$$

Using the density computed in Figure 6.4, we perform a spline extrapolation to the data and obtain the compressibility. This estimation is crosschecked via the direct evaluation of the field configurations using

$$\frac{dn}{d(\beta\mu)} = \frac{1}{V^2}\langle N^2 - \langle N\rangle^2\rangle\,, \tag{6.19}$$

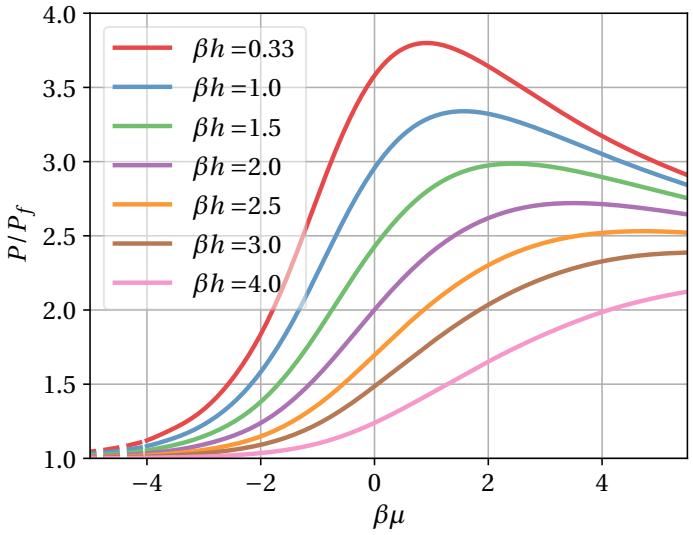

Figure 6.7: Pressure in units of the non-interacting pressure for $\beta\epsilon_b = 2$. The dashed lines on the left indicate the tails integrated from the virial expansion in the high temperature regime.

where $N = N_\uparrow + N_\downarrow$. Generally, both methods yield excellent agreement, although some deviations occur as $\beta\mu$ increases due to the finite lattice size. When these deviations exceed the statistical errors of the direct estimation, we adjust the error bars to span twice the range between them [134]. In Figure 6.6, we show the compressibility for $\beta\epsilon_b = 2$ and various external field values. As βh increases, the compressibility shows an enhancement compared to its non-interacting counterpart, contrasting the situation for the balanced gas and at small βh, where it is suppressed. This enhancement can be attributed to the stronger suppression of the density (see Figure 6.4 inset) compared to the particle number fluctuations (6.19). Crucially, we again agree with the virial expansion results up to $\beta\mu \approx -2$, including the upwards bending behavior for large external fields. The Padé approximation also correctly predicts a peak and subsequent drop in the compressibility for $\beta h = 3$ and $\beta h = 4$. However, both height and position do not agree with the lattice results.

6.3.4 Systematic error estimates

We conduct a series of checks to estimate the systematic errors of the results presented in the previous section. Specifically, we examine the system's density at three different values of the dimensionless chemical potential $\beta\mu = -4, 0, 4$ for an asymmetry of $\beta h = 3$ and $\beta\epsilon_b = 2$. This approach allows us to analyze the system as it transitions from a high-temperature dilute regime to a low-temperature quantum degenerate state.

There are two limits to consider: the thermodynamic limit and the continuum limit. The thermodynamic limit is approached by keeping the lattice parameters fixed while increasing the number of points. To take the continuum limit, one decreases β, which leads to a decrease in filling and consequently decreases the bare coupling g. This reduction in g also makes the error induced by the Trotter-Suzuki decomposition smaller. For our computations, which use $\Delta t = 0.05$, we find this error to be small compared to those introduced by spatial lattice effects and typically of the order

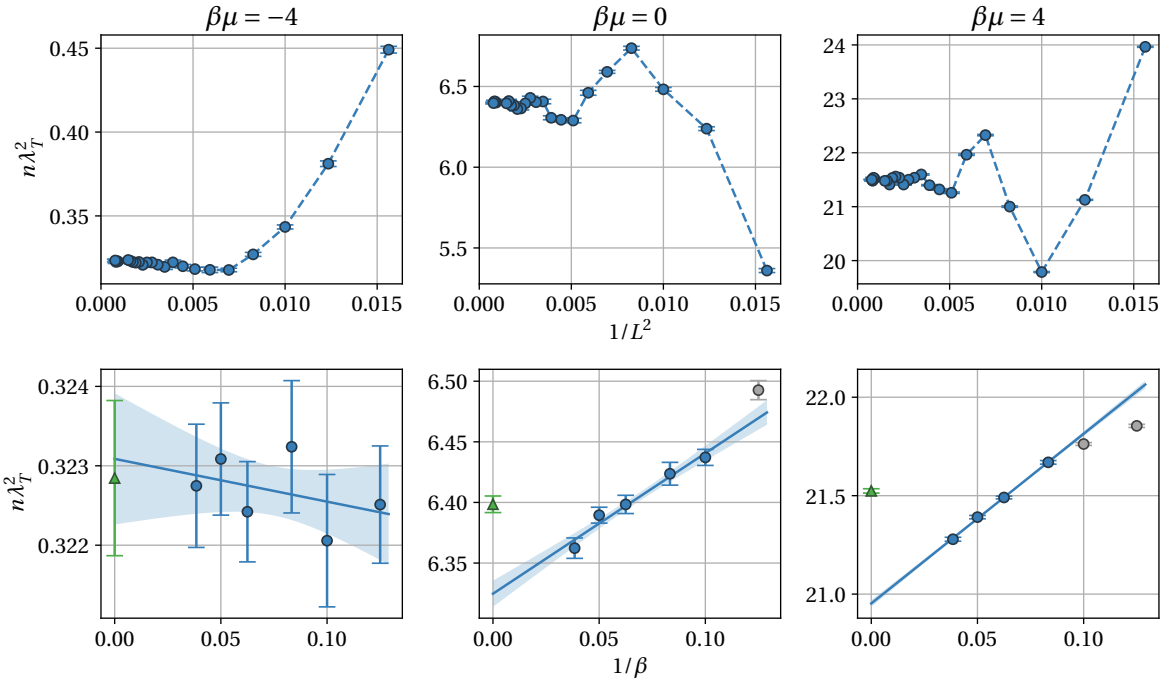

Figure 6.8: Extrapolation of the density for $\beta e_b = 2$, $\beta h = 3$ in the thermodynamic limit for $\beta = 16$ (top) and the continuum limit (bottom). The solid lines are linear fits, where only the colored data points are included. The shaded regions indicate error estimates of the fit. The green triangles are the lattice results shown in Figure 6.4.

of the statistical uncertainty.

The results of an extrapolation to large lattices and dilute systems are shown in Figure 6.8. Achieving the thermodynamic limit is particularly challenging for fermions due to strong shell effects caused by the sharp momentum space distribution. This can be seen in the top panel of Figure 6.8, where we plot the density for various lattice sizes with $\beta = 16$. We find that the density oscillates as the number of lattice points increases, with the amplitude decreasing as the lattice size grows. These effects are most prominent at smaller fillings, where the density of available momentum states is smaller. Since linear extrapolation is not feasible in this scenario, we use the largest lattice computed to estimate the continuum limit. We use $V = 41^2$ for the most dilute systems, while smaller lattices suffice to reach the thermodynamic limit at larger β.

In the bottom panels of Figure 6.8, the density is shown for various values of β on the largest lattices used. Its behavior is well described by a linear extrapolation in $1/\beta$, which we use as an estimate for the infinite volume result. Note that this extrapolation is effectively equivalent to an extrapolation in the total filling of the lattice when sufficiently close to the thermodynamic limit. In the semiclassical regime ($\beta\mu = -4$), the density remains constant within the error margin, as expected, since the filling is small even for the lowest β considered. Thus, the density computed on the finite lattice is expected to agree with the continuum result, as indicated by the excellent agreement with the virial expansion prediction. We observe a stronger dependence in the quan-

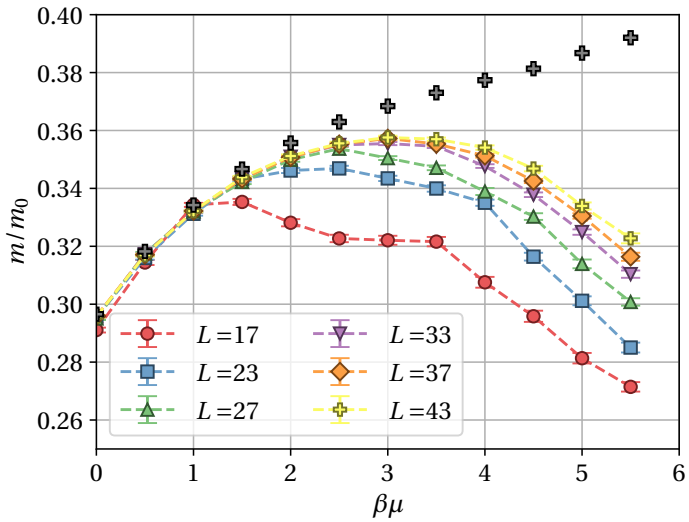

Figure 6.9: Magnetization, normalized by the asymptotic magnetization of the non-interacting system $m_{0,\infty}$ for $\beta\epsilon_b = 2$, $\beta h = 1$ and $\beta = 16$. We compare lattice volumes between $V = 17^2$ and $V = 43^2$. The gray crosses indicate quenched results for the largest lattice.

tum degenerate regime, where the linear extrapolation indicates a decrease in the dimensionless density as the lattice system becomes more dilute. Comparing with the results given in Figure 6.4, we find the extrapolated value to carry a relative deviation of 1% at $\beta\mu = 0$ and 2.5% at $\beta\mu = 4$.

For two reasons taking a continuum limit of the grand-canonical description becomes increasingly difficult for large $\beta\mu$. First, the Fermi surface becomes sharper if the lattice density remains fixed, leading to a stronger dependence on the total volume, even for a spin-balanced system. Secondly, the sign problem becomes severe at low temperatures, rendering the reweighting procedure increasingly ineffective.

6.3.5 Magnetization

In contrast to the density, the total magnetization of the system,

$$M = N_\uparrow - N_\downarrow \,, \tag{6.20}$$

can display a relatively strong dependence on the volume, mainly when the imbalance is small, but the temperature is low. This sensitivity partly arises from the non-monotonic behavior of the density discussed in Section 6.3.4 and is further enhanced by taking the difference between the two spin species. While the issue is somewhat obscured in the total density due to the large number of fermions below the Fermi surface, these fermions do not significantly affect the magnetization when the imbalance is slight.

Figure 6.9 shows the magnetization for increasing central chemical potential and $\beta h = 1$ across lattices ranging from $L = 17$ to 43. The values are normalized by the asymptotic magnetization of the free system $m_{0,\infty}$, where $\beta\mu \to \infty$ while βh remains fixed. We observe that the volume dependence of the magnetization is more pronounced than that of the density. Specifically, the 17^2

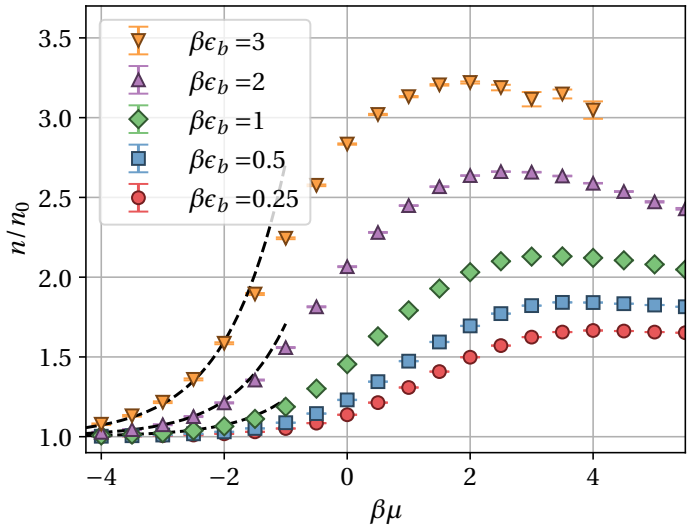

Figure 6.10: Density normalized by the non-interacting density for $\beta h = 3$ and $\beta\epsilon_b = 0.25, 0.5, 1, 2, 3$. The dashed lines indicate results from a Padé approximation based on the fifth order virial expansion.

lattice exhibits strong fluctuations, which smooth out as the volume increases. The magnetization seems to decrease after about $\beta\mu = 3$, possibly due to pairing becoming energetically favored as the temperature drops, eventually leading to a balanced conventional superfluid. However, due to the large volume dependence, we cannot draw a definitive conclusion. Further tests suggest an enhancement of the down-bending behavior as β increases while keeping the dimensionless parameters fixed. It is important to note that the back-bending occurs only for low imbalances and is not present for the computations at $\beta\epsilon_b = 2, 3, 4$. For the 43^2 lattice, we also observe that the phase quenched estimator of the magnetization starts to deviate significantly from the reweighted one at about $\beta\mu = 2$, which does not show a decrease but rather a monotonic increase.

Overall, we expect the stronger dependence on volume to make the magnetization less reliable than the total density when studied in the grand canonical context. In particular, the susceptibility becomes difficult to estimate, as it is directly related to the magnetization of the derivative with the external field,

$$\chi = \frac{1}{V}\frac{dM}{dh}. \tag{6.21}$$

One possible solution is to work in the canonical ensemble or to treat the grand canonical parameters as tuning parameters, aiming for a fixed density and imbalance.

6.3.6 Coupling scan

In Figure 6.10, we show the density for a set of different dimensionless couplings ranging from $\beta\epsilon_b = 0.25$ to $\beta\epsilon_b = 3$ for $\beta h = 3$. We use the same lattice parameters as in Figure 6.4. In agreement with the results for the balanced gas [134, 165–167], the density is a smooth curve and becomes enhanced with increasing interaction strength. However, the sign problem becomes severe for

the largest coupling considered, indicated by a significant rise in statistical error when $\beta\mu$ exceeds 2.5. The Padé approximants show good agreement in the regime of low chemical potential but begin to deviate beyond $\beta\mu = 2$. These approximants are provided only for couplings greater than $\beta\epsilon_b = 1$, as they become increasingly unstable for smaller couplings, offering little improvement over the regular fifth-order virial expansion. Overall, this shows that while we can obtain reliable results in the crossover regime, the sign problem becomes prohibitive on the BEC side and at low temperatures.

6.4 Conclusion

This chapter explored the thermodynamic properties of a spin-imbalanced gas of fermions with attractive contact interactions in two spatial dimensions. To study these properties, we performed lattice Monte Carlo simulations of the grand canonical partition function, which is known to be affected by the sign problem. However, we demonstrated that the sign problem is mild for a wide range of parameters. Moreover, we argued that some observables, such as the density, are only mildly impacted, even when the average sign is small. This suggests that phase-quenched estimators can still yield accurate results. Due to the mildness of the sign problem, we were able to perform computations on larger lattices, facilitated by a model space truncation approach, and employed standard reweighting procedures to obtain unbiased estimators for physical observables.

We calculated the density and polarization equations of state, compressibility, and pressure. All these observables show excellent agreement with results derived from a Padé approximant based on fifth-order virial coefficients up to roughly $\beta\mu \approx -2$. In particular, we find an enhancement in compressibility compared to the non-interacting system as the external field increases. All results presented here are quantities directly accessible to cold atoms experiments and make for falsifiable predictions. In fact, experiments measuring the density, pressure, and compressibility of spin-balanced systems have already been conducted [166, 167].

Overall, the weak sign problem opens up opportunities for further investigations of the spin-imbalanced Fermi gas. While reaching temperatures low enough to search for an inhomogeneous superfluid phase may be challenging, the smallest temperatures we considered are well within the regime where the balance gas displays pseudogap effects. It could be interesting to explore whether similar effects occur in the imbalanced gas and, if so, how the imbalance influences them. Although we expect the sign problem to be more pronounced in higher dimensions, the reweighting approach can be easily generalized. This may present an opportunity to study the unitary gas, which has previously been done via a complex Langevin approach [32, 33], and could provide a helpful crosscheck and insights into the validity of the latter.

Chapter 7

Inverse Renormalization group architectures for normalizing flows

In the previous chapters, we have shown how the sign problem can make lattice simulations intractable and discussed approaches to alleviate this issue. A similarly challenging problem in lattice field theory is critical slowing down, which occurs when correlation lengths become large. Since this happens in the continuum limit regardless of the system specifics, the problem is of general interest. It affects any lattice simulation even away from the critical point of the physical theory. Fortunately, the critical exponent associated with critical slowing down is not universal but depends on the sampling algorithm. This is made evident, for example, by the staggering success of cluster algorithm for spin models, which can reduce autocorrelation times by several orders of magnitude and show better scaling behavior than Metropolis Monte Carlo [180]. With the ever-rising popularity of machine learning, it is natural to ask if these methods can similarly aid critical slowing down in lattice field theory.

Normalizing flows are a class of generative models that aim to find a map between a trivial (most of the time Gaussian) distribution and some complicated target distribution. While normalizing flows, in the context of machine learning, have become popular roughly within the last decade [181–183], related ideas have been proposed in the context of lattice field theory before. In particular, the trivializing maps introduced by Lüscher [184] can be seen as a direct precursor.

In this chapter, we will give a general introduction to normalizing flows and discuss the popular Real NVP and continuous flow architectures. Moreover, we show how normalizing flows may be used to upscale lattice configurations in the context of inverse renormalization group transformations. This can substantially decrease the computational footprint of the model compared to the architectures proposed in the literature. Additionally, since the large-scale features are already sampled on the coarse lattice, critical slowing down could be avoided with such an approach.

7.1 Normalizing flows for lattice field theory

A normalizing flow model aims to map a simple, known prior distribution to a more complicated posterior, or target, one. For our purposes, we are interested in posterior distributions, where the (unnormalized) probability density is known, but sampling may be hard due to, e.g., large autocorrelation times. Notably, the construction of flows does not depend on the choice of prior distribu-

tion, and a good initial guess can significantly reduce training time and model performance.

The general idea is to find some changes of variables, such that the probability density of the target distribution $p(\phi)$, where $\phi \in \mathbb{R}^N$, may be written in terms of the prior distribution $p_p(z)$ as

$$\begin{aligned} p(\phi) &\simeq p_p(f^{-1}(\phi)) \left| \det\left(\frac{\partial f^{-1}(\phi)}{\partial \phi} \right) \right| \\ &= p_p(f^{-1}(\phi)) \left| \det J[f] \right|^{-1} = q(\phi)\,. \end{aligned} \tag{7.1}$$

Here, $z \in \mathbb{R}$ and $f^{-1}(\phi) : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^{\mathbb{N}}$ is an invertible function. It describes the *normalizing* direction of the flow and maps samples from the target distribution to the prior, typically *normal* distribution. Consequently, f maps samples from the prior to the target distribution and is called the *generative* direction.

Since critical slowing down poses a significant challenge to conventional lattice simulations, the allure of a map that can generate samples is clear. Configurations are drawn from the prior, trivial distribution. Each sample is independent by construction. As a result, the samples in the target distribution are also independent, effectively eliminating autocorrelation and resolving critical slowing down.

In practice, the function f is defined by a machine learning architecture and must allow for the computation of the Jacobi determinant in a tractable manner. However, training these functions comes with a significant computational cost, which increases exponentially close to criticality based on early experiments [47]. Thus, while critical slowing down is diminished in sampling, it is transferred to the training process.

It is important to note that the learned map will never precisely match the target distribution, and the relation in (7.1) holds only approximately. We thus denote the distribution actually learned by the flow as $q(\phi)$, which is an approximation to the target distribution $p(\phi)$. A Markov Chain introduced on the drawn samples resolves this problem and corrects possible deviations. The probability of accepting a new configuration ϕ' based on the current configuration ϕ is computed in the usual Metropolis-Hastings fashion,

$$p_{\text{accept}}(\phi'|\phi) = \min\left(1, \frac{p(\phi')\,q(\phi)}{q(\phi')\,p(\phi)}\right). \tag{7.2}$$

If the mapping is exact, the acceptance rate will be one, rendering the Markov Chain unnecessary. The acceptance rate can serve as a measure of model quality and an optimization target during training. However, it tends to be quite noisy in practice. Instead, the *reverse* Kullback-Leibler (KL) divergence, which is the relative entropy between the target and the approximate distribution,

$$D_{KL}(q||p) = \int q(\phi) \log\left(\frac{q(\phi)}{p(\phi)}\right) d\phi\,, \tag{7.3}$$

is typically used. The KL divergence, is minimized precisely when the two distributions are equal, making it an effective measure of flow quality. It is estimated stochastically during training by drawing samples from the approximate distribution q, representing a variational, unsupervised training approach. This is a distinct advantage over the *forward* KL divergence, which requires samples from the target distribution instead. Importantly, the overall normalization of the target distribution, $Z = \int \exp(-S)$, does not affect the minimization procedure. As long as it remains constant during training, the KL divergence value only shifts by a fixed amount.

7.1.1 Real NVP flows

The first and still one of the most popular normalizing flow architectures for lattice field theory are real non-volume-preserving (real NVP) flows [43, 183]. They construct the complicated function f by a set of affine coupling layers, which scale and offset a subset of the field component in each step. Assuming we have N field component in ϕ, the lattice is split into two halves of size $N/2$, we call ϕ_a the active and ϕ_b the frozen subset. A single affine coupling layer g_i is defined as the transformation

$$g_i(\phi_a) = \phi_a \odot \exp(s_i(\phi_b)) + t_i(\phi_b)\,, \tag{7.4}$$

which leaves the frozen subset unchanged. Here, the scaling and offset functions s_i and t_i both map $\mathbb{R}^{N/2}$ to itself and are conditioned on the frozen half of the fields. The operator $\odot$ denotes element-wise multiplication. In order to achieve a high level of expressiveness, s_i and t_i are usually defined by unique neural networks for each coupling layer. Importantly, these neural networks do not necessarily need to be invertible themselves. The inverse transformation can be calculated with the same computational cost as the forward pass and is expressed as

$$g_i^{-1}(\phi_a) = (\phi_a - t_i(\phi_b)) \odot \exp(-s_i(\phi_b)). \tag{7.5}$$

This simple form makes the Jacobian determinant of the transformation easy to compute via

$$\left|\det\left(\frac{\partial g_i(\phi_a)}{\partial \phi_a}\right)\right| = \exp\left(\sum_j [s_i(\phi_b)]_j\right), \tag{7.6}$$

where the sum runs over all $N/2$ components of the scaling function.

Layers are stacked sequentially, alternating between frozen and active field components. The choice of how to separate the fields is arbitrary, and different decompositions may be used for different layers. A popular approach is a simple checkerboard breakup, allowing all fields to be conditioned on their nearest neighbors.

The scaling and offset functions are typically parameterized by deep convolutional neural networks. This has the advantage of encoding the translational symmetry of the lattice into the flow. Including symmetries is crucial to the success of flow-based methods and can significantly reduce model complexity and training time. Affine coupling layers that are equivariant under $SU(N)$ gauge symmetry and the $Z2$ symmetry of scalar ϕ^4-theory are also available [47].

7.1.2 Continuous flows

The real NVP flows in the last section are a discrete set of transformations that map from the prior to the target distribution in a finite number of steps. On the other hand, continuous flows are defined by a differential equation, describing a continuous path between the two distributions. The concept of continuous flows first emerged in 2018 [185], along with the introduction of neural ordinary differential equations (NODEs). Using the adjoint method, NODEs enable efficient backpropagation through the solution of ordinary differential equations (ODEs). This was a crucial development, as regular backpropagation through an ODE solver is computationally expensive and can be numerically unstable.

In the most generic form, the field is determined by an ODE, which is defined for the flow time $t \in [0, T]$ as

$$\frac{d\phi(t)}{dt} = g(\phi(t), t). \tag{7.7}$$

In the machine learning sense, this can be seen as a type of shallow architecture, but we are free to parametrize the vector field g as we see fit. The only constraint is that it must be Lipschitz continuous. Notably, the function itself is not required to be invertible. The complete transformation is obtained by integrating the ODE from $t = 0$ to $t = T$

$$\phi(T) = \phi(0) + \int_0^T g(\phi(t), t)\, dt. \tag{7.8}$$

It is immediately clear that the transformation is inverted by simply reversing the direction of the flow time. In order to evaluate the weight of the transformed field, the probability density is also tracked during the integration. This is done by integrating another ODE, given by

$$\frac{\log p(\phi(t))}{dt} = -\left(\nabla_\phi \cdot g\right)(\phi(t), t). \tag{7.9}$$

We require only the gradient of g with respect to the field, which is readily computed analytically or via automatic differentiation, depending on the architecture.

In lattice field theory, continuous flows have been used, e.g., for scalar ϕ^4 models [186], gauge theories [187] and Nambu-Goto strings [188].

In the following, we will consider scalar theories and architectures based on the one proposed in [186]. Therein, the vector field g is parametrized by a basis expansion in the field components and a Fourier expansion in the flow time. In total, the vector field is given by

$$g_x(\phi(t), t) = \sum_{y,d,f} W_{xydf} K(t)_d H(\phi(t)_y)_f\,, \tag{7.10}$$

where d and f run over the number of field and time basis functions D and F respectively, and y indexes the N lattice sites. The function $H : \mathbb{R} \to \mathbb{R}^F$ governs the basis expansion, while $K : [0, T] \to \mathbb{R}^D$ is the Fourier expansion. In principle, both of these functions may contain learnable parameters in addition to the N^2DF learnable weights of the tensor W. The latter can be interpreted as a convolutional kernel spanning the whole lattice. The effective range at which lattice points interact can be controlled by fixing the specific weights connecting them to zero. We will call the range of this kernel $\mathcal{K}$ in the following.

In some cases, training can be improved by factorizing the weight tensor and adding more parameters in the form of bond dimensions

$$W_{xydf} = \sum_{d',f'} \tilde{W}_{xyd'f'} U_{d'd} V_{f'f}\,, \tag{7.11}$$

where, d' and f' index the bond dimensions. The resulting matrices U and V are real-valued and of size $D' \times D$ and $F' \times F$, respectively.

In our basis expansion, we use a combination of a polynomial expansion up to order P and learnable frequency terms of the form $\sin(\omega_f t)$ and $\cos(\omega_f t)$. In practical computations, including higher-order terms in the polynomial expansion can lead to numerical instabilities. Therefore, in most cases, we only include the linear term. If the theory is symmetric under a global $\phi \to -\phi$ transformation, the cosine and even polynomial terms can be omitted.

Furthermore, making the model fully equivariant under all symmetries of the lattice actions is crucial for the success of continuous flows and drastically reduces the number of parameters required. In [186], the translational, rotational, and mirror symmetries encountered in scalar theories are discussed in detail. We use a similar construction in the following sections but adjust to symmetry based on upscaling.

7.2 Inverse RG with normalizing flows

The renormalization group is a powerful tool in statistical physics and quantum field theory that facilitates the study of systems at various length scales [189]. In its most basic form, a renormalization group transformation involves "block spinning", a process in which a lattice is coarse-grained step by step, reducing the number of degrees of freedom. One example of such a coarse graining step is simple averaging over blocks of spins or fields. This averaging procedure, however, is not invertible for a given configuration.

Efforts to develop an (approximate) *inverse* transformation have mainly focused on statistical spin systems [190, 191]. More recently, the idea has sprung over to lattice field theories, where convolutional neural networks were used to map from small to large lattices [192]. The authors take a supervised approach, training the network on fine configurations and the corresponding coarse-grained counterparts. The resulting map may then be used to increase the size of a lattice configuration iteratively. Unfortunately, this procedure does not allow for sampling in the coarse theory, as the weight of the coarse configurations is not know analytically.

Related concepts in lattice field theory and machine learning include multi-grid and multi-level sampling approaches [193, 194], as well as a range of super-resolution techniques commonly used in image processing [195].

7.2.1 Upscaling configurations with continuous flows

Since the weight of coarse-grained configurations is generally unknown, and there is a mismatch in size between fine and coarse-grained lattices, we cannot directly use coarse samples to train the flow. Instead, we propose a procedure where the flow effectively links two lattice models, each potentially having different parameters and sizes. This approach requires us to address the issue of bijectivity between the original and coarse field spaces, as the number of degrees of freedom differs.

To tackle the mismatch in dimensionality, we employ an injective layer, which allows us to compute the density change as

$$p(\psi) = p_p(U^{-1}(\psi)) \left|\det\left(J^T[U]J[U]\right)\right|^{-1/2}_{U^{-1}(\psi)}, \tag{7.12}$$

where $U : \mathbb{R}^n \to \mathbb{R}^N$ is an injective function with $n < N$, and the Jacobain is a $N \times n$ matrix. Note that U^{-1} in only defined on a subset $M_N \subset \mathbb{R}^N$.

Our approach utilizes a fixed transformation that duplicates field components along all dimensions, resulting in a single upscaling step $L \to 2L$. Here, L is the linear extent of the lattice. The

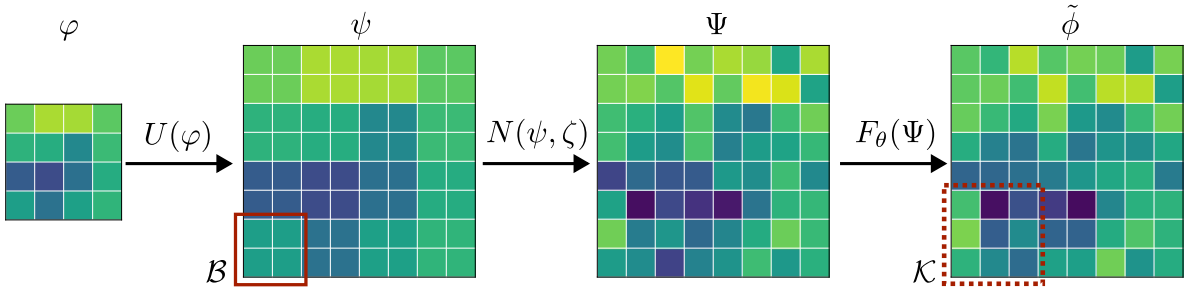

Figure 7.1: Illustration of the full upscaling procedure. The original sample is naively upscaled by a factor of 2, and noise is added to each block. The flow then maps the noisy, upscaled sample to the target distribution.

Jacobian for this upsampling operation also used in [196] is given by

$$J = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ \vdots & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ 0 & \vdots & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \cdots & 1 \\ 0 & 0 & \cdots & 1 \end{pmatrix}_{(2L)^d \times L^d} . \tag{7.13}$$

Since the transformation is independent of the field, the determinant remains constant with a value of $\log|\det J^T J| = L^d \log(2^d)$, where d is the dimensionality of the lattice. This constant shift in the KL divergence can be disregarded during training. However, it may be relevant when calculating specific observables, such as the change in the partition function between prior and posterior distribution.

As stated above, renormalization group transformations are non-invertible, which means the coarse configuration is ignorant regarding the fine-scale information lost in the coarsening step. We parametrize this ignorance by introducing a degree of stochasticity to the sample in the form of correlated noise for each block $\mathcal{B}$ in Figure 7.1. This is easily achieved by drawing a multivariate random variable ζ of size $N_\mathcal{B} - 1$ from a Gaussian distribution, where $N_\mathcal{B}$ is the number of sites in a single block, and the covariance matrix is given by

$$\Sigma = \mathrm{I}_{N_\mathcal{B}-1} - \frac{1}{N_\mathcal{B}}. \tag{7.14}$$

We require the noise added to each block to be homogeneous and to sum to zero. The latter ensures that a simple averaging over blocks leads back to the original coarse sample and means that the operation can be understood as the left inverse of both the upscaling and noise application. With

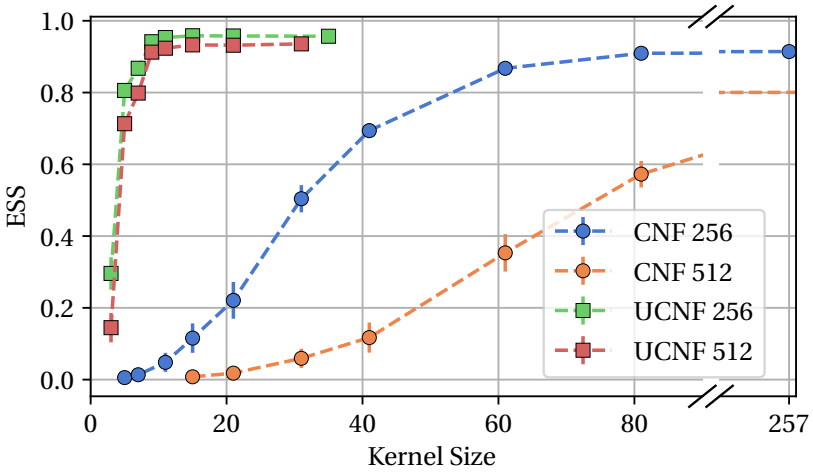

Figure 7.2: Effective sample size (ESS) as a function of the kernel size in one dimension for target lattices of size $L = 256$ and $L = 512$. We compare the upscaling flow (UCNF, squares) to the original continuous flow architecture trained from a Gaussian prior (CNF, circles). The couplings are tuned so that a correlation length of $\xi/L \simeq 0.078$ is achieved in all cases.

these requirements, the final component of the noise for each block is trivially computed as

$$\zeta_{N_\mathcal{B}} = -\sum_{i=1}^{N_\mathcal{B}-1} \zeta_i. \tag{7.15}$$

Samples are drawn independently for each block, which means that we have essentially added $(2L)^d - L^d$ degrees of freedom via noise. The operation that adds the noise is denoted as $N(\psi, \zeta)$, where $\psi \in M_{(2L)^d} \simeq \mathbb{R}^{L^d}$ and $\zeta \in \mathbb{R}^{(2L)^d - L^d}$. The configuration on the coarse lattice and the noise now match the degrees of freedom on the fine lattice, enabling us to learn a bijective map between them. Indeed, with the described procedure, any configuration on the fine lattice can be broken up into a coarsened field and a noise part, and their densities can be tracked exactly.

We complete the model by incorporating a normalizing flow F_θ that maps from the upscaled sample with noise to the target distribution. This flow is continuous and similar to the one discussed earlier, with some adjustments in terms of symmetries and kernel size. The upscaling procedure described so far destroys the sample's translational symmetry, and we rely on the flow to restore it. However, we can still use the symmetry of the original sample, which manifests in a 2-step-translational symmetry, and mirror and rotational symmetries in the upscaled field with noise. Consequently, the upscaling flows require more parameters than the original continuous flow architecture, all else equal. To counteract this increase, the range at which lattice points interact with each other can be reduced using the weight tensor W in (7.10). This is possible since we want the flow to learn short-scale behavior while inheriting the long-range behavior from the coarse sample.

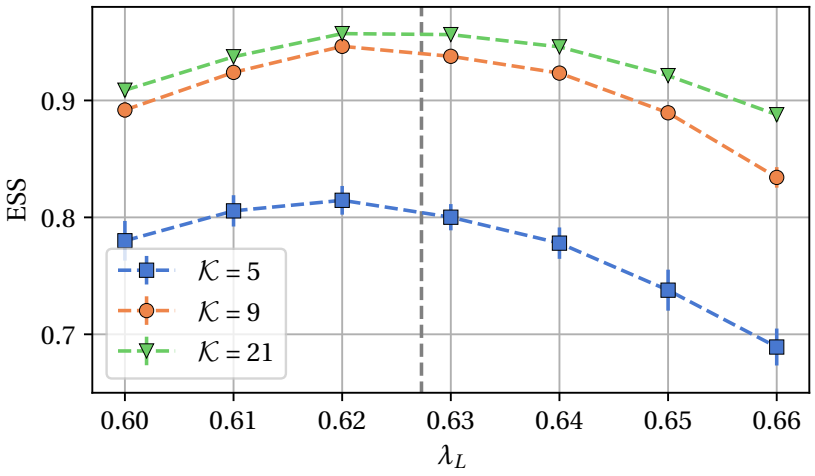

Figure 7.3: Effective sample size (ESS) as a function of the coupling of the coarse system. The dashed gray line indicates the value where the correlation lengths between the coarse and fine systems agree in units of the volume.

7.2.2 Experiments

To test the proposed upscaling architecture, we use it on a simple ϕ^4-theory in one and two dimensions. The lattice action is given by

$$S_L[\hat{\phi}] = \sum_x \left[\frac{\hat{m}_L^2}{2}\hat{\phi}_x^2 + \frac{\hat{\lambda}_L}{4}\hat{\phi}_x^4 + \frac{1}{2}\sum_{\mu>0}\left(2\hat{\phi}_x^2 - \hat{\phi}_x\hat{\phi}_{x+\hat{\mu}} - \hat{\phi}_x\hat{\phi}_{x-\hat{\mu}}\right)\right], \tag{7.16}$$

where the lattice spacing was absorbed into the fields and parameters, yielding the dimensionless quantities

$$\hat{\phi}_x = \phi(a_L x)\, a_L^{d/2-1} \tag{7.17}$$

$$\hat{m}_L = m\, a_L \tag{7.18}$$

$$\hat{\lambda}_L = \lambda\, a_L^{4-d}. \tag{7.19}$$

For now, the bare parameters of the fine S_{2L} and coarse S_L levels are tuned to maintain consistent physics and enable effective upscaling. Unless stated otherwise, we use $D = 17$ components to expand the time kernel and $F = 20$ for the fields. The latter includes a linear term and 19 learnable frequencies. Moreover, the bond dimensions are set to $D' = 20$ and $F' = 20$.

In our first experiment, we consider ϕ^4–theory in $d = 1$ dimensions and use the proposed architecture to learn a continuous normalizing flow between lattices of size L and $L' = 2L$. Since the coarse theory reproduces the infrared (IR) features of the fine level quite well, a significantly smaller kernel $\mathcal{K}$ can be used when learning between theories compared to learning with a Gaussian distribution as prior. Results for lattices of size $L' = 256$, $L' = 512$ and correlation length $\xi/L = 0.078$ are shown in Figure 7.2. To evaluate and compare the performance of the different models, the

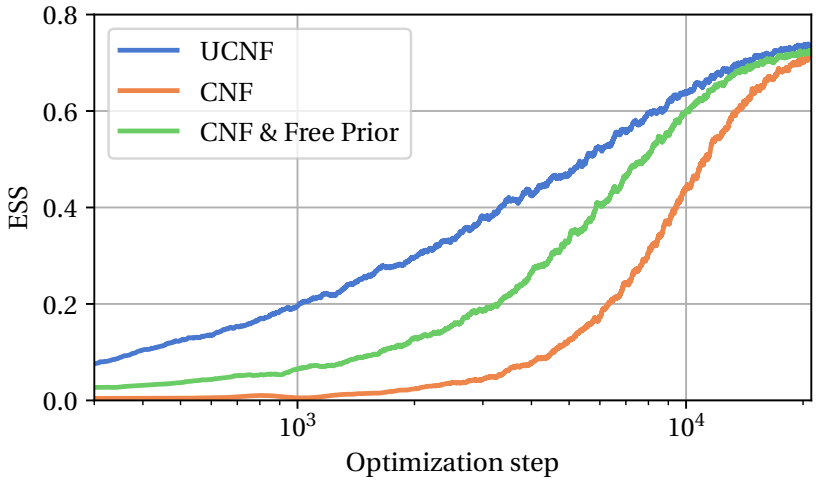

Figure 7.4: Effective sample size on a $V = 64^2$ lattice as a function of the number of training steps. We compare the upscaling architecture to regular continuous flows trained from Gaussian and free theory priors.

effective sample size (ESS) per sample after training is estimated as

$$\mathrm{ESS} = \frac{\left(\frac{1}{N}\sum_i p[\tilde{\phi}_i]/q[\tilde{\phi}_i]\right)^2}{\frac{1}{N}\sum_i \left(p[\tilde{\phi}_i]/q[\tilde{\phi}_i]\right)^2}. \tag{7.20}$$

It is defined so that ESS = 1 for a perfect match between the target and flowed distribution. The model with a coarse prior achieves an ESS of > 0.9 for kernels as small as $\mathcal{K} = 11$ and is well above 0.5 even for $\mathcal{K} = 5$. With a Gaussian prior, performance is considerably worse. This is a direct consequence of the correlation length growing large in lattice units as the volume is increased ($\xi = 20$ and $\xi = 40$ respectively). Performance improves steadily with increasing kernel size and becomes larger than 0.5 for $\mathcal{K} \gtrsim 2\xi$, which is expected since correlating sites separated by more than half the kernel size is challenging with the single layer continuous flow architecture. Overall performance is slightly better for the upsampling model, with a final ESS of > 0.9 compared to 0.8 for the baseline model, even when the kernel spans the entire lattice.

We conducted tests to gauge the impact of varying parameters at the coarse level with fixed kernel sizes of $\mathcal{K} = 5, 9, 21$ in Figure 7.3. The model is sensitive to these changes, and the ESS peaks when the correlation length (measured in terms of volume units) aligns between the coarse and target levels. This indicates the need for a careful tuning and robust matching method to ensure optimal performance. We have performed additional experiments with different levels of external noise, which we found to have a minor impact on overall performance. However, the noise should have reasonable intensity, roughly comparable to the fluctuations between nearest neighbors on the coarse lattice. Also, the upsampling models generally perform better than the baseline ones when the number of frequency basis components is decreased.

In a second experiment, we consider the theory in two spatial dimensions. This case has been studied extensively using various normalizing flow architectures and is an excellent benchmark. In

contrast to the one-dimensional theory, there is a second-order phase transition in two dimensions associated with the breaking of $Z2$ symmetry. This makes the problem more challenging due to significant autocorrelation times at criticality. We use $\hat{m}_L = 1$ and tune the coupling so that the correlation length becomes $\xi / L = 0.25$. Overall, the behavior is very similar to the one-dimensional case. We can drastically reduce the kernel size without a significant loss in performance. Moreover, we find the training to be slightly faster, in terms of optimization steps, than the benchmark when learning from a Gaussian, as evidenced in Figure 7.4. We also show the regular CNF when training starts from a non-interacting theory with unit mass. This trains slightly faster than the Gaussian prior but cannot benefit from the reduced kernel size. Unfortunately, the upscaling model does not perform significantly better than the other approaches after training, with the final ESS remaining broadly consistent across all models.

7.2.3 Summary & Outlook

We have introduced a novel architecture for upscaling lattice configurations in scalar field theories. Our approach can be viewed as an inversion of a block-spinning transformation and involves an upscaling step that encodes the uncertainty of fine-scale information as noise. We then reconstruct this short-range information from the noise using a continuous normalizing flow. Our architecture achieves performance comparable to the best benchmark normalizing flows currently available for both one- and two-dimensional ϕ^4 theories.

However, a notable limitation of our method is the requirement to tune the correlation lengths between the two lattices, which demands prior knowledge of the underlying physical properties. This knowledge is typically obtained through conventional lattice simulations. The tuning is a significant disadvantage, and ongoing work aims to infer the appropriate lattice coupling of the theory on the fine lattice during training.

We are exploring the potential of applying the flow iteratively to upscale configurations progressively to finer grids. In this envisioned scenario, higher momentum modes would be added step by step, thereby avoiding the issue of critical slowing down.

Moreover, the approach should extend to fermionic theories, such as the ones studied in previous chapters. The baseline computational cost is higher there, and the potential gains from normalizing flows and inverse RG techniques are even more substantial.

Chapter 8

Conclusion, Summary & Outlook

In this work, we used stochastic approaches to explore the physics of fermionic ultracold atom systems in reduced dimensions. This was motivated by the rapid experimental progress in the field of ultracold atoms, which has made it possible to gain insights into strongly correlated systems in a controlled environment. Lattice-based methods, like those used here, are well-suited for investigating cold atom physics because they provide non-perturbative access, even when particle numbers are large and interactions are strong. Our work specifically contributes to understanding the role of trapping potentials in lattice approaches and the effects of population imbalances.

In Chapter 2, we briefly introduced lattice field theory and stochastic methods. Potential pitfalls of lattice approaches, including the sign problem and critical slowing down, were also discussed, and solutions were outlined. We reviewed in detail how non-relativistic fermions are put on a lattice and how bosonic auxiliary fields can be introduced, and we particularly focused on the occurrence of a sign problem. The chapter also discussed the projective reweighting approach used for simulations in the canonical ensemble and the ground state later on.

One key simulation method we used was complex Langevin, which we introduced inChapter 3. The approach allows for computations even in the presence of a sign problem but comes with some issues, as discussed. We investigated the applicability of the methods to scalar field theories with an external magnetic field. These models are attractive due to the occurrence of Lee-Yang zeros, but they carry a severe sign problem. In simple toy models, the complex Langevin approach failed, likely due to the stochastic process's inability to correctly handle the phases of different fixed points contributing to the result. Notably, the complete three-dimensional theory showed phase transition like behavior but also some signs of wrong convergence. However, we could not make a definitive statement about its correctness. For this, a benchmark calculation would be necessary, which is made difficult by the sign problem and beyond the scope of this work.

In Chapter 4, we discussed the technical aspects of DQMC computations for non-relativistic fermions. Numerical difficulties arise at low temperatures, which can be addressed by stabilization techniques based on matrix decompositions. This chapter also details the truncation scheme used in subsequent sections, which is crucial for efficient computations, particularly for trapped systems.

The physics of trapped fermions in one dimension were explored in Chapter 5, beginning with complex Langevin computations. To start, we discussed the general applicability of the approach. We argued that it is unsuitable for studying imbalanced Fermions with attractive interactions, at

least in the density-channel Hubbard-Stratonovic transformation. On the other hand, repulsive interactions were found to be treatable, although some slow decay in sampling distributions is always present. To further investigate the balanced system and the effect of trapping potentials, we turned to a more standard DQMC approach. Using this method, we identified the dispersion as a crucial ingredient to the success of simulations. We presented results for density profiles and correlations with up to 80 particles. Moreover, energy observables were compared to results from experiments and exact diagonalization, to which we found good agreement. Studying the imbalanced system, we encountered no sign problem for the parameter ranges considered, which is surprising, as there is no reason for the weight to be purely positive a priori. Clear signs of FFLO-type pairing appeared when computing density-density correlations in the imbalanced system. This result opens the door for further exploration of imbalanced gases through Monte Carlo methods, offering the potential for unbiased predictions.

The absence of a sign problem in one dimension led us to explore the imbalanced (and untrapped) gas in two spatial dimensions in Chapter 6. After briefly reviewing the balanced case, we computed the density equation of state in the presence of an external magnetic field. We found that the sign problem is relatively mild for a wide range of parameters and that the density is only mildly affected, even when the sign problem is severe. We also computed various other observables in the BEC-BCS crossover regime above T_c, including the pressure, magnetization, and compressibility. The latter showed an enhancement in the crossover regime compared to its population-balanced counterpart. These results constitute predictions that experiments can test directly. No signs of a transition to a FFLO phase were evident for the parameters considered.

Finally, in Chapter 7, we introduced a novel normalizing flow architecture for upscaling field configuration, effectively acting as an inverse renormalization group transformation. This was made possible by a naïve upscaling step and the use of noise that encodes the small-scale features. A continuous normalizing then learns a map between this noisy configuration and the theory of interest. We showed that our approach significantly reduces the required model size compared to learning a standard continuous flow with a Gaussian distribution as prior. Moreover, training was faster and more stable, reaching a similar final performance to the benchmark results.

Several direct extensions to this work immediately suggest themselves. For example, a detailed exploration of different auxiliary field decompositions could be performed in the case of complex Langevin simulations for cold atoms. This would help to alleviate the conceptual issues we discussed and potentially lead to a better understanding of the spin-imbalanced system with attractive interactions. The work done on harmonic traps in one dimension should extend straightforwardly to more complicated trapping potentials and higher dimensions. A multitude of interesting phenomena in the two-dimensional trapped case could be studied, including the emergence of superfluid behavior with an increasing number of particles and even rotating systems. The normalizing flow method for inverse RG is still under construction, and the envisaged approach is generic, allowing for applications to an extensive range of systems. One of the most exciting prospects is the possibility of exploring models with long-range interactions. These are typically computationally challenging due to large correlation lengths and the non-local nature of the action.

Acknowledgements

Firstly, I want to thank my supervisor, Jan Pawlowski, for his guidance throughout the last few years, collaboration, and advice on the projects presented in this thesis. I greatly appreciated working with him, our fruitful discussions, and the fantastic research environment he created in the Strongly Correlated Systems group.

Secondly, I'd like to thank Maurits Haverkort for the interesting discussions and for agreeing to be the second examiner of this thesis.

Very special thanks go to Felipe Attanasio for his mentorship and friendship during large parts of my time in Heidelberg. I am grateful for the many he taught me, both in terms of physics and how to write efficient and readable code. I also want to thank him for our many hours of discussions, which were always helpful and enjoyable.

I am particularly indebted to Felipe Attanasio, Renzo Kapust, and Finn Temmen for each volunteering to read parts of this manuscript and for providing valuable feedback.

Working with my colleagues and friends in the (extended) Strongly Correlated Systems group was a great pleasure. I want to thank all of them for the fun discussions, both physics-related and beyond, and all the lunches and BBQs we had together. Here, I want to mention the few to whom I have grown closest: Felipe Attanasio, Friederike Ihssen, Renzo Kapust, and Finn Temmen.

Lastly, I could not have done this without the support, love, and encouragement of my family and fiancée, for which I am beyond grateful.

Appendix A

Diverging variance - a simple example

We consider a simple, one-dimensional model to illustrate the effect of a divergence in the variance on Monte Carlo simulations. One way of realizing such a divergence is by choosing a distribution that decays only polynomially. For our purposes, let

$$p(x) = \frac{1}{\mathcal{N}} \frac{1}{(x+1)^5}, \qquad x \in [0, \infty), \tag{A.1}$$

where $\mathcal{N}$ is a normalization constant. The second moment of the distribution exists and is given by

$$\langle x^2 \rangle = \int_0^\infty dx\, x^2\, p(x) = \frac{1}{3}. \tag{A.2}$$

The variance, however, does not exist due to a logarithmic divergence in the integral. Consequently, naive Monte Carlo estimations of even the second moment are difficult, as no reliable estimate of the error can be given. Indeed, tails in $p(x)$, which are heavily suppressed, contribute significantly to the second moment, making a large number of samples necessary.

As an analog to the method used in the main text, we introduce a modified distribution,

$$p_\alpha(x) = \frac{1}{\mathcal{N}_\alpha} \frac{1}{(x+1)^{5-\alpha}}, \qquad \alpha > 0. \tag{A.3}$$

Although this distribution exhibits an even more severe variance problem, we can use it to obtain better estimates of the moments of our initial distribution. We draw samples from $p_\alpha(x)$, and use a reweighting step to move back to $p(x)$

$$\langle O(x) \rangle = \int_0^\infty dx\, O(x)\, \frac{p(x)}{p_\alpha(x)}\, p_\alpha(x) = \langle O(x) \frac{p(x)}{p_\alpha(x)} \rangle_\alpha. \tag{A.4}$$

This new observable on the modified distribution can now have a finite variance, depending on the parameter α, allowing us to compute a reliable error estimate. Analytically, the variance of the second moment is now given by

$$\langle (x \frac{p(x)}{p_\alpha(x)})^2 \rangle_\alpha - \langle x \frac{p(x)}{p_\alpha(x)} \rangle_\alpha^2 = \frac{384\,\Gamma(\alpha)}{(4-\alpha)\,\Gamma(5+\alpha)}, \tag{A.5}$$

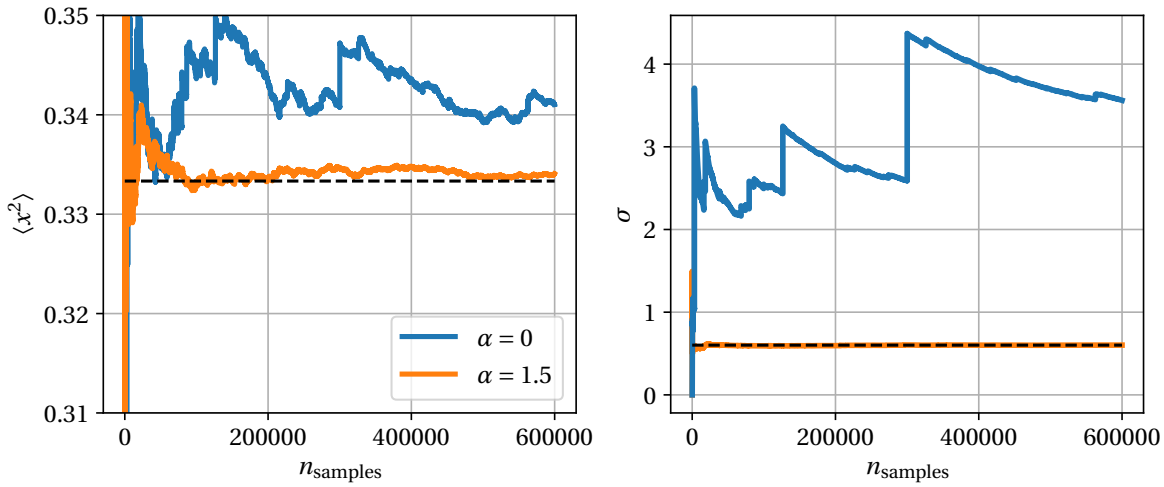

Figure A.1: Left: Comparison of a Monte-Carlo estimate on $\langle x^2 \rangle$ with $\alpha = 0, 1$. The dashed line indicates the analytical result. Right: The variance of the respective estimators converges for $\alpha = 1$ but not for $\alpha = 0$. The dashed line indicates the analytical result for $\alpha = 1$.

which diverges for $\alpha \to 0$ as expected, but is finite for $4 > \alpha > 0$. In fig. A.1, we compare the Monte-Carlo estimate of $\langle x^2 \rangle$ for $\alpha = 0, 1$. The left panel shows the running average of up to 10^6 samples. The naive estimator produces some notable jumps, which are a result of the diverging variance. In contrast, the reweighted estimator is much smoother and converges with the analytical result relatively quickly. The right panel shows the variance of the respective estimators. While the variance of the naive estimator does not show convergence, the reweighted estimator converges to the analytical result.

Some care needs to be taken when using metropolis-type sampling on a distribution that decays only polynomially. In the tails, the probability density only varies very slowly with x, which means large steps have to be taken to ensure a reasonable autocorrelation time. This starkly contrasts exponentially decaying distributions, where large steps in the tails are generally suppressed. To address this issue, steps are proposed via a Gaussian with position-dependent variance $\sigma_s = \beta(1 + x)$, where β is a parameter tuned to give the desired acceptance rate. This ensures that the acceptance rate is roughly constant over the entire range of x and that the autocorrelation time is not dominated by the tails of the distribution.

Appendix B

Trapped fermions

B.1 Approach to the continuum

To demonstrate the proximity of our computations to the continuum theory, we compare the density profiles of the $N = 20$ balanced system with $g/\sqrt{\omega} = -3$ at various filling values. To this end, we keep the extent of the lattice fixed in unit of the harmonic oscillator $N_x/L_T = 20$, while varying the number of lattice sites between $N_x = 24$ and 200. Note, that this procedure requires a retuning of the coupling for each lattice as described in Section 6.1.1. The results are shown in Figure B.1. We find the density profile around the center of the trap to be well converged, even for lattices with only 40 point. In contrast, smaller lattice sizes show clear signs of saturation. As can be seen in the inset, the $Nx = 80$ and 200 lattices show good agreement in the tails over eight orders of magnitude,

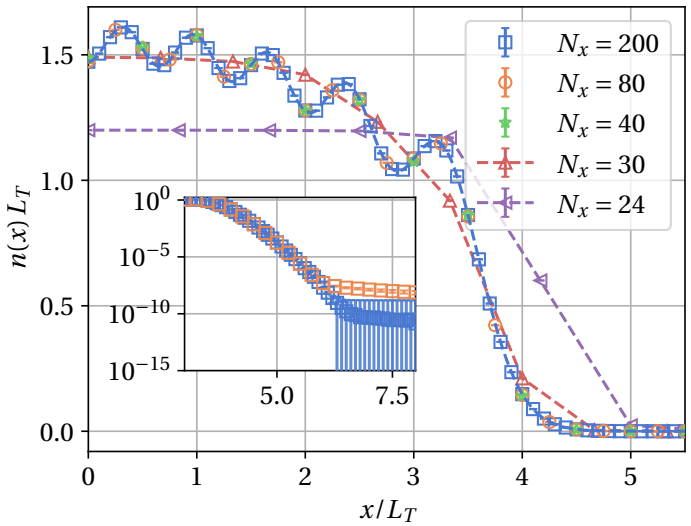

Figure B.1: Density profile for the $N = 20$ balanced system at various lattice sizes. The inset shows the tails of the density profile on a log scale.

after which they become indistinguishable from zero. We do not show the smaller lattices in the inset to avoid visual clutter.

Appendix C

The virial expansion

C.1 Basics

In the main text, several thermodynamic quantities were compared to results from virial expansion. This appendix briefly overviews the method and derivations for the aforementioned quantities.

The virial expansion uses the fact that in the dilute or high-temperature regime, thermodynamic systems are typically well described by considering only few-body terms. This is possible due to the increase in the inter-particle spacing or decrease in thermal wavelength. The grand canonical partition function is given as an expansion in the fugacity $z = e^{\beta\mu}$ as

$$\begin{aligned} Z &= \mathrm{Tr}\, e^{\beta\hat{H}-\beta\mu\hat{N}} \\ &= \sum_{N=0}^{\infty} z^N Z_N , \end{aligned} \tag{C.1}$$

Here Z_N is the canonical partition function of the system with N particles. Clearly, the range of validity of such an expansion is constrained by the value of $\beta\mu$, which governs the temperature and density. For $\beta\mu \to -\infty$ the main contribution is from the single particle canonical system, while higher particle contributions appear as the value is increased. The expansion is usually written in terms of the grand potential $-\beta\Omega = \log Z$, which is obtained by a straightforward expansion of the logarithm,

$$-\beta\Omega = Z_1 \sum_{n=0}^{\infty} b_n z^n . \tag{C.2}$$

Here, the b_n are called virial coefficients and can be found in terms of the N-body canonical partition functions by comparing eq. (C.1) and eq. (C.2). This yields $b_1 = 1$ for the first coefficient, while

the next three are given by

$$Z_1 b_2 = Z_2 - \frac{Z_1^2}{2}, \tag{C.3}$$

$$Z_1 b_3 = Z_3 - b_2 Z_1^2 - \frac{Z_1^3}{3!}, \tag{C.4}$$

$$Z_1 b_4 = Z_4 - (b_3 + \frac{b_2^2}{2}) Z_1^2 - b_2 \frac{Z_1^3}{2!} - \frac{Z_1^4}{4!}.. \tag{C.5}$$

As the non-interacting canonical systems can be solved exactly for any particle number, it is convenient to define the change in the virial coefficients due to interactions as

$$\Delta b_n = b_n - b_n^{(0)}. \tag{C.6}$$

where $b_n^{(0)}$ is the n-th virial coefficient of the non-interacting system in d dimensions

$$b_n^{(0)} = (-1)^{n+1} n^{-(d+1)/2}. \tag{C.7}$$

The relative coefficients can be obtained from the change in the partition functions by recasting eqs. (C.3) to (C.5) to depend on Δb_n and ΔQ_n, see e.g. [179].

C.2 Two flavour fermi gases

In the case of two flavor fermi gases, which we are interested in, there are two independent fugacity parameters, $z_\uparrow$ and $z_\downarrow$, for the two fermion species. This leads to a double expansion of the partition function and the grand potential, which takes the form

$$Z = \sum_{n,m=0}^{\infty} z_\uparrow^n z_\downarrow^m Z_{n,m}, \tag{C.8}$$

$$-\beta\Omega = Z_1 \sum_{n,m=0}^{\infty} z_\uparrow^n z_\downarrow^m b_{n,m}. \tag{C.9}$$

In the special case of the spin-balanced gas, the fugacity parameters are equal, and the expansion simplifies to the one described above with

$$b_2 = b_{1,1}, \tag{C.10}$$

$$b_3 = b_{2,1} + b_{1,2}, \tag{C.11}$$

$$b_4 = b_{3,1} + b_{2,2} + b_{1,3}. \tag{C.12}$$

The two-body system is given by the celebrated Beth-Uhlenbeck result, which, in two dimensions, reads

$$\Delta b_2 = e^{(\beta e_b)^2} - \int_0^\infty \frac{dy}{y} \frac{2e^{-(e_b \beta y)^2}}{\pi + 4\ln^2 y}, \tag{C.13}$$

where e_b is the two-body binding energy used to set the scales. The third coefficient is more difficult to obtain but has been computed in [197] via a diagrammatic approach. A later study using time discretization with a subsequent continuum extrapolation found good agreement with this diagrammatic result while also providing coefficients up to fifth order [179]. The latter study also introduced Padé approximants and Borel resummation to extend the range of convergence of the expansion.

C.3 Physical observables from the virial expansion

In order to obtain observables that can be compared to the lattice Monte Carlo results from the main text, we generally need to compute derivatives of the grand potential. To this end, the grand potential is split into a free and interaction part in analogy to the virial coefficients above

$$-\beta\Omega = -\beta\left(\Omega^{(0)} + \Delta\Omega\right) \tag{C.14}$$

$$-\beta\Delta\Omega = Z_1 \sum_{n,m=0}^{\infty} z_\uparrow^n z_\downarrow^m \Delta b_{n,m} \tag{C.15}$$

Perhaps the simplest observables to compute are the density and magnetization, which are obtained via derivatives with respect to the chemical potential $\mu = \left(\mu_\uparrow + \mu_\downarrow\right)/2$, and the external magnetic field $h = \left(\mu_\uparrow - \mu_\downarrow\right)/2$. One finds the density

$$n = n^{(0)} + \Delta n, \tag{C.16}$$

$$\begin{aligned}
\Delta n &= -\frac{1}{\beta}\frac{\partial\Delta\Omega}{\partial\beta\mu} \\
&= \frac{Z_1}{V}\sum_{n,m}(n+m)\, z_\uparrow^n z_\downarrow^m \Delta b_{n,m} \\
&= \frac{Z_1}{V}\Big[2\, z_\uparrow z_\downarrow \Delta b_{1,1} + 3\,(z_\uparrow z_\downarrow^2 + z_\uparrow^2 z_\downarrow)\Delta b_{1,2} + 4\left(z_\uparrow^2 z_\downarrow^2 \Delta b_{2,2} + z_\uparrow^1 z_\downarrow^3 \Delta b_{1,3} + z_\uparrow^3 z_\downarrow^1 \Delta b_{3,1}\right) \\
&\qquad + 5\left(z_\uparrow^3 z_\downarrow^2 \Delta b_{3,2} + z_\uparrow^2 z_\downarrow^3 \Delta b_{2,3} + z_\uparrow^1 z_\downarrow^4 \Delta b_{1,4} + z_\uparrow^4 z_\downarrow^1 \Delta b_{4,1}\right)\Big],
\end{aligned} \tag{C.17}$$

where $n^{(0)}$ is the density of the non-interacting system. The last line gives the expansion up to fifth order as used in the main text. The magnetization is given by

$$m = m^{(0)} + \Delta m, \tag{C.18}$$

$$\begin{aligned}
\Delta m &= -\frac{1}{\beta}\frac{\partial\Delta\Omega}{\partial h} \\
&= \frac{Z_1}{V}\sum_{n,m}(n-m)\, z_\uparrow^n z_\downarrow^m \Delta b_{n,m} \\
&= \frac{Z_1}{V}\Big[(z_\uparrow^2 z_\downarrow \Delta b_{2,1} - z_\uparrow z_\downarrow^2 \Delta b_{1,2} + z_\uparrow^3 z_\downarrow^2 \Delta b_{3,2} - z_\uparrow^2 z_\downarrow^3 \Delta b_{2,3}) + 2\left(z_\uparrow^1 z_\downarrow^3 + z_\uparrow^3 z_\downarrow^1\right)\Delta b_{3,1} \\
&\qquad + 3\left(z_\uparrow^4 z_\downarrow^1 - z_\uparrow^1 z_\downarrow^4\right)\Delta b_{4,1}\Big],
\end{aligned} \tag{C.19}$$

where $m^{(0)}$ is the non-interacting magnetization, and the last line again gives the expansion up to fifth order.

The rest of the observables can be obtained analogously by using the derivative and integral relations given in Section 6.3.

C.3.1 Padé approximants

To further enhance the range from which results can be obtained, Padé approximants can be used. This was done in [179] for the fifth-order virial expansion and spin imbalanced fermions in one

and three dimensions, respectively. In two dimensions no such analysis was performed. We therefore use the results from in the same work and compute Padé approximants for the density and magnetization.

The Padé approximant of order $[n/m]$ is defined as

$$R_{[n/m]}(z) = \frac{P(z)}{Q(z)}, \tag{C.20}$$

where $P(z)$ and $Q(z)$ are polynomials of degree n and m respectively. The coefficients of these polynomials are determined by requiring that the Taylor expansion of $R_{[n/m]}(z)$ around $z = 0$ matches the virial expansion up to order z^{n+m}, yielding an 'optimal' expansion in terms of the rational functions of chosen order. In comparison to the Taylor expansion, Padé approximants are often able to capture more of the asymptotic behavior of a function and are, therefore, better suited to extrapolate the virial expansion to higher orders. In practice, we compute the approximant for each observable separately while setting βh to a fixed value, leading to an expansion parameter

$$z = e^{\beta\mu} = e^{\beta(\mu_\uparrow+\mu_\downarrow)/2}, \tag{C.21}$$

which is just the central chemical potential.

Bibliography

[1] Felipe Attanasio, Marc Bauer, Renzo Kapust, and Jan M. Pawlowski. "Harmonically trapped fermions in one dimension: A finite-temperature lattice Monte Carlo study". In: *Phys. Rev. A* 109.3 (2024), p. 033305. DOI: 10.1103/PhysRevA.109.033305. arXiv: 2309.06766 [cond-mat.quant-gas].

[2] Felipe Attanasio, Marc Bauer, Lukas Kades, and Jan M. Pawlowski. "Searching for Yang-Lee zeros in O(N) models". In: *PoS* LATTICE2021 (2022), p. 223. DOI: 10.22323/1.396.0223. arXiv: 2111.12645 [hep-lat].

[3] Felipe Attanasio, Marc Bauer, and Jan M. Pawlowski. "Density profiles and correlations of harmonically trapped ultracold fermions via complex Langevin". In: *PoS* LATTICE2022 (2023), p. 363. DOI: 10.22323/1.430.0363. arXiv: 2212.11314 [cond-mat.quant-gas].

[4] Felipe Attanasio, Marc Bauer, and Jan M. Pawlowski. "Population Imbalanced Fermi Gases in the 2D BEC-BCS crossover". Manuscript in preparation.

[5] Marc Bauer, Renzo Kapust, Jan M. Pawlowski, and Finn L. Temmen. "Normalizing Flows for Inverse Renormalisation Group Steps". Manuscript in preparation.

[6] Felipe Attanasio, Marc Bauer, Jelle Dijkstra, Timoteo Lee, Jan M. Pawlowski, and Wolfram Pernice. "Speeding up Fermionic Lattice Calculations with Photonic Accelerated Inverters". In: (Jan. 2024). arXiv: 2401.14200 [hep-lat].

[7] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. "Sign problem in the numerical simulation of many-electron systems". In: *Phys. Rev. B* 41 (13 May 1990), pp. 9301–9307. DOI: 10.1103/PhysRevB.41.9301.

[8] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. "Observation of Bose-Einstein condensation in a dilute atomic vapor". In: *Science* 269 (1995), pp. 198–201. DOI: 10.1126/science.269.5221.198.

[9] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. "Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions". In: *Phys. Rev. Lett.* 75 (9 Aug. 1995), pp. 1687–1690. DOI: 10.1103/PhysRevLett.75.1687.

[10] K. B. Davis, M. -O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. "Bose-Einstein Condensation in a Gas of Sodium Atoms". In: *Phys. Rev. Lett.* 75 (22 Nov. 1995), pp. 3969–3973. DOI: 10.1103/PhysRevLett.75.3969.

[11] B. DeMarco, J. L. Bohn, J. P. Burke, M. Holland, and D. S. Jin. “Measurement of *p*-Wave Threshold Law Using Evaporatively Cooled Fermionic Atoms”. In: *Phys. Rev. Lett.* 82 (21 May 1999), pp. 4208–4211. DOI: 10.1103/PhysRevLett.82.4208.

[12] F. Schreck, L. Khaykovich, K. L. Corwin, G. Ferrari, T. Bourdel, J. Cubizolles, and C. Salomon. “Quasipure Bose-Einstein Condensate Immersed in a Fermi Sea”. In: *Phys. Rev. Lett.* 87 (8 Aug. 2001), p. 080403. DOI: 10.1103/PhysRevLett.87.080403.

[13] C. A. Stan, M. W. Zwierlein, C. H. Schunck, S. M. F. Raupach, and W. Ketterle. “Observation of Feshbach Resonances between Two Different Atomic Species”. In: *Phys. Rev. Lett.* 93 (14 Sept. 2004), p. 143001. DOI: 10.1103/PhysRevLett.93.143001.

[14] S. Inouye, J. Goldwin, M. L. Olsen, C. Ticknor, J. L. Bohn, and D. S. Jin. “Observation of Heteronuclear Feshbach Resonances in a Mixture of Bosons and Fermions”. In: *Phys. Rev. Lett.* 93 (18 Oct. 2004), p. 183201. DOI: 10.1103/PhysRevLett.93.183201.

[15] K. Osterloh, M. Baig, L. Santos, P. Zoller, and M. Lewenstein. “Cold Atoms in Non-Abelian Gauge Potentials: From the Hofstadter "Moth" to Lattice Gauge Theory”. In: *Phys. Rev. Lett.* 95 (1 June 2005), p. 010403. DOI: 10.1103/PhysRevLett.95.010403.

[16] C. Monroe, W. Swann, H. Robinson, and C. Wieman. “Very cold trapped atoms in a vapor cell”. In: *Phys. Rev. Lett.* 65 (13 Sept. 1990), pp. 1571–1574. DOI: 10.1103/PhysRevLett.65.1571.

[17] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. “Trapping of Neutral Sodium Atoms with Radiation Pressure”. In: *Phys. Rev. Lett.* 59 (23 Dec. 1987), pp. 2631–2634. DOI: 10.1103/PhysRevLett.59.2631.

[18] Naoto Masuhara, John M. Doyle, Jon C. Sandberg, Daniel Kleppner, Thomas J. Greytak, Harald F. Hess, and Greg P. Kochanski. “Evaporative Cooling of Spin-Polarized Atomic Hydrogen”. In: *Phys. Rev. Lett.* 61 (8 Aug. 1988), pp. 935–938. DOI: 10.1103/PhysRevLett.61.935.

[19] Cheng Chin, Rudolf Grimm, Paul Julienne, and Eite Tiesinga. “Feshbach resonances in ultracold gases”. In: *Rev. Mod. Phys.* 82 (2 Apr. 2010), pp. 1225–1286. DOI: 10.1103/RevModPhys.82.1225.

[20] W. Ketterle and M. W. Zwierlein. “Making, probing and understanding ultracold Fermi gases”. In: *Nuovo Cimento Rivista Serie* 31.5-6 (May 2008), pp. 247–422. DOI: 10.1393/ncr/i2008-10033-1. arXiv: 0801.2500 [cond-mat.other].

[21] Giancarlo Calvanese Strinati, Pierbiagio Pieri, Gerd Röpke, Peter Schuck, and Michael Urban. “The BCS–BEC crossover: From ultra-cold Fermi gases to nuclear systems”. In: *Physics Reports* 738 (2018). The BCS–BEC crossover: From ultra-cold Fermi gases to nuclear systems, pp. 1–76. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2018.02.004.

[22] Bhamidipati Chandrasekhar. “A NOTE ON THE MAXIMUM CRITICAL FIELD OF HIGH-FIELD SUPERCONDUCTORS”. In: *Applied Physics Letters* 1 (1962), pp. 7–8.

[23] A. M. Clogston. “Upper Limit for the Critical Field in Hard Superconductors”. In: *Phys. Rev. Lett.* 9 (6 Sept. 1962), pp. 266–267. DOI: 10.1103/PhysRevLett.9.266.

[24] Peter Fulde and Richard A. Ferrell. “Superconductivity in a Strong Spin-Exchange Field”. In: *Phys. Rev.* 135 (3A Aug. 1964), A550–A563. DOI: 10.1103/PhysRev.135.A550.

[25] A. I. Larkin and Y. N. Ovchinnikov. "Nonuniform state of superconductors". In: *Zh. Eksp. Teor. Fiz.* 47 (1964), pp. 1136–1146.

[26] C. H. Schunck, Y. Shin, A. Schirotzek, M. W. Zwierlein, and W. Ketterle. "Pairing Without Superfluidity: The Ground State of an Imbalanced Fermi Mixture". In: *Science* 316.5826 (2007), pp. 867–870. DOI: 10.1126/science.1140749. eprint: https://www.science.org/doi/pdf/10.1126/science.1140749.

[27] Martin W. Zwierlein, André Schirotzek, Christian H. Schunck, and Wolfgang Ketterle. "Fermionic Superfluidity with Imbalanced Spin Populations". In: *Science* 311.5760 (2006), pp. 492–496. DOI: 10.1126/science.1122318. eprint: https://www.science.org/doi/pdf/10.1126/science.1122318.

[28] Guthrie B. Partridge, Wenhui Li, Ramsey I. Kamar, Yean-an Liao, and Randall G. Hulet. "Pairing and Phase Separation in a Polarized Fermi Gas". In: *Science* 311 (2006), pp. 503–505. DOI: 10.1126/science.1122876. arXiv: cond-mat/0511752.

[29] G. B. Partridge, Wenhui Li, Y. A. Liao, R. G. Hulet, M. Haque, and H. T. C. Stoof. "Deformation of a Trapped Fermi Gas with Unequal Spin Populations". In: *Phys. Rev. Lett.* 97 (2006), p. 190407. DOI: 10.1103/PhysRevLett.97.190407. arXiv: cond-mat/0608455.

[30] W. Ong, Chingyun Cheng, I. Arakelyan, and J. E. Thomas. "Spin-Imbalanced Quasi-Two-Dimensional Fermi Gases". In: *Phys. Rev. Lett.* 114 (11 Mar. 2015), p. 110403. DOI: 10.1103/PhysRevLett.114.110403.

[31] Debayan Mitra, Peter T. Brown, Peter Schauß, Stanimir S. Kondov, and Waseem S. Bakr. "Phase Separation and Pair Condensation in a Spin-Imbalanced 2D Fermi Gas". In: *Phys. Rev. Lett.* 117 (9 Aug. 2016), p. 093601. DOI: 10.1103/PhysRevLett.117.093601.

[32] Lukas Rammelmüller, Andrew C. Loheac, Joaquín E. Drut, and Jens Braun. "Finite-temperature equation of state of polarized fermions at unitarity". In: *Phys. Rev. Lett.* 121.17 (2018), p. 173001. DOI: 10.1103/PhysRevLett.121.173001. arXiv: 1807.04664 [cond-mat.quant-gas].

[33] Lukas Rammelmüller, Yaqi Hou, Joaquín E. Drut, and Jens Braun. "Pairing and the spin susceptibility of the polarized unitary Fermi gas in the normal phase". In: *Phys. Rev. A* 103.4 (2021), p. 043330. DOI: 10.1103/PhysRevA.103.043330. arXiv: 2102.05911 [cond-mat.quant-gas].

[34] Felipe Attanasio, Lukas Rammelmüller, Joaquín E. Drut, and Jens Braun. "Pairing patterns in polarized unitary Fermi gases above the superfluid transition". In: *Phys. Rev. A* 105 (6 June 2022), p. 063317. DOI: 10.1103/PhysRevA.105.063317.

[35] M. J. Wolak, B. Grémaud, R. T. Scalettar, and G. G. Batrouni. "Pairing in a two-dimensional Fermi gas with population imbalance". In: *Phys. Rev. A* 86 (2 Aug. 2012), p. 023630. DOI: 10.1103/PhysRevA.86.023630.

[36] Heinz J. Rothe. *Lattice Gauge Theories : An Introduction (Fourth Edition)*. Vol. 43. World Scientific Publishing Company, 2012. ISBN: 978-981-4365-87-1, 978-981-4365-85-7. DOI: 10.1142/8229.

[37] I. Montvay and G. Munster. *Quantum fields on a lattice*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, Mar. 1997. ISBN: 978-0-521-59917-7, 978-0-511-87919-7. DOI: 10.1017/CBO9780511470783.

[38] Christof Gattringer and Christian B. Lang. *Quantum chromodynamics on the lattice*. Vol. 788. Berlin: Springer, 2010. ISBN: 978-3-642-01849-7, 978-3-642-01850-3. DOI: 10.1007/978-3-642-01850-3.

[39] A. Baumgaertner et al. *The Monte Carlo method in condensed matter physics*. Ed. by K. Binder. Vol. 71. Topics in applied physics. Berlin, Germany: Springer, 1992.

[40] Joaquin E. Drut and Amy N. Nicholson. "Lattice methods for strongly interacting many-body systems". In: *J. Phys. G* 40 (2013), p. 043101. DOI: 10.1088/0954-3899/40/4/043101. arXiv: 1208.6556 [cond-mat.stat-mech].

[41] Hao Shi, Simone Chiesa, and Shiwei Zhang. "Ground-state properties of strongly interacting Fermi gases in two dimensions". In: *Phys. Rev. A* 92 (3 Sept. 2015), p. 033603. DOI: 10.1103/PhysRevA.92.033603.

[42] Yuan-Yao He, Hao Shi, and Shiwei Zhang. "Precision Many-Body Study of the Berezinskii-Kosterlitz-Thouless Transition and Temperature-Dependent Properties in the Two-Dimensional Fermi Gas". In: *Phys. Rev. Lett.* 129 (7 Aug. 2022), p. 076403. DOI: 10.1103/PhysRevLett.129.076403.

[43] M. S. Albergo, G. Kanwar, and P. E. Shanahan. "Flow-based generative models for Markov chain Monte Carlo in lattice field theory". In: *Phys. Rev. D* 100.3 (2019), p. 034515. DOI: 10.1103/PhysRevD.100.034515. arXiv: 1904.12072 [hep-lat].

[44] Michael S. Albergo, Denis Boyda, Daniel C. Hackett, Gurtej Kanwar, Kyle Cranmer, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan. "Introduction to Normalizing Flows for Lattice Field Theory". In: (Jan. 2021). arXiv: 2101.08176 [hep-lat].

[45] Martin Luscher. "Trivializing maps, the Wilson flow and the HMC algorithm". In: *Commun. Math. Phys.* 293 (2010), pp. 899–919. DOI: 10.1007/s00220-009-0953-7. arXiv: 0907.5491 [hep-lat].

[46] Ryan Abbott et al. "Aspects of scaling and scalability for flow-based sampling of lattice QCD". In: *Eur. Phys. J. A* 59.11 (2023), p. 257. DOI: 10.1140/epja/s10050-023-01154-w. arXiv: 2211.07541 [hep-lat].

[47] Luigi Del Debbio, Joe Marsh Rossney, and Michael Wilson. "Efficient modeling of trivializing maps for lattice $\phi 4$ theory using normalizing flows: A first look at scalability". In: *Phys. Rev. D* 104.9 (2021), p. 094507. DOI: 10.1103/PhysRevD.104.094507. arXiv: 2105.12481 [hep-lat].

[48] M. J. Wolak, V. G. Rousseau, C. Miniatura, B. Grémaud, R. T. Scalettar, and G. G. Batrouni. "Finite-temperature quantum Monte Carlo study of the one-dimensional polarized Fermi gas". In: *Phys. Rev. A* 82 (1 July 2010), p. 013614. DOI: 10.1103/PhysRevA.82.013614.

[49] M. Suzuki. "Generalized Trotter's Formula and Systematic Approximants of Exponential Operators and Inner Derivations with Applications to Many Body Problems". In: *Commun. Math. Phys.* 51 (1976), pp. 183–190. DOI: 10.1007/BF01609348.

[50] Lukas Rammelmüller, Joaquín E. Drut, and Jens Braun. "Pairing patterns in one-dimensional spin- and mass-imbalanced Fermi gases". In: *SciPost Phys.* 9 (2020), p. 014. DOI: 10.21468/SciPostPhys.9.1.014. arXiv: 2003.06853 [cond-mat.quant-gas].

[51] J. Hubbard. “Calculation of Partition Functions”. In: *Phys. Rev. Lett.* 3 (2 July 1959), pp. 77–78. DOI: 10.1103/PhysRevLett.3.77.

[52] Aurel Bulgac, Joaquin E. Drut, and Piotr Magierski. “Quantum Monte Carlo simulations of the BCS-BEC crossover at finite temperature”. In: *Phys. Rev. A* 78 (2008), p. 023625. DOI: 10.1103/PhysRevA.78.023625. arXiv: 0803.3238 [cond-mat.stat-mech].

[53] Florian Ehmann, Joaquín E. Drut, and Jens Braun. “A lattice pairing-field approach to ultracold Fermi gases”. In: (Dec. 2022). arXiv: 2212.12298 [cond-mat.quant-gas].

[54] Lukas Rammelmüller. “Exploring imbalanced Fermi gases with stochastic quantization”. en. PhD thesis. Darmstadt: Technische Universität, 2020. DOI: https://doi.org/10.25534/tuprints-00011308.

[55] Fakher Assaad and H. Evertz. “World-line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons”. In: *Lecture Notes in Physics* 739 (Jan. 2008). DOI: 10.1007/978-3-540-74686-7_10.

[56] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar. “Monte Carlo calculations of coupled boson-fermion systems. I”. In: *Phys. Rev. D* 24 (8 Oct. 1981), pp. 2278–2286. DOI: 10.1103/PhysRevD.24.2278.

[57] J. E. Hirsch. “Two-dimensional Hubbard model: Numerical simulation study”. In: *Phys. Rev. B* 31 (7 Apr. 1985), pp. 4403–4419. DOI: 10.1103/PhysRevB.31.4403.

[58] Stephan Humeniuk. “Quantum Monte Carlo studies of strongly correlated systems for quantum simulators”. In: (2018).

[59] J Carlson, Stefano Gandolfi, Kevin E Schmidt, and Shiwei Zhang. “Auxiliary-field quantum Monte Carlo method for strongly paired fermions”. In: *Physical Review A* 84.6 (2011), p. 061602.

[60] Casey E. Berger, Joaquín E. Drut, and William J. Porter. “Hard-Wall and Non-Uniform Lattice Monte Carlo Approaches to One-Dimensional Fermi Gases in a Harmonic Trap”. In: *Comput. Phys. Commun.* 208 (2016), pp. 103–108. DOI: 10.1016/j.cpc.2016.08.005. arXiv: 1510.02976 [cond-mat.quant-gas].

[61] C. E. Berger, E. R. Anderson, and J. E. Drut. “Energy, contact, and density profiles of one-dimensional fermions in a harmonic trap via non-uniform lattice Monte Carlo”. In: *Phys. Rev. A* 91 (2015), p. 053618. DOI: 10.1103/PhysRevA.91.053618. arXiv: 1410.8103 [cond-mat.quant-gas].

[62] G. H. Lang, C. W. Johnson, S. E. Koonin, and W. E. Ormand. “Monte Carlo evaluation of path integrals for the nuclear shell model”. In: *Phys. Rev. C* 48 (1993), pp. 1518–1545. DOI: 10.1103/PhysRevC.48.1518. arXiv: nucl-th/9305009.

[63] W. E. Ormand, D. J. Dean, C. W. Johnson, G. H. Lang, and S. E. Koonin. “Demonstration of the auxiliary-field Monte Carlo approach for sd-shell nuclei”. In: *Phys. Rev. C* 49 (1994), pp. 1422–1427. DOI: 10.1103/PhysRevC.49.1422. arXiv: nucl-th/9307018.

[64] S. Rombouts and K. Heyde. “An Accurate and Efficient Algorithm for the Computation of the Characteristic Polynomial of a General Square Matrix”. In: *Journal of Computational Physics* 140.2 (1998), pp. 453–458. ISSN: 0021-9991. DOI: https://doi.org/10.1006/jcph.1998.5909.

[65] C. N. Gilbreth and Y. Alhassid. "Stabilizing Canonical-Ensemble Calculations in the Auxiliary-Field Monte Carlo Method". In: *Comput. Phys. Commun.* 188 (2015), pp. 1–6. DOI: 10.1016/j.cpc.2014.09.002. arXiv: 1402.3585 [physics.comp-ph].

[66] C. R. Shill and J. E. Drut. "Particle Projection Using a Complex Langevin Method". In: *EPJ Web Conf.* 175 (2018). Ed. by M. Della Morte, P. Fritzsch, E. Gámiz Sánchez, and C. Pena Ruano, p. 03003. DOI: 10.1051/epjconf/201817503003. arXiv: 1710.03247 [hep-lat].

[67] Edward Nelson. "Derivation of the Schrödinger Equation from Newtonian Mechanics". In: *Phys. Rev.* 150 (4 Oct. 1966), pp. 1079–1085. DOI: 10.1103/PhysRev.150.1079.

[68] Poul H. Damgaard and Helmuth Hüffel. "Stochastic quantization". In: *Physics Reports* 152.5 (1987), pp. 227–398. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(87)90144-X.

[69] G. Parisi and Yong-shi Wu. "Perturbation Theory Without Gauge Fixing". In: *Sci. Sin.* 24 (1981), p. 483.

[70] Mikio Namiki. *Stochastic quantization*. Vol. 9. Springer Science & Business Media, 1992.

[71] G. G. Batrouni. "Variations on the Langevin equation for lattice QCD with fermions". In: *Phys. Rev. D* 33 (6 Mar. 1986), pp. 1815–1818. DOI: 10.1103/PhysRevD.33.1815.

[72] Daniel Alvestad, Rasmus Larsen, and Alexander Rothkopf. "Stable solvers for real-time Complex Langevin". In: *JHEP* 08 (2021), p. 138. DOI: 10.1007/JHEP08(2021)138. arXiv: 2105.02735 [hep-lat].

[73] John R. Klauder. "Coherent-state Langevin equations for canonical quantum systems with applications to the quantized Hall effect". In: *Phys. Rev. A* 29 (4 Apr. 1984), pp. 2036–2047. DOI: 10.1103/PhysRevA.29.2036.

[74] F. Karsch and H. W. Wyld. "Complex Langevin Simulation of the SU(3) Spin Model with Nonzero Chemical Potential". In: *Phys. Rev. Lett.* 55 (21 Nov. 1985), pp. 2242–2245. DOI: 10.1103/PhysRevLett.55.2242.

[75] H. Q. Lin and J. E. Hirsch. "Monte Carlo versus Langevin methods for nonpositive definite weights". In: *Phys. Rev. B* 34 (3 Aug. 1986), pp. 1964–1967. DOI: 10.1103/PhysRevB.34.1964.

[76] John R Klauder and Wesley P Petersen. "Spectrum of certain non-self-adjoint operators and solutions of Langevin equations with complex drift". In: *Journal of statistical physics* 39 (1985), pp. 53–72.

[77] Felipe Attanasio, Benjamin Jäger, and Felix P. G. Ziegler. "QCD equation of state via the complex Langevin method". In: (Mar. 2022). arXiv: 2203.13144 [hep-lat].

[78] Gert Aarts, Felipe Attanasio, Benjamin Jäger, and Dénes Sexty. "The QCD phase diagram in the limit of heavy quarks using complex Langevin dynamics". In: *JHEP* 09 (2016), p. 087. DOI: 10.1007/JHEP09(2016)087. arXiv: 1606.05561 [hep-lat].

[79] Gert Aarts, Erhard Seiler, Denes Sexty, and Ion-Olimpiu Stamatescu. "Complex Langevin dynamics and zeroes of the fermion determinant". In: *JHEP* 05 (2017). [Erratum: JHEP 01, 128 (2018)], p. 044. DOI: 10.1007/JHEP05(2017)044. arXiv: 1701.02322 [hep-lat].

[80] Denes Sexty. "Progress in complex Langevin simulations of full QCD at non-zero density". In: *Nucl. Phys. A* 931 (2014). Ed. by Peter Braun-Munzinger, Bengt Friman, and Johanna Stachel, pp. 856–860. DOI: 10.1016/j.nuclphysa.2014.09.029. arXiv: 1408.6767 [hep-lat].

[81] Andrew C. Loheac and Joaquin E. Drut. "Third-order perturbative lattice and complex Langevin analyses of the finite-temperature equation of state of nonrelativistic fermions in one dimension". In: *Phys. Rev. D* 95.9 (2017), p. 094502. DOI: 10.1103/PhysRevD.95.094502. arXiv: 1702.04666 [hep-lat].

[82] Felipe Attanasio and Joaquín E. Drut. "Thermodynamics of spin-orbit-coupled bosons in two dimensions from the complex Langevin method". In: *Phys. Rev. A* 101.3 (2020), p. 033617. DOI: 10.1103/PhysRevA.101.033617. arXiv: 1908.02715 [cond-mat.quant-gas].

[83] Tomoya Hayata and Arata Yamamoto. "Complex Langevin simulation of quantum vortices in a Bose-Einstein condensate". In: *Phys. Rev. A* 92.4 (2015), p. 043628. DOI: 10.1103/PhysRevA.92.043628. arXiv: 1411.5195 [cond-mat.quant-gas].

[84] Philipp Heinen and Thomas Gasenzer. "Simulating the Berezinskii-Kosterlitz-Thouless transition with the complex Langevin algorithm". In: *Phys. Rev. A* 108.5 (2023), p. 053311. DOI: 10.1103/PhysRevA.108.053311. arXiv: 2304.05699 [cond-mat.quant-gas].

[85] Philipp Heinen and Thomas Gasenzer. "Complex Langevin approach to interacting Bose gases". In: *Phys. Rev. A* 106.6 (2022), p. 063308. DOI: 10.1103/PhysRevA.106.063308. arXiv: 2204.10661 [cond-mat.quant-gas].

[86] Gert Aarts, Frank A. James, Erhard Seiler, and Ion-Olimpiu Stamatescu. "Adaptive stepsize and instabilities in complex Langevin dynamics". In: *Phys. Lett. B* 687 (2010), pp. 154–159. DOI: 10.1016/j.physletb.2010.03.012. arXiv: 0912.0617 [hep-lat].

[87] Gert Aarts, Frank A. James, Erhard Seiler, and Ion-Olimpiu Stamatescu. "Complex Langevin: Etiology and Diagnostics of its Main Problem". In: *Eur. Phys. J. C* 71 (2011), p. 1756. DOI: 10.1140/epjc/s10052-011-1756-5. arXiv: 1101.3270 [hep-lat].

[88] Keitaro Nagata, Jun Nishimura, and Shinji Shimasaki. "Argument for justification of the complex Langevin method and the condition for correct convergence". In: *Phys. Rev. D* 94.11 (2016), p. 114515. DOI: 10.1103/PhysRevD.94.114515. arXiv: 1606.07627 [hep-lat].

[89] Felipe Attanasio and Benjamin Jäger. "Dynamical stabilisation of complex Langevin simulations of QCD". In: *Eur. Phys. J. C* 79.1 (2019), p. 16. DOI: 10.1140/epjc/s10052-018-6512-7. arXiv: 1808.04400 [hep-lat].

[90] Gert Aarts, Erhard Seiler, and Ion-Olimpiu Stamatescu. "The Complex Langevin method: When can it be trusted?" In: *Phys. Rev. D* 81 (2010), p. 054508. DOI: 10.1103/PhysRevD.81.054508. arXiv: 0912.3360 [hep-lat].

[91] Gert Aarts and Frank A. James. "On the convergence of complex Langevin dynamics: The Three-dimensional XY model at finite chemical potential". In: *JHEP* 08 (2010), p. 020. DOI: 10.1007/JHEP08(2010)020. arXiv: 1005.3468 [hep-lat].

[92] Manuel Scherzer, Erhard Seiler, Dénes Sexty, and Ion-Olimpiu Stamatescu. "Complex Langevin and boundary terms". In: *Phys. Rev. D* 99.1 (2019), p. 014512. DOI: 10.1103/PhysRevD.99.014512. arXiv: 1808.05187 [hep-lat].

[93] M. Scherzer, E. Seiler, D. Sexty, and I. -O. Stamatescu. "Controlling Complex Langevin simulations of lattice models by boundary term analysis". In: *Phys. Rev. D* 101.1 (2020), p. 014501. DOI: 10.1103/PhysRevD.101.014501. arXiv: 1910.09427 [hep-lat].

[94] Erhard Seiler, Dénes Sexty, and Ion-Olimpiu Stamatescu. "Complex Langevin: Correctness criteria, boundary terms, and spectrum". In: *Phys. Rev. D* 109.1 (2024), p. 014509. DOI: 10.1103/PhysRevD.109.014509. arXiv: 2304.00563 [hep-lat].

[95] C. N. Yang and T. D. Lee. "Statistical Theory of Equations of State and Phase Transitions. I. Theory of Condensation". In: *Phys. Rev.* 87 (3 Aug. 1952), pp. 404–409. DOI: 10.1103/PhysRev.87.404.

[96] T. D. Lee and C. N. Yang. "Statistical Theory of Equations of State and Phase Transitions. II. Lattice Gas and Ising Model". In: *Phys. Rev.* 87 (3 Aug. 1952), pp. 410–419. DOI: 10.1103/PhysRev.87.410.

[97] M. A. Stephanov. "QCD critical point and complex chemical potential singularities". In: *Phys. Rev. D* 73 (9 May 2006), p. 094508. DOI: 10.1103/PhysRevD.73.094508.

[98] Matteo Giordano and Attila Pásztor. "Reliable estimation of the radius of convergence in finite density QCD". In: *Phys. Rev. D* 99.11 (2019), p. 114510. DOI: 10.1103/PhysRevD.99.114510. arXiv: 1904.01974 [hep-lat].

[99] Swagato Mukherjee and Vladimir Skokov. "Universality driven analytic structure of the QCD crossover: Radius of convergence in the baryon chemical potential". In: *Phys. Rev. D* 103 (7 Apr. 2021), p. L071501. DOI: 10.1103/PhysRevD.103.L071501.

[100] Matteo Giordano, Kornel Kapas, Sandor D. Katz, Daniel Nogradi, and Attila Pasztor. "Radius of convergence in lattice QCD at finite μ_B with rooted staggered fermions". In: *Phys. Rev. D* 101.7 (2020). [Erratum: Phys.Rev.D 104, 119901 (2021)], p. 074511. DOI: 10.1103/PhysRevD.101.074511. arXiv: 1911.00043 [hep-lat].

[101] Andrew Connelly, Gregory Johnson, Fabian Rennecke, and Vladimir Skokov. "Universal Location of the Yang-Lee Edge Singularity in $O(N)$ Theories". In: *Phys. Rev. Lett.* 125.19 (2020), p. 191602. DOI: 10.1103/PhysRevLett.125.191602. arXiv: 2006.12541 [cond-mat.stat-mech].

[102] Gokce Basar. "Universality, Lee-Yang Singularities, and Series Expansions". In: *Phys. Rev. Lett.* 127.17 (2021), p. 171603. DOI: 10.1103/PhysRevLett.127.171603. arXiv: 2105.08080 [hep-th].

[103] Guido Nicotra, Petros Dimopoulos, Lorenzo Dini, Francesco Di Renzo, Jishnu Goswami, Christian Schmidt, Simran Singh, Kevin Zambello, and Felix Ziesche. "Lee-Yang edge singularities in 2+1 flavor QCD with imaginary chemical potential." In: *PoS* LATTICE2021 (2022), p. 260. DOI: 10.22323/1.396.0260. arXiv: 2111.05630 [hep-lat].

[104] L. L. Salcedo. "Does the complex Langevin method give unbiased results?" In: *Phys. Rev. D* 94.11 (2016), p. 114505. DOI: 10.1103/PhysRevD.94.114505. arXiv: 1611.06390 [hep-lat].

[105] Edward Witten. "A New Look At The Path Integral Of Quantum Mechanics". In: *Surveys in Differential Geometry* 15 (Sept. 2010). DOI: 10.4310/SDG.2010.v15.n1.a11.

[106] Marco Cristoforetti, Francesco Di Renzo, and Luigi Scorzato. "New approach to the sign problem in quantum field theories: High density QCD on a Lefschetz thimble". In: *Phys. Rev. D* 86 (7 Oct. 2012), p. 074506. DOI: 10.1103/PhysRevD.86.074506.

[107] J. Bloch, J. Glesaaen, J. J. M. Verbaarschot, and S. Zafeiropoulos. "Complex Langevin Simulation of a Random Matrix Model at Nonzero Chemical Potential". In: *JHEP* 03 (2018), p. 015. DOI: 10.1007/JHEP03(2018)015. arXiv: 1712.07514 [hep-lat].

[108] Jun Nishimura and Shinji Shimasaki. "New Insights into the Problem with a Singular Drift Term in the Complex Langevin Method". In: *Phys. Rev. D* 92.1 (2015), p. 011501. DOI: 10.1103/PhysRevD.92.011501. arXiv: 1504.08359 [hep-lat].

[109] Daniel Alvestad, Rasmus Larsen, and Alexander Rothkopf. "Towards learning optimized kernels for complex Langevin". In: *JHEP* 04 (2023), p. 057. DOI: 10.1007/JHEP04(2023)057. arXiv: 2211.15625 [hep-lat].

[110] Daniel Alvestad, Alexander Rothkopf, and Dénes Sexty. "Lattice real-time simulations with learned optimal kernels". In: *Phys. Rev. D* 109.3 (2024), p. L031502. DOI: 10.1103/PhysRevD.109.L031502. arXiv: 2310.08053 [hep-lat].

[111] Kirill Boguslavski, Paul Hotzy, and David I. Müller. "Real-time correlators in 3+1D thermal lattice gauge theory". In: (Dec. 2023). arXiv: 2312.03063 [hep-lat].

[112] Kirill Boguslavski, Paul Hotzy, and David I. Müller. "Stabilizing complex Langevin for real-time gauge theories with an anisotropic kernel". In: *JHEP* 06 (2023), p. 011. DOI: 10.1007/JHEP06(2023)011. arXiv: 2212.08602 [hep-lat].

[113] E.Y. LOH and J.E. GUBERNATIS. "CHAPTER 4 - Stable Numerical Simulations of Models of Interacting Electrons in Condensed-Matter Physics". In: *Electronic Phase Transitions*. Ed. by W. HANKE and Yu.V. KOPAEV. Vol. 32. Modern Problems in Condensed Matter Sciences. Elsevier, 1992, pp. 177–235. DOI: https://doi.org/10.1016/B978-0-444-88885-3.50009-3.

[114] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar. "Numerical study of the two-dimensional Hubbard model". In: *Phys. Rev. B* 40 (1989), pp. 506–516. DOI: 10.1103/PhysRevB.40.506.

[115] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, R. L. Sugar, and S. R. White. "Stable Matrix-Multiplication Algorithms for Low-Temperature Numerical Simulations of Fermions". In: *Interacting Electrons in Reduced Dimensions*. Ed. by Dionys Baeriswyl and David K. Campbell. Boston, MA: Springer US, 1989, pp. 55–60. ISBN: 978-1-4613-0565-1. DOI: 10.1007/978-1-4613-0565-1_8.

[116] Carsten Bauer. "Fast and stable determinant quantum Monte Carlo". In: *SciPost Phys. Core* 2 (2020), p. 011. DOI: 10.21468/SciPostPhysCore.2.2.011.

[117] E. Y. LOH, J. E. GUBERNATIS, R. T. SCALETTAR, S. R. WHITE, D. J. SCALAPINO, and R. L. SUGAR. "NUMERICAL STABILITY AND THE SIGN PROBLEM IN THE DETERMINANT QUANTUM MONTE CARLO METHOD". In: *International Journal of Modern Physics C* 16.08 (2005), pp. 1319–1327. DOI: 10.1142/S0129183105007911. eprint: https://doi.org/10.1142/S0129183105007911.

[118] Cagin Yunus and William Detmold. "Infinite variance in Monte Carlo sampling of lattice field theories". In: *Phys. Rev. D* 106.9 (2022), p. 094506. DOI: 10.1103/PhysRevD.106.094506. arXiv: 2205.01001 [hep-lat].

[119] Andrei Alexandru, Paulo F. Bedaque, Andrea Carosso, and Hyunwoo Oh. "Infinite variance problem in fermion models". In: *Phys. Rev. D* 107.9 (2023), p. 094502. DOI: 10.1103/PhysRevD.107.094502. arXiv: 2211.06419 [hep-lat].

[120] Hao Shi and Shiwei Zhang. "Infinite Variance in Fermion Quantum Monte Carlo Calculations". In: *Phys. Rev. E* 93.3 (2016), p. 033303. DOI: 10.1103/PhysRevE.93.033303. arXiv: 1511.04084 [physics.comp-ph].

[121] C. N. Gilbreth, S. Jensen, and Y. Alhassid. "Reducing the complexity of finite-temperature auxiliary-field quantum Monte Carlo". In: *Comput. Phys. Commun.* 264 (2021), p. 107952. DOI: 10.1016/j.cpc.2021.107952. arXiv: 1907.10596 [physics.comp-ph].

[122] Yuan-Yao He, Hao Shi, and Shiwei Zhang. "Reaching the Continuum Limit in Finite-Temperature Ab Initio Field-Theory Computations in Many-Fermion Systems". In: *Phys. Rev. Lett.* 123.13 (2019), p. 136402. DOI: 10.1103/PhysRevLett.123.136402. arXiv: 1906.02247 [cond-mat.str-el].

[123] Joonho Lee, Miguel A. Morales, and Fionn D. Malone. "A phaseless auxiliary-field quantum Monte Carlo perspective on the uniform electron gas at finite temperatures: Issues, observations, and benchmark study". In: *The Journal of Chemical Physics* 154.6 (Feb. 2021), p. 064109. ISSN: 0021-9606. DOI: 10.1063/5.0041378. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0041378/15583713/064109_1_online.pdf.

[124] Maksim Ulybyshev and Fakher Assaad. "Mitigating spikes in fermion Monte Carlo methods by reshuffling measurements". In: *Phys. Rev. E* 106.2 (2022), p. 025318. DOI: 10.1103/PhysRevE.106.025318. arXiv: 2112.05260 [cond-mat.str-el].

[125] Alexander L. Gaunt, Tobias F. Schmidutz, Igor Gotlibovych, Robert P. Smith, and Zoran Hadzibabic. "Bose-Einstein Condensation of Atoms in a Uniform Potential". In: *Phys. Rev. Lett.* 110 (20 May 2013), p. 200406. DOI: 10.1103/PhysRevLett.110.200406.

[126] Lauriane Chomaz, Laura Corman, Tom Bienaimé, Rémi Desbuquois, Christof Weitenberg, Sylvain Nascimbène, Jérôme Beugnon, and Jean Dalibard. "Emergence of coherence via transverse condensation in a uniform quasi-two-dimensional Bose gas". In: *Nature Communications* 6.1 (Jan. 2015). ISSN: 2041-1723. DOI: 10.1038/ncomms7162.

[127] Klaus Hueck, Niclas Luick, Lennart Sobirey, Jonas Siegl, Thomas Lompe, and Henning Moritz. "Two-Dimensional Homogeneous Fermi Gases". In: *Phys. Rev. Lett.* 120 (6 Feb. 2018), p. 060402. DOI: 10.1103/PhysRevLett.120.060402.

[128] Nir Navon, Robert P. Smith, and Zoran Hadzibabic. "Quantum gases in optical boxes". In: *Nature Phys.* 17.12 (2021), pp. 1334–1341. DOI: 10.1038/s41567-021-01403-z. arXiv: 2106.09716 [cond-mat.quant-gas].

[129] Christopher Oliver, Aaron Smith, Thomas Easton, Grazia Salerno, Vera Guarrera, Nathan Goldman, Giovanni Barontini, and Hannah M. Price. "Bloch oscillations along a synthetic dimension of atomic trap states". In: *Phys. Rev. Res.* 5 (3 July 2023), p. 033001. DOI: 10.1103/PhysRevResearch.5.033001.

[130] Philipp Lunt, Paul Hill, Johannes Reiter, Philipp M. Preiss, Maciej Gałka, and Selim Jochim. "Realization of a Laughlin state of two rapidly rotating fermions". In: (Feb. 2024). arXiv: 2402.14814 [cond-mat.quant-gas].

[131] Marvin Holten, Luca Bayha, Keerthan Subramanian, Sandra Brandstetter, Carl Heintze, Philipp Lunt, Philipp M. Preiss, and Selim Jochim. "Observation of Cooper pairs in a mesoscopic two-dimensional Fermi gas". In: *Nature* 606.7913 (2022), pp. 287–291. DOI: 10.1038/s41586-022-04678-1. arXiv: 2109.11511 [cond-mat.quant-gas].

[132] Luca Bayha, Marvin Holten, Ralf Klemt, Keerthan Subramanian, Johannes Bjerlin, Stephanie M. Reimann, Georg M. Bruun, Philipp M. Preiss, and Selim Jochim. "Observing the emergence of a quantum phase transition shell by shell". In: *Nature* 587.7835 (Nov. 2020), pp. 583–587. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2936-y.

[133] Sandra Brandstetter et al. "Emergent hydrodynamic behaviour of few strongly interacting fermions". In: (Aug. 2023). arXiv: 2308.09699 [cond-mat.quant-gas].

[134] E. R. Anderson and J. E. Drut. "Pressure, compressibility, and contact of the two-dimensional attractive Fermi gas". In: *Phys. Rev. Lett.* 115.11 (2015), p. 115301. DOI: 10.1103/PhysRevLett.115.115301. arXiv: 1505.01525 [cond-mat.quant-gas].

[135] R. Rossi, T. Ohgoe, E. Kozik, N. Prokof'ev, B. Svistunov, K. Van Houcke, and F. Werner. "Contact and Momentum Distribution of the Unitary Fermi Gas". In: *Phys. Rev. Lett.* 121 (13 Sept. 2018), p. 130406. DOI: 10.1103/PhysRevLett.121.130406.

[136] S. Ramachandran, S. Jensen, and Y. Alhassid. "Pseudogap effects in the strongly correlated regime of the two-dimensional Fermi gas". In: (Dec. 2022). arXiv: 2212.14880 [cond-mat.quant-gas].

[137] Tomasz Sowi ński, Tobias Grass, Omjyoti Dutta, and Maciej Lewenstein. "Few interacting fermions in a one-dimensional harmonic trap". In: *Phys. Rev. A* 88 (3 Sept. 2013), p. 033607. DOI: 10.1103/PhysRevA.88.033607.

[138] Pino D'Amico and Massimo Rontani. "Pairing of a few Fermi atoms in one dimension". In: *Phys. Rev. A* 91 (4 Apr. 2015), p. 043610. DOI: 10.1103/PhysRevA.91.043610.

[139] Lukas Rammelmüller, David Huber, Matija Čufar, Joachim Brand, Hans-Werner Hammer, and Artem G. Volosniev. "Magnetic impurity in a one-dimensional few-fermion system". In: *SciPost Phys.* 14.1 (2023), p. 006. DOI: 10.21468/SciPostPhys.14.1.006. arXiv: 2204.01606 [cond-mat.quant-gas].

[140] Tomasz Grining, Michał Tomza, Michał Lesiuk, Michał Przybytek, Monika Musiał, Robert Moszynski, Maciej Lewenstein, and Pietro Massignan. "Crossover between few and many fermions in a harmonic trap". In: *Phys. Rev. A* 92 (6 Dec. 2015), p. 061601. DOI: 10.1103/PhysRevA.92.061601.

[141] Tomasz Grining, Michał Tomza, Michał Lesiuk, Michał Przybytek, Monika Musiał, Pietro Massignan, Maciej Lewenstein, and Robert Moszynski. "Many interacting fermions in a one-dimensional harmonic trap: a quantum-chemical treatment". In: *New Journal of Physics* 17.11 (Oct. 2015), p. 115001. DOI: 10.1088/1367-2630/17/11/115001.

[142] Michele Casula, D. M. Ceperley, and Erich J. Mueller. "Quantum Monte Carlo study of one-dimensional trapped fermions with attractive contact interactions". In: *Phys. Rev. A* 78 (3 Sept. 2008), p. 033607. DOI: 10.1103/PhysRevA.78.033607.

[143] Daniel Pecak and Tomasz Sowinski. "Unconventional pairing in few-fermion systems at finite temperature". In: *Scientific Reports* 12.1 (Oct. 2022). ISSN: 2045-2322. DOI: 10.1038/s41598-022-22411-w.

[144] Thomas Busch, Berthold-Georg Englert, Kazimierz Rzażewski, and Martin Wilkens. "Two Cold Atoms in a Harmonic Trap". In: *Found. Phys.* 28.4 (1998), pp. 549–559. DOI: 10.1023/a:1018705520999.

[145] Pino D'Amico and Massimo Rontani. "Three interacting atoms in a one-dimensional trap: a benchmark system for computational approaches". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 47.6 (Feb. 2014), p. 065303. DOI: 10.1088/0953-4075/47/6/065303.

[146] Jon Flower, Steve W. Otto, and Sean Callahan. "Complex Langevin equations and lattice gauge theory". In: *Phys. Rev. D* 34 (2 July 1986), pp. 598–604. DOI: 10.1103/PhysRevD.34.598.

[147] Hersh Singh and Shailesh Chandrasekharan. "Few-body physics on a spacetime lattice in the worldline approach". In: *Phys. Rev. D* 99.7 (2019), p. 074511. DOI: 10.1103/PhysRevD.99.074511. arXiv: 1812.05080 [hep-lat].

[148] Lukas Rammelmüller, William J. Porter, Joaquín E. Drut, and Jens Braun. "Surmounting the sign problem in non-relativistic calculations: a case study with mass-imbalanced fermions". In: *Phys. Rev. D* 96.9 (2017), p. 094506. DOI: 10.1103/PhysRevD.96.094506. arXiv: 1708.03149 [cond-mat.quant-gas].

[149] Richard T. Scalettar, Reinhard M. Noack, and Rajiv R. P. Singh. "Ergodicity at large couplings with the determinant Monte Carlo algorithm". In: *Phys. Rev. B* 44 (19 Nov. 1991), pp. 10502–10507. DOI: 10.1103/PhysRevB.44.10502.

[150] G. Zürn, A. N. Wenz, S. Murmann, A. Bergschneider, T. Lompe, and S. Jochim. "Pairing in Few-Fermion Systems with Attractive Interactions". In: *Phys. Rev. Lett.* 111 (17 Oct. 2013), p. 175302. DOI: 10.1103/PhysRevLett.111.175302.

[151] Adam Richie-Halford, Joaquín E. Drut, and Aurel Bulgac. "Emergence of a pseudogap in the BCS-BEC crossover". In: *Phys. Rev. Lett.* 125.6 (2020), p. 060403. DOI: 10.1103/PhysRevLett.125.060403. arXiv: 2004.05014 [cond-mat.quant-gas].

[152] C. N. Gilbreth and Y. Alhassid. "Pair condensation in a finite trapped Fermi gas". In: *Phys. Rev. A* 88.6 (2013), p. 063643. DOI: 10.1103/PhysRevA.88.063643. arXiv: 1210.4131 [cond-mat.quant-gas].

[153] T. Duguet, P. Bonche, P.-H. Heenen, and J. Meyer. "Pairing correlations. II. Microscopic analysis of odd-even mass staggering in nuclei". In: *Phys. Rev. C* 65 (1 Dec. 2001), p. 014311. DOI: 10.1103/PhysRevC.65.014311.

[154] Olga Goulko and Matthew Wingate. "Thermodynamics of balanced and slightly spin-imbalanced Fermi gases at unitarity". In: *Physical Review A* 82.5 (Nov. 2010). ISSN: 1094-1622. DOI: 10.1103/physreva.82.053621.

[155] Andrei Alexandru, Paulo F. Bedaque, and Neill C. Warrington. "Spin polarized nonrelativistic fermions in 1+1 dimensions". In: *Phys. Rev. D* 98.5 (2018), p. 054514. DOI: 10.1103/PhysRevD.98.054514. arXiv: 1805.00125 [hep-lat].

[156] Takahiro M. Doi, Hiroyuki Tajima, and Shoichiro Tsutsui. "Complex Langevin study for polarons in a one-dimensional two-component Fermi gas with attractive contact interactions". In: *Phys. Rev. Res.* 3.3 (2021), p. 033180. DOI: 10.1103/PhysRevResearch.3.033180. arXiv: 2105.11072 [cond-mat.quant-gas].

[157] Jacek Dobrzyniecki, Giuliano Orso, and Tomasz Sowiński. "Unconventional pairing in few-fermion systems tuned by external confinement". In: *Phys. Rev. Res.* 3.4 (2021), p. 043105. DOI: 10.1103/PhysRevResearch.3.043105. arXiv: 2105.12519 [cond-mat.quant-gas].

[158] D. N. Basov and T. Timusk. "Electrodynamics of high-T_c superconductors". In: *Rev. Mod. Phys.* 77 (2 Aug. 2005), pp. 721–779. DOI: 10.1103/RevModPhys.77.721.

[159] Masatoshi Sato, Yoshiro Takahashi, and Satoshi Fujimoto. "Non-Abelian Topological Order in *s*-Wave Superfluids of Ultracold Fermionic Atoms". In: *Phys. Rev. Lett.* 103 (2 July 2009), p. 020401. DOI: 10.1103/PhysRevLett.103.020401.

[160] P. A. Murthy, I. Boettcher, L. Bayha, M. Holzmann, D. Kedar, M. Neidig, M. G. Ries, A. N. Wenz, G. Zürn, and S. Jochim. "Observation of the Berezinskii-Kosterlitz-Thouless Phase Transition in an Ultracold Fermi as". In: *Phys. Rev. Lett.* 115 (1 June 2015), p. 010401. DOI: 10.1103/PhysRevLett.115.010401.

[161] M. G. Ries, A. N. Wenz, G. Zürn, L. Bayha, I. Boettcher, D. Kedar, P. A. Murthy, M. Neidig, T. Lompe, and S. Jochim. "Observation of Pair Condensation in the Quasi-2D BEC-BCS Crossover". In: *Phys. Rev. Lett.* 114 (23 June 2015), p. 230401. DOI: 10.1103/PhysRevLett.114.230401.

[162] Lennart Sobirey, Niclas Luick, Markus Bohlen, Hauke Biss, Henning Moritz, and Thomas Lompe. "Observation of superfluidity in a strongly correlated two-dimensional Fermi gas". In: *Science* 372.6544 (2021), pp. 844–846. DOI: 10.1126/science.abc8793. eprint: https://www.science.org/doi/pdf/10.1126/science.abc8793.

[163] Brendan C. Mulkerin, Lianyi He, Paul Dyke, Chris J. Vale, Xia-Ji Liu, and Hui Hu. "Superfluid density and critical velocity near the Berezinskii-Kosterlitz-Thouless transition in a two-dimensional strongly interacting Fermi gas". In: *Phys. Rev. A* 96 (5 Nov. 2017), p. 053608. DOI: 10.1103/PhysRevA.96.053608.

[164] G. Bighin and L. Salasnich. "Finite-temperature quantum fluctuations in two-dimensional Fermi superfluids". In: *Phys. Rev. B* 93 (1 Jan. 2016), p. 014519. DOI: 10.1103/PhysRevB.93.014519.

[165] Marianne Bauer, Meera M. Parish, and Tilman Enss. "Universal Equation of State and Pseudogap in the Two-Dimensional Fermi Gas". In: *Phys. Rev. Lett.* 112 (13 Apr. 2014), p. 135302. DOI: 10.1103/PhysRevLett.112.135302.

[166] I. Boettcher, L. Bayha, D. Kedar, P. A. Murthy, M. Neidig, M. G. Ries, A. N. Wenz, G. Zürn, S. Jochim, and T. Enss. "Equation of State of Ultracold Fermions in the 2D BEC-BCS Crossover Region". In: *Phys. Rev. Lett.* 116 (4 Jan. 2016), p. 045303. DOI: 10.1103/PhysRevLett.116.045303.

[167] K. Fenech, P. Dyke, T. Peppler, M. G. Lingham, S. Hoinka, H. Hu, and C. J. Vale. "Thermodynamics of an Attractive 2D Fermi Gas". In: *Phys. Rev. Lett.* 116 (4 Jan. 2016), p. 045302. DOI: 10.1103/PhysRevLett.116.045302.

[168] Rodrigo A. Fontenele, Natanael C. Costa, Raimundo R. dos Santos, and Thereza Paiva. "Two-dimensional attractive Hubbard model and the BCS-BEC crossover". In: *Phys. Rev. B* 105 (18 May 2022), p. 184502. DOI: 10.1103/PhysRevB.105.184502.

[169] Thereza Paiva, Raimundo R. dos Santos, R. T. Scalettar, and P. J. H. Denteneer. "Critical temperature for the two-dimensional attractive Hubbard model". In: *Phys. Rev. B* 69 (18 May 2004), p. 184501. DOI: 10.1103/PhysRevB.69.184501.

[170] Lukas Rammelmüller, William J. Porter, and Joaquín E. Drut. "Ground state of the two-dimensional attractive Fermi gas: Essential properties from few to many body". In: *Phys. Rev. A* 93 (3 Mar. 2016), p. 033639. DOI: 10.1103/PhysRevA.93.033639.

[171] Jesper Levinsen and Meera M. Parish. "Strongly interacting two-dimensional Fermi gases". In: (Aug. 2014). DOI: 10.1142/9789814667746_0001. arXiv: 1408.2737 [cond-mat.quant-gas].

[172] Félix Werner and Yvan Castin. "General relations for quantum gases in two and three dimensions: Two-component fermions". In: *Phys. Rev. A* 86 (1 July 2012), p. 013626. DOI: 10.1103/PhysRevA.86.013626.

[173] M. D. Hoffman, P. Javernick, A. C. Loheac, W. J. Porter, E. R. Anderson, and J. E. Drut. "Universality in one-dimensional fermions at finite temperature: Density, pressure, compressibility, and contact". In: *Phys. Rev. A* 91.3 (2015), p. 033618. DOI: 10.1103/PhysRevA.91.033618. arXiv: 1410.7370 [cond-mat.quant-gas].

[174] Cole Miles, Benjamin Cohen-Stead, Owen Bradley, Steven Johnston, Richard Scalettar, and Kipton Barros. "Dynamical tuning of the chemical potential to achieve a target particle number in grand canonical Monte Carlo simulations". In: *Phys. Rev. E* 105 (4 Apr. 2022), p. 045311. DOI: 10.1103/PhysRevE.105.045311.

[175] G. J. Conduit, P. H. Conlon, and B. D. Simons. "Superfluidity at the BEC-BCS crossover in two-dimensional Fermi gases with population and mass imbalance". In: *Phys. Rev. A* 77 (2008), p. 053617. DOI: 10.1103/PhysRevA.77.053617. arXiv: 0805.2726 [cond-mat.str-el].

[176] Heron Caldas, A. L. Mota, R. L. S. Farias, and L. A. Souza. "Superfluidity in Two-Dimensional Imbalanced Fermi Gases". In: *J. Stat. Mech.* 1210 (2012), P10019. DOI: 10.1088/1742-5468/2012/10/P10019. arXiv: 1108.5407 [cond-mat.other].

[177] Jia-jia Du, Cheng Chen, and Jun-jun Liang. "Asymmetric two-component Fermi gas in two dimensions". In: *Phys. Rev. A* 80 (2009), p. 023601. DOI: 10.1103/PhysRevA.80.023601.

[178] H. Biss. *Probing Excitation Spectra of Ultracold 2D and 3D Fermionic Superfluids*. Universität Hamburg, 2023.

[179] Y. Hou and J. E. Drut. "Virial expansion of attractively interacting Fermi gases in one, two, and three dimensions, up to fifth order". In: *Phys. Rev. A* 102 (3 Sept. 2020), p. 033319. DOI: 10.1103/PhysRevA.102.033319.

[180] Robert H. Swendsen and Jian-Sheng Wang. "Nonuniversal critical dynamics in Monte Carlo simulations". In: *Phys. Rev. Lett.* 58 (2 Jan. 1987), pp. 86–88. DOI: 10.1103/PhysRevLett.58.86.

[181] Danilo Jimenez Rezende and Shakir Mohamed. "Variational Inference with Normalizing Flows". In: May 2015. arXiv: 1505.05770 [stat.ML].

[182] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. "Normalizing Flows for Probabilistic Modeling and Inference". In: *J. Machine Learning Res.* 22.1 (2021), pp. 2617–2680. DOI: 10.5555/3546258.3546315. arXiv: 1912.02762 [stat.ML].

[183] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. "Density estimation using Real NVP". In: (May 2016). arXiv: 1605.08803 [cs.LG].

[184] Martin Lüscher. "Trivializing Maps, the Wilson Flow and the HMC Algorithm". In: *Communications in Mathematical Physics* 293.3 (Nov. 2009), pp. 899–919. ISSN: 1432-0916. DOI: 10.1007/s00220-009-0953-7.

[185] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. "Neural Ordinary Differential Equations". In: (June 2018). arXiv: 1806.07366 [cs.LG].

[186] Mathis Gerdes, Pim de Haan, Corrado Rainone, Roberto Bondesan, and Miranda C. N. Cheng. "Learning lattice quantum field theories with equivariant continuous flows". In: *SciPost Phys.* 15.6 (2023), p. 238. DOI: 10.21468/SciPostPhys.15.6.238. arXiv: 2207.00283 [hep-lat].

[187] Simone Bacchio, Pan Kessel, Stefan Schaefer, and Lorenz Vaitl. "Learning trivializing gradient flows for lattice gauge theories". In: *Phys. Rev. D* 107.5 (2023), p. L051504. DOI: 10.1103/PhysRevD.107.L051504. arXiv: 2212.08469 [hep-lat].

[188] Michele Caselle, Elia Cellini, and Alessandro Nada. "Sampling the lattice Nambu-Goto string using Continuous Normalizing Flows". In: *JHEP* 02 (2024), p. 048. DOI: 10.1007/JHEP02(2024)048. arXiv: 2307.01107 [hep-lat].

[189] Kenneth G. Wilson. "Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture". In: *Phys. Rev. B* 4 (9 Nov. 1971), pp. 3174–3183. DOI: 10.1103/PhysRevB.4.3174.

[190] Dorit Ron, Robert H. Swendsen, and Achi Brandt. "Inverse Monte Carlo Renormalization Group Transformations for Critical Phenomena". In: *Phys. Rev. Lett.* 89 (27 Dec. 2002), p. 275701. DOI: 10.1103/PhysRevLett.89.275701.

[191] Stavros Efthymiou, Matthew J. S. Beach, and Roger G. Melko. "Super-resolving the Ising model with convolutional neural networks". In: *Phys. Rev. B* 99 (7 Feb. 2019), p. 075113. DOI: 10.1103/PhysRevB.99.075113.

[192] Dimitrios Bachtis, Gert Aarts, Francesco Di Renzo, and Biagio Lucini. "Inverse Renormalization Group in Quantum Field Theory". In: *Phys. Rev. Lett.* 128 (8 Feb. 2022), p. 081603. DOI: 10.1103/PhysRevLett.128.081603.

[193] Jonathan Goodman and Alan D. Sokal. "Multigrid Monte Carlo Method for Lattice Field Theories". In: *Phys. Rev. Lett.* 56 (10 Mar. 1986), pp. 1015–1018. DOI: 10.1103/PhysRevLett.56.1015.

[194] Karl Jansen, Eike H. Müller, and Robert Scheichl. "Multilevel Monte Carlo algorithm for quantum mechanics on a lattice". In: *Phys. Rev. D* 102 (11 Dec. 2020), p. 114512. DOI: 10.1103/PhysRevD.102.114512.

[195] Zhihao Wang, Jian Chen, and Steven C. H. Hoi. “Deep Learning for Image Super-Resolution: A Survey”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43.10 (2021), pp. 3365–3387. DOI: 10.1109/TPAMI.2020.2982166.

[196] Stavros Efthymiou, Matthew J. S. Beach, and Roger G. Melko. “Super-resolving the Ising model with convolutional neural networks”. In: *Phys. Rev. B* 99 (7 Feb. 2019), p. 075113. DOI: 10.1103/PhysRevB.99.075113.

[197] Vudtiwat Ngampruetikorn, Jesper Levinsen, and Meera M. Parish. “Pair Correlations in the Two-Dimensional Fermi Gas”. In: *Phys. Rev. Lett.* 111 (26 Dec. 2013), p. 265301. DOI: 10.1103/PhysRevLett.111.265301.