

Studies of the sign problem in determinant quantum Monte Carlo




João Pedro da Costa Ferreira

Master in Physics

Department of Physics and Astronomy

Faculty of Sciences of the University of Porto

2026



Studies of the sign problem in determinant quantum Monte Carlo

João Pedro da Costa Ferreira

Dissertation carried out as part of the Master in Physics
Department of Physics and Astronomy
2026

Supervisor

João Manuel Viana Parente Lopes, Assistant Professor, FCUP

Co-supervisor

Vitor Manuel Pereira, Assistant Professor, FCUP

“ One requiteth a teacher badly if one remain merely a scholar. ”

Friedrich Nietzsche

À minha mãe.

Agradecimentos

Uma dissertação pode ter um só nome na capa, mas nunca é feita verdadeiramente por uma só pessoa. É a culminação de vários anos de aprendizagem, de dúvidas, de desafios e de crescimento, e seria impossível ter chegado até aqui sem todas as pessoas que, de diferentes formas, fizeram parte deste percurso. A todas elas, mesmo às que, por limitação de espaço ou de memória, não aparecerem nestas linhas, quero deixar aqui o meu agradecimento.

Ao meu orientador, Prof. João Viana Lopes, agradeço não só a orientação ao longo deste e de outros trabalhos, como também o sentido de persistência e o pensamento crítico que incutiu em mim, bem como todas as oportunidades que me proporcionou. Por tudo isto e muito mais, obrigado. Agradeço também ao meu coorientador, Prof. Vitor Pereira, por todas as questões e discussões que me ajudaram a olhar para os problemas de outra perspetiva. Ao Henrique, por toda a ajuda e disponibilidade ao longo deste ano, sem as quais muitos dos obstáculos que foram surgindo teriam sido mais difíceis de ultrapassar.

A todos aqueles com quem me cruzei no DFA ao longo destes anos, entre professores, investigadores e colegas, agradeço tudo aquilo que me ensinaram dentro e fora das salas de aula, desde física até à técnica do apagador.

Aos Discípulos e associados – Tomás, Rodrigo, Rui, Ricardo, Diogo, Henrique, Carlos e muitos outros que não cabem aqui –, por todas as tainadas e viagens; por todos os momentos em que a física ficou, felizmente, em segundo plano; e também por aqueles em que acabámos inevitavelmente a falar dela. Por tudo o que já vivemos e pelo muito que ainda espero que venha, obrigado.

Ao Pedro, à Lara e ao Zé, com quem partilhei muito mais do que uma casa durante estes anos. Entre jantares, conversas a estenderem-se por noites dentro, noites de cinema, momentos bons e outros menos bons, acabaram por se tornar uma segunda família para mim. Obrigado por me aturarem.

À minha mãe, ao meu irmão e a toda a minha família, pelo carinho e pelo apoio constantes e por celebrarem comigo cada etapa deste percurso. Obrigado por me ajudarem a tornar-me quem sou hoje.

Resumo

O método de Monte Carlo quântico do determinante (DQMC) é um método numérico poderoso, amplamente utilizado no estudo de sistemas com interações [1, 2], mas cuja aplicação a férmions interagentes é frequentemente limitada pelo problema do sinal. Quando os pesos determinantis W podem assumir valores negativos, as observáveis têm de ser obtidas por reponderação e o custo estatístico cresce rapidamente à medida que o sinal médio diminui [3]. Embora este quantifique a severidade do problema, não descreve a estrutura do espaço de configurações, a origem das flutuações do sinal ou a sua dependência do campo auxiliar.

Nesta dissertação, estudamos esta estrutura no modelo de Hubbard repulsivo. Introduzimos a variável de energia efetiva $x = -\log |W|$ e, juntamente com o sinal da configuração $s = \text{sgn}(W)$, utilizamo-la para organizar e caracterizar o espaço de configurações. Analisamos depois a ligação entre os dois setores de sinal através de inversões de um único spin do campo auxiliar e de transformações coletivas, incluindo fatias espaciais, linhas de mundo num sítio fixo, segmentos temporais contíguos e inversões coletivas aleatórias.

No regime de sinal médio reduzido, verificamos que as configurações com peso positivo e negativo aparecem em números quase iguais na região de energia efetiva mais visitada pela simulação, embora continuem pouco ligadas por movimentos que invertem um único spin. Já as transformações coletivas, sobretudo as que envolvem segmentos contíguos ao longo do eixo do tempo imaginário, permitem encontrar ligações muito mais fortes entre os dois setores de sinal. A comparação com inversões coletivas aleatórias e com a dinâmica habitual de DQMC mostra ainda que estes movimentos contíguos são mais eficazes a transitar de um setor de sinal ao outro para segmentos de tamanho intermédio, melhorando deste modo a eficiência da amostragem.

Estes resultados revelam que, subjacente ao problema do sinal em DQMC, existe um espaço de configurações com uma estrutura altamente não trivial, cuja natureza não é capturada por uma métrica simples como o sinal médio, geralmente empregado isoladamente. Embora não seja esperada uma solução genérica para o problema do sinal [4], explorar esta estrutura poderá ajudar a reduzir ineficiências adicionais de amostragem associadas a transições lentas entre os dois setores de sinal.

Palavras-chave: Monte Carlo quântico do determinante, problema do sinal, espaço de configurações do campo auxiliar, modelo de Hubbard

Abstract

Determinant quantum Monte Carlo (DQMC) is a powerful numerical method that has been extensively used to study interacting systems [1, 2], but its applicability to interacting fermions is often limited by the sign problem. Whenever the determinantal weights W are possibly non-positive, observables must be obtained through reweighting, and the statistical cost grows rapidly as the average sign decreases [3]. Although the average sign quantifies the severity of this problem, it does not describe the structure of the configuration space, how fluctuations of the sign arise, or their dependence on the auxiliary field.

In this dissertation, we study this internal structure for the repulsive Hubbard model. We introduce the effective-energy variable $x = -\log |W|$ and, together with the configuration sign $s = \text{sgn}(W)$, use it to organize and characterize the configuration space. We then probe the connectivity between the two sign sectors using single-spin and collective moves on the configuration space, including spatial slices, complete fixed-site worldlines, contiguous temporal segments and random collective moves.

We find that, in the low average-sign regime, positive- and negative-weight configurations occur with nearly equal multiplicity over the effective-energy region most strongly sampled by the simulation, while remaining only sparsely connected by single-spin moves on the auxiliary field. Collective transformations, in particular contiguous modifications along the imaginary-time axis, reveal substantially stronger connections between the two sign sectors. Comparisons with random collective moves and the standard DQMC dynamics further show that these contiguous inversions are more effective at transitioning between the two sign sectors for intermediate segment lengths, thereby improving the sampling efficiency.

These results reveal that, underlying the DQMC sign problem, lies a configuration space with a highly nontrivial structure, whose nature is not captured by a simple metric as the average sign alone usually employed. While a generic solution of the sign problem is not expected [4], exploiting this structure may help reduce additional sampling inefficiencies associated with slow transitions between the two sign sectors.

Keywords: determinant quantum Monte Carlo, sign problem, auxiliary-field configuration space, Hubbard model

Table of Contents

List of Figures.....	viii
List of Abbreviations.....	xiii
1. Introduction.....	1
1.1. Motivation.....	1
1.2. Thesis outline.....	5
2. The Hubbard model and determinant quantum Monte Carlo.....	6
2.1. The repulsive Hubbard model.....	6
2.2. Auxiliary-field formulation.....	8
2.2.1. Imaginary-time discretization.....	9
2.2.2. Hubbard-Stratonovich transformation.....	10
2.2.3. Fermionic trace and determinantal weights.....	12
2.2.4. Checkerboard decomposition.....	14
2.3. Monte Carlo sampling and local updates.....	16
2.3.1. Markov-chain Monte Carlo sampling.....	17
2.3.2. Equal-time Green functions.....	18
2.3.3. Fast local updates.....	20
2.4. The fermion sign problem.....	22
2.5. Numerical stabilization.....	25
2.5.1. Stabilized matrix products.....	25
2.5.2. Stable propagation through imaginary time.....	27
2.6. Simulation procedure and validation.....	28
3. Effective-energy and sign-resolved structure of configuration space.....	32
3.1. The effective-energy variable.....	32
3.1.1. Scaling properties of the effective energy.....	35
3.2. Sign-resolved density of states and entropy.....	42
3.3. Wang-Landau determination of the entropy.....	44
3.4. Entropy results and thermodynamic scaling.....	52
3.4.1. Dependence on inverse temperature.....	53
3.4.2. Dependence on system size.....	57
4. Local geometry under single-spin excitations.....	61
4.1. Geometric and dynamical diagnostics.....	61
4.2. Local sign-changing connectivity.....	65
4.3. Effective-energy displacements and dynamical accessibility.....	68

5. Collective excitations	73
5.1. Collective excitation families	73
5.2. Spatial and temporal collective excitations	77
5.3. Dependence on the segment length	80
5.4. Contiguous and random fixed-site excitations	83
5.4.1. Direct random fixed-site excitations	83
5.4.2. Fixed-shell reconstruction from single-spin excitations	85
5.4.3. Effect of temporal contiguity	89
5.5. Collective excitations and local-update dynamics	91
6. Conclusions and outlook	97
References	101
A. DQMC derivations	108
A.1. Trace-determinant identity	108
A.2. Equal-time Green-function representation	110
A.3. Local update determinant ratio	113
A.4. Sherman-Morrison Green function update	115
A.5. Absence of the sign problem at half filling	116
B. Determinant ratios for collective excitations	119
B.1. General collective determinant ratio	119
B.2. Complete time-slice excitations	119
B.3. Complete fixed-site worldline excitations	121
B.4. Contiguous fixed-site segment excitations	124

List of Figures

- 2.1. Comparison between the equal-time spin correlations obtained with DQMC and ED for a four-site system. The left panel shows the spin correlations as a function of the distance from a reference site, with the exact result compared with the S^z and S^x correlations measured in DQMC. The right panel shows the corresponding relative deviations from the exact values, with the error bars representing the statistical uncertainties of the DQMC estimates.30
- 2.2. Average sign $\langle s \rangle$ as a function of the average band filling $\langle n \rangle$, comparing the present DQMC implementation with the QUEST code. The two calculations show excellent agreement over the full filling range. The average sign is equal to unity close to the empty- and half-filled limits and decreases at intermediate fillings, reaching its minimum around $\langle n \rangle \simeq 0.85$. This comparison both validates the treatment of the determinant sign in the present implementation and identifies a parameter region in which the sign problem is sufficiently pronounced for the analysis developed in the following chapters.31
- 3.1. Representative Monte Carlo time series of the logarithm of the determinantal weight magnitude at $\beta = 4$ and $\beta = 8$. Blue and red segments correspond to configurations with positive and negative determinantal weights, respectively. The average signs measured during the simulations are displayed above each panel. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$33
- 3.2. Representative normalized distributions of the effective energy for configurations with positive and negative sign. The left and right panels correspond to $\beta = 4$ and $\beta = 8$, respectively. Blue and red histograms denote configurations with $W > 0$ and $W < 0$, while the dashed lines show Gaussian fits to the corresponding distributions. The simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$34
- 3.3. Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different imaginary-time steps. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. The remaining simulation parameters are $N = 16$, $\beta = 4$, $U = 6$, and $\mu = -1.5$37
- 3.4. Sign-resolved mean of the intensive effective energy as a function of the inverse imaginary-time step. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively. The vertical bars show one standard deviation of each effective-energy distribution rather than the statistical uncertainty of the estimated mean.38
- 3.5. Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different inverse temperatures. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. We note the absence of distributions of the negative weights for $\beta = 1$ and 2, meaning that no negative-signed configurations were visited during the simulation. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$39
- 3.6. Scaling of the mean effective energy with inverse temperature. The plotted quantity satisfies $\beta\langle\bar{x}\rangle = \langle x \rangle/N$. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively, while the dashed line shows a common linear fit. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$39

- 3.7. Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different lattice sizes. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$40
- 3.8. Scaling of the mean effective energy with the number of lattice sites. The plotted quantity satisfies $N\langle\bar{x}\rangle = \langle x \rangle/\beta$. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively, while the dashed line shows a linear fit. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$40
- 3.9. Variance of the intensive effective energy as a function of the inverse lattice size. Blue and red points correspond to positive- and negative-weight configurations, respectively. The dashed lines show fits of the form $\sigma_{\bar{x}}^2 \propto 1/N$ excluding the smallest lattice, for which finite-size effects are most pronounced. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$41
- 3.10. Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ obtained using the generalized sampling weight $P_{\lambda}(h) = e^{-\lambda x(h)}$. The upper and lower panels correspond to configurations with positive and negative determinantal weights, respectively. The color indicates the value of λ . Note that the two panels use different vertical scales. The remaining simulation parameters are $N = 16$, $\beta = 8$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$47
- 3.11. Evolution of the sign-resolved Wang-Landau bias during the later iterations of the algorithm. The left and right panels show $\omega_{+}(x)/(\beta N)$ and $\omega_{-}(x)/(\beta N)$, respectively, as functions of the intensive effective energy $\bar{x} = x/(\beta N)$. The color indicates the Wang-Landau iteration.51
- 3.12. Convergence of the sign-resolved Wang-Landau bias. The plotted quantity is the maximum absolute change in the intensive bias between consecutive Wang-Landau iterations, as defined in Eq. (3.53). Blue and red points correspond to the positive- and negative-sign sectors, respectively. The logarithmic vertical scale shows the progressive reduction of the corrections as the modification factor is decreased.51
- 3.13. Sign-resolved entropies as functions of the intensive effective energy $\bar{x} = x/(\beta N)$ for several inverse temperatures. The left and right panels show $S_{+}(x)/(\beta N)$ and $S_{-}(x)/(\beta N)$, respectively. For each inverse temperature, the same additive constant is subtracted from both sign sectors so that the largest entropy is zero. No negative-sign entropy is shown for $\beta = 1$, since no negative-weight configurations were observed in the interval studied. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$53
- 3.14. Difference between the positive- and negative-sign entropies, divided by the space-time volume, as a function of the intensive effective energy for several inverse temperatures. The dashed horizontal line marks equal sign-resolved entropies.54
- 3.15. Local average sign as a function of the intensive effective energy for several inverse temperatures. The curves are reconstructed from the sign-resolved entropy difference using Eq. (3.29).55
- 3.16. Sign-resolved canonical profiles reconstructed from $e^{S_{\sigma}(x)-x}$ for several inverse temperatures. Solid and dashed lines correspond to the positive- and negative-sign sectors, respectively. Each curve is divided by its own maximum, so the figure compares only the positions and shapes of the distributions.56

3.17.	Sign-resolved entropies as functions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different lattice sizes. The left and right panels show $S_+(x)/(\beta N)$ and $S_-(x)/(\beta N)$, respectively. For each lattice size, the same additive constant is subtracted from both sign sectors so that the largest entropy is zero. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$	57
3.18.	Difference between the positive- and negative-sign entropies, divided by the space-time volume, as a function of the intensive effective energy for different lattice sizes. The dashed horizontal line marks equal sign-resolved entropies.	58
3.19.	Local average sign as a function of the intensive effective energy for different lattice sizes. The curves are reconstructed from the sign-resolved entropy difference using Eq. (3.29).	59
3.20.	Peak-normalized sign-resolved canonical profiles reconstructed from $e^{S_\sigma(x)-x}$ for different lattice sizes. Solid and dashed lines correspond to the positive- and negative-sign sectors, respectively. Each curve is divided by its own maximum, so the figure compares only the positions and shapes of the profiles.	60
4.1.	Distribution $\rho(p x, \sigma)$ of the configuration-dependent sign-changing fraction $p(h)$ for single-spin excitations. The left panel shows initially positive configurations and the right panel initially negative configurations. The dashed red curve shows the local average sign $\sigma(x)$ and is included to identify the central low-sign region. The insets show fixed- x cuts of the distributions at representative values of the intensive effective energy. The parameters used are $N = 16$, $\beta = 8$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$, corresponding to $L = 128$	66
4.2.	Mean sign-changing fraction $\langle p \rangle_{x, \sigma}$ for single-spin excitations as a function of the intensive effective energy. The blue curve corresponds to positive-to-negative excitations and the red curve to negative-to-positive excitations. The dashed black curve shows the local average sign $\sigma(x)$	67
4.3.	Sign-conserving component of the displacement distribution $\rho(\Delta x, \sigma' x, \sigma)$ for single-spin excitations. The left panel shows positive-to-positive excitations and the right panel negative-to-negative excitations. The dashed red curve shows the local average sign $\sigma(x)$. The insets show fixed- x cuts at representative values of the intensive effective energy.	69
4.4.	Sign-changing component of the displacement distribution $\rho(\Delta x, \sigma' x, \sigma)$ for single-spin excitations. The left panel shows positive-to-negative excitations and the right panel negative-to-positive excitations. The dashed red curve shows the local average sign $\sigma(x)$. The insets show fixed- x cuts at representative values of the intensive effective energy.	70
4.5.	Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for single-spin excitations. The blue and red curves correspond to positive-to-negative and negative-to-positive transitions, respectively. The dashed black curve shows the local average sign $\sigma(x)$	71
4.6.	Conditional acceptance probability $A_\sigma(x) = P_{\sigma \rightarrow -\sigma}(x)/\langle p \rangle_{x, \sigma}$ for sign-changing single-spin excitations. The blue and red curves correspond to initially positive and initially negative configurations, respectively. The dashed black curve shows the local average sign $\sigma(x)$	71
5.1.	Mean sign-changing fraction $\langle p \rangle_{x, \sigma}$ for fixed-time spatial-slice and fixed-site world-line excitations. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$	78

- 5.2. Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for fixed-time spatial-slice and fixed-site worldline excitations. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$78
- 5.3. Conditional acceptance probability $A_{\sigma}(x)$ for sign-changing fixed-time spatial-slice and fixed-site worldline excitations. The left and right panels correspond to initially positive and initially negative configurations, respectively. The dashed black curve shows the local average sign $\sigma(x)$79
- 5.4. Mean sign-changing fraction $\langle p \rangle_{x,\sigma}$ for contiguous fixed-site segments of different lengths M . The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$. The cases $M = 1$ and $M = 128$ correspond to the single-spin and complete-worldline limits, respectively.80
- 5.5. Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for contiguous fixed-site segments of different lengths M . The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$81
- 5.6. Dependence of the sign-changing response on the length M of a contiguous fixed-site segment at $\bar{x}_{\text{peak}} = -3.23$. The upper panel shows the mean sign-changing fraction $\langle p \rangle_{x,\sigma}$, the middle panel the accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$, and the lower panel the conditional acceptance probability $A_{\sigma}(x)$. Blue and red curves correspond to initially positive and initially negative configurations, respectively. The horizontal dashed lines indicate the corresponding values for the fixed-time spatial-slice excitation.82
- 5.7. Mean sign-changing fraction $\langle p \rangle_{x,\sigma}$ obtained by directly sampling random fixed-site excitations containing M flipped auxiliary-field spins. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$. The case $M = 1$ corresponds to the single-spin limit.84
- 5.8. Comparison between the directly sampled and reconstructed mean sign-changing fractions for random fixed-site excitations. Solid curves show direct measurements obtained by sampling random sets of M distinct time slices at one spatial site, while colored dashed curves show the fixed-shell reconstruction obtained from the single-spin displacement distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$87
- 5.9. Comparison between the directly sampled fully random excitations and the fixed-shell reconstruction. Solid curves show direct measurements obtained by selecting M distinct auxiliary-field spins uniformly from the complete space-time lattice, while dashed curves show the reconstruction obtained from the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. Colors indicate the number M of flipped spins.89
- 5.10. Comparison between the mean sign-changing fractions of contiguous fixed-site segments and the fixed-shell random baseline. Solid curves show contiguous segments, while dashed curves show the reconstructed random fixed-site excitations obtained from the fixed-shell convolution of the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.
90

- 5.11. Comparison between the accepted sign-changing probabilities of contiguous fixed-site segments and the fixed-shell random baseline. Solid curves show contiguous segments, while dashed curves show the reconstructed random fixed-site excitations obtained from the fixed-shell convolution of the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$. 91
- 5.12. Lagged sign-transition probability $P_M(-\sigma | x, \sigma)$ measured from the ordinary single-spin DQMC Markov chain. The left and right panels correspond to positive-to-negative and negative-to-positive transitions, respectively. The colorbar indicates the number M of elementary Metropolis proposals separating the initial and final configurations. The dashed red curve shows the large- M limit of Eq. (5.37) and the dashed black curve shows the canonical distribution of Eq. (3.32).94
- 5.13. Comparison between sign changes generated by the ordinary DQMC dynamics and by contiguous collective excitations at the canonically most relevant effective energy $\bar{x}_{\text{peak}} = -3.23$. The blue curve shows the lagged probability $P_M(-\sigma | x, \sigma)$ of finding the opposite sign after M sequential single-spin Metropolis proposals, while the red curve shows the accepted sign-changing probability for a single contiguous segment containing M flipped auxiliary-field spins. Since the transition probabilities $P_{+\rightarrow-}$ and $P_{-\rightarrow+}$ were closely matched, only their average is shown. The black dashed line shows the reference linear scaling $MP_{\sigma\rightarrow-\sigma}(x; M = 1)$95

List of Abbreviations

AF	auxiliary field
DQMC	determinant quantum Monte Carlo
ED	exact diagonalization
HS	Hubbard-Stratonovich
MC	Monte Carlo
PBC	periodic boundary conditions
QMC	quantum Monte Carlo

1. Introduction

1.1. Motivation

Interacting quantum many-body systems are at the origin of many of the most interesting physical phenomena, being responsible for the emergence of phase transitions, magnetism and superconductivity [5]. In theoretical physics, a paradigmatic example of such a system is the single-band Hubbard model [6], which, despite its simplicity in the description of electron-electron interactions via a competition between kinetic delocalization and local Coulomb repulsion between electrons, exhibits phenomena such as strong correlations and antiferromagnetism [7]. For this reason, the Hubbard model has become one of the standard theoretical settings for investigating interacting fermionic systems and for benchmarking numerical methods designed to study them.

Obtaining quantitative results for interacting many-body systems is, however, challenging both analytically and numerically. Although exact analytical solutions are known for particular models or limiting cases, the same is not true in general due to the difficulty introduced by interactions. Mean-field theory provides a way to approximately compute properties of these systems; however, it relies on uncontrolled approximations and may neglect important correlation effects [8, 9]. An alternative route lies in the use of numerically exact methods or methods with controllable errors. One such method is exact diagonalization (ED), which is restricted by the exponential growth of the many-body Hilbert space with system size [10]. Quantum Monte Carlo (QMC) methods adopt a different approach by replacing the explicit treatment of the full Hilbert space with a stochastic sampling procedure, allowing substantially larger systems to be studied, with the added advantage that statistical errors can be clearly estimated and systematically improved, in principle [11, 12].

This dissertation focuses on the determinant quantum Monte Carlo (DQMC) method. It provides a formalism to reformulate the finite-temperature partition function of an interacting fermion system through the introduction of a Hubbard-Stratonovich (HS) auxiliary field, so that, for each fixed auxiliary field (AF) configuration, the fermionic degrees of freedom appear quadratically in the effective action and can be treated analytically [13]. Although the early AF formulation used continuous auxiliary fields [1, 14], later applications to lattice models with discrete auxiliary fields provided, under favorable circumstances, a reliable and efficient way of analyzing such systems [2, 15, 16].

In the implementation considered here, the auxiliary field is discrete and can be qualitatively pictured as a collection of classical Ising spins distributed over a lattice defined along both space and imaginary time dimensions. A given choice of spin orientations defines a configuration of the auxiliary field. The key technical feature of the DQMC formalism is how it rewrites the partition function and observable expectation values as a sum over these spin configurations of weights that can be evaluated in terms of determinants of matrices. This sum, however, cannot be evaluated explicitly since it contains an exponentially large number of terms. What the formalism does allow is the sampling of these many configurations through a Markov-chain Monte Carlo (MC) procedure. It is from these sampled configurations that the necessary Green functions are computed, which are then used to evaluate physical observables with controllable statistical uncertainty [17].

Whereas the weights in classical Monte Carlo and many bosonic QMC formulations are always nonnegative [18, 19], the weights that arise in DQMC are not constrained to be non-negative for fermionic systems, except in select models or fine-tuned parameters [20]. This non-positivity prevents the determinantal weights from being interpreted directly as a probability. The usual approach is then to sample configurations according to the absolute value of their weight and include the corresponding sign through reweighting. As cancellations between positive and negative weights become severe, the average sign decreases and the statistical uncertainty of measured observables increases rapidly; consequently, the sampling effort also increases substantially in order to maintain a given statistical accuracy. The existence of non-positive “statistical weights” and the challenges this poses to feasible statistical sampling is that is generally known as “the sign problem” in the context of QMC applied to fermionic systems [3]. In practice, the average sign often decreases exponentially with system size and inverse temperature, but it also depends strongly on model-dependent parameters such as the band filling and lattice geometry [21]. The sign problem can therefore severely restrict the range of temperatures, system sizes and physical regimes that are accessible to otherwise numerically exact DQMC simulations. For the Hubbard model, a particular example is the hole-doped low-temperature regime, which is the experimentally relevant range in its application as a candidate model for cuprate superconductors. The reduced average sign in this regime has severely limited the numerical results accessible to QMC methods in this region of the phase diagram [22, 23].

The presence and severity of the sign problem can depend on both the single-particle basis used to represent the Hamiltonian and the particular choice of the auxiliary field used to decouple the interaction. For instance, fermionic creation and annihilation operators can be rewritten in terms of Majorana operators and, for suitable classes of models, the HS decoupling can be chosen such that the two Majorana sectors contribute with complex-conjugate weights, yielding a non-negative total weight [24, 25]. Furthermore, in some cases, symmetries can guarantee that the determinantal weights remain non-negative, leading to formulations free of the sign problem [20]. Even when the sign problem cannot be eliminated completely, its severity can be reduced by optimizing the single-particle basis used [26]. In other QMC methods, there have been attempts at incorporating sign cancellations more explicitly into the sampling procedure. A notable example is the meron-cluster approach, where non-local clusters whose inversion changes the fermionic sign can be identified and used to cancel opposite-sign contributions exactly for certain models, and simulations can be restricted to configurations with non-negative weight [27]. Crucially, these methods are necessarily model- or representation-dependent. In fact, a general, model-independent polynomial-time solution of the sign problem is not to be expected as it is a NP-hard computational problem, and, as such, practical progress has relied on approaches tailored for specific models [4].

Although the sign problem is usually characterized through the decay of the average sign, several studies have shown that it is related to nontrivial structure at the level of the individual auxiliary-field configurations. In DQMC, the sign of the determinant can be understood as a geometric property of the full imaginary-time single-particle evolution, with negative-sign configurations being associated with a nontrivial geometric phase [28]. At the statistical level, the occurrence of negative-signed configurations is also strongly affected by the regions of configuration space selected by absolute-weight sampling, showing that the severity of the sign problem cannot be understood solely from the properties of arbitrary auxiliary-field configurations [29]. Studies focusing on the auxiliary-field configurations have also revealed nontrivial organization within configuration space, for instance through Hamming-distance measures that probe both their distribution and the efficiency with which they are explored dynamically [30]. Machine-learning techniques have likewise been applied directly to auxiliary-field configurations, showing that these configurations contain information about physical phases, while also becoming substantially harder to interpret in regimes affected by the sign problem [31].

Further evidence for a relation between configuration space and the sign has followed

from approaches that condition statistical information on the sign or analyze quantities derived directly from the fermionic matrices. For instance, by obtaining the sign-resolved distributions of physical observables it has been noted that positive- and negative-weight configurations form statistically distinct populations, and that small differences survive even as the average sign approaches zero [32]. At the level of the fermionic matrices, the singular values of these matrices contain information about the organization of the sampled auxiliary fields and changes in the underlying phase transition, even when the fluctuating determinant phase is discarded [33]. The determinant magnitude has also been used directly as a coordinate for flat-histogram sampling through the Wang-Landau algorithm, allowing a reduction in autocorrelation barriers by increasing the range explored in configuration-weight space [34]. Taken together, these results suggest that the sign problem is profoundly related to a structured configuration space whose geometry, statistical weight, and fermionic matrix properties are mutually related, while leaving open a more direct characterization of how these positive- and negative-sign configurations are distributed and connected within that space.

The way in which auxiliary-field configuration space is explored can itself introduce additional difficulties. In hybrid Monte Carlo formulations with continuous auxiliary fields, zeros of the fermionic determinant can form barriers that separate the two sign sectors of configuration space and obstruct ergodic sampling, which amounts to an additional dynamical difficulty introduced by the sign problem [35]. In DQMC, updates involving the inversion of an entire fixed-site worldline have also been shown to help overcome ergodicity problems at strong coupling [36]. These results indicate that the connectivity between elements of the configuration space also depends on the transformations used to explore it, motivating a more direct study of how configurations belonging to opposite sign sectors are connected by local and collective auxiliary-field moves.

Motivated by these observations, this dissertation investigates the relation between the determinant sign and the structure of the auxiliary-field configuration space in the repulsive Hubbard model. We begin by organizing the configurations according to the magnitude of their determinantal weight and their sign, which allows us to characterize the populations of the two sign sectors as a function of the magnitude of the weights. We also probe how well the two sign sectors are connected through different transformations of the auxiliary field, beginning with the elementary single-spin flip and extending to collective moves. Particular attention is given to the comparison between collective moves and the sign-changing behavior of the standard local-update DQMC dynamics. In this way, the aim is

to distinguish the balance between the number of positive- and negative-weight configurations from the connectivity between them under a given class of auxiliary-field transformations, together with how easily can the Markov-chain dynamics move between these two populations. The latter question is relevant whenever the sign itself has a long autocorrelation time, since inefficient transition rates among sectors of configuration space with opposite sign introduces an additional sampling cost beyond that arising from the sign problem itself [35].

1.2. Thesis outline

The remainder of this dissertation is organized as follows. Chapter 2 introduces the Hubbard model and the DQMC formalism using a discrete HS decoupling in the spin channel, the evaluation of the fermionic determinants and Green functions, the MC sampling procedure, the origin of the sign problem, and an overview of the numerical stabilization techniques used.

Chapter 3 defines the effective-energy variable used throughout this work, studies the statistical organization of positive- and negative-weight configurations through sign-resolved distributions and entropies, and analyzes the thermodynamic scaling of the latter.

Chapter 4 investigates a local structure of the configuration space by considering single-spin updates of the auxiliary field. The probability of these updates reaching the opposite sign sector and the corresponding Metropolis acceptance probabilities are characterized as functions of the effective energy, which provide a direct measure of the local connectivity between the sign sectors.

Chapter 5 extends this analysis to collective moves involving multiple auxiliary-field spins. A comparison is made between different spatial and imaginary-time structures, including fixed-time-slice, fixed-site, contiguous and random updates. The resulting sign-changing probabilities are also compared with the behavior generated by the standard local-update DQMC dynamics, allowing the effects of collective transformations and sequential standard sampling to be distinguished.

Finally, Chapter 6 summarizes the main results of this dissertation and discusses their implications for the understanding of the configuration-space structure underlying the sign problem, together with possible directions for future work.

2. The Hubbard model and determinant quantum Monte Carlo

In this chapter, we introduce the repulsive Hubbard model and formulate the determinant quantum Monte Carlo method used throughout this thesis. We begin by defining the Hamiltonian and the conventions adopted here. We then discretize the imaginary-time direction and apply a HS transformation, reducing the interacting problem to a sum over configurations of an auxiliary field. After tracing out the fermionic degrees of freedom, we obtain the determinantal configuration weights and introduce the checkerboard decomposition used to apply the kinetic propagator efficiently. We next describe the Markov-chain Monte Carlo procedure used to sample the auxiliary-field configurations, together with the equal-time Green functions, local determinant ratios, and rank-one updates required for an efficient implementation. Having established the sampling algorithm, we discuss the fermion sign problem and the reweighting procedure required when the determinantal weights can become negative. Finally, we address the numerical instabilities caused by long products of imaginary-time propagators, summarize the stabilized simulation procedure, and describe the statistical analysis and validation tests used to verify the implementation.

2.1. The repulsive Hubbard model

The Hubbard model was introduced as a minimal description of electrons occupying a narrow energy band [6]. Starting from electrons moving in a periodic potential, the electronic states are expanded in a basis of localized Wannier orbitals, and the description is restricted to one relevant orbital per lattice site. Retaining the dominant hopping matrix elements between different sites and the on-site Coulomb interaction then reduces the many-electron problem to a competition between kinetic delocalization and local repulsion. Despite this simple formulation, the model displays a broad range of correlated phenomena, including metallic behavior, antiferromagnetic correlations, and Mott insulating phases [5, 7, 22]. Its combination of a compact formulation and nontrivial behavior has made it a standard benchmark for numerical methods for interacting fermions [11].

We denote by $\hat{c}_{i\sigma}^\dagger$ and $\hat{c}_{i\sigma}$ the creation and annihilation operators for a fermion at site i with spin $\sigma \in \{\uparrow, \downarrow\}$, and define the local occupation operator by

$$\hat{n}_{i\sigma} = \hat{c}_{i\sigma}^\dagger \hat{c}_{i\sigma}. \quad (2.1)$$

In the grand-canonical ensemble, the single-band Hubbard Hamiltonian is

$$\hat{H} = - \sum_{i,j,\sigma} t_{ij} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} - \mu_{\text{phys}} \sum_{i,\sigma} \hat{n}_{i\sigma}. \quad (2.2)$$

The hopping amplitudes t_{ij} , which are often restricted to a few neighboring sites, determine the kinetic motion of the electrons through the lattice. The parameter $U > 0$ is the on-site repulsion, which penalizes double occupancy, while μ_{phys} is the chemical potential controlling the band filling. The competition between kinetic delocalization and local repulsion is primarily controlled by the dimensionless ratio U/t , where t denotes a characteristic hopping amplitude.

In the weak-coupling regime, $U/t \ll 1$, the electrons are more itinerant and the kinetic delocalization dominates. In the strong-coupling regime, $U/t \gg 1$, double occupancy is strongly suppressed. Near half filling, this suppresses charge fluctuations and leaves approximately one electron per site, producing local magnetic moments and Mott insulating behavior. Although real hopping is inhibited, it generates an antiferromagnetic exchange interaction leading to antiferromagnetic correlations between neighboring moments [7]. Between the weak- and strong-coupling limits, neither the kinetic nor the interaction term can generally be treated perturbatively. The precise behavior also depends on the band filling, lattice geometry, dimensionality, and temperature. Throughout this work, we consider a square lattice with nearest-neighbor hopping, periodic boundary conditions (PBC), and $N = N_x N_y$ lattice sites. The hopping amplitude is taken as the unit of energy, $t = 1$, and the Boltzmann constant is set to unity, so that the temperature is specified through the inverse temperature $\beta = 1/T$.

Since many numerical studies are performed at or near the half-filled case $\langle \hat{n}_i \rangle = 1$ [15, 16, 37], it is convenient to write the interaction in a particle-hole-symmetric form. Using

$$U \left(\hat{n}_{i\uparrow} - \frac{1}{2} \right) \left(\hat{n}_{i\downarrow} - \frac{1}{2} \right) = U \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} - \frac{U}{2} (\hat{n}_{i\uparrow} + \hat{n}_{i\downarrow}) + \frac{U}{4}, \quad (2.3)$$

we rewrite the Hamiltonian, up to an irrelevant additive constant, as

$$\hat{H} = \hat{T} + \hat{V}, \quad (2.4)$$

where

$$\hat{T} = - \sum_{i,j,\sigma} t_{ij} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} - \mu \sum_{i,\sigma} \hat{n}_{i\sigma}, \quad (2.5)$$

$$\hat{V} = U \sum_i \left(\hat{n}_{i\uparrow} - \frac{1}{2} \right) \left(\hat{n}_{i\downarrow} - \frac{1}{2} \right), \quad (2.6)$$

and

$$\mu = \mu_{\text{phys}} - \frac{U}{2}. \quad (2.7)$$

With this convention, $\mu = 0$ is the particle-hole-symmetric point and corresponds to half filling. We also introduce the single-particle matrix K associated with the quadratic part of the Hamiltonian,

$$\hat{T} = \sum_{i,j,\sigma} \hat{c}_{i\sigma}^\dagger K_{ij} \hat{c}_{j\sigma}, \quad K_{ij} = -t_{ij} - \mu \delta_{ij}. \quad (2.8)$$

The hopping is spin independent, so the same matrix K is used for both spin sectors. The separation in Eq. (2.4) between a quadratic term and a local quartic interaction is the starting point of the determinant quantum Monte Carlo formulation.

2.2. Auxiliary-field formulation

A central difficulty in applying Monte Carlo methods to an interacting quantum system is that the thermal partition function involves the trace of an interacting many-body propagator acting on a Fock space. Although the trace may formally be written as a sum over many-body basis states, its matrix elements do not directly provide a set of simple, computable statistical weights suitable for Monte Carlo sampling. Furthermore, since the dimension of the fermionic Fock space grows exponentially with the number of lattice sites, explicitly constructing and diagonalizing the full Hamiltonian becomes impractical except for small systems [17]. Determinant quantum Monte Carlo avoids this explicit many-body treatment by transforming the interacting problem into a sum over configurations of an auxiliary field [1, 14, 38]. For each such configuration of this field, the fermionic problem becomes quadratic, which allows the fermionic degrees of freedom to be traced out analytically. Their contribution is then represented by determinants of matrices acting on the single-particle space. These determinants provide the weights used to sample the auxiliary-field configurations through Markov-chain Monte Carlo. Although the auxiliary-field configuration space remains exponentially large, it need not be enumerated explicitly,

and observables can instead be estimated from the sampled configurations with a controllable statistical error. In this section, we construct this representation by successively discretizing imaginary time, decoupling the interaction through a HS transformation, and expressing the resulting fermionic trace as a determinant.

2.2.1 Imaginary-time discretization

At finite temperature, the thermodynamic properties of the system are determined by the grand-canonical partition function

$$Z = \text{Tr} \left[e^{-\beta \hat{H}} \right], \quad (2.9)$$

where $\beta = 1/T$ and the trace is taken over the complete fermionic Fock space. Although Eq. (2.9) is formally compact, it is not yet suitable for Monte Carlo sampling. In particular, the Hamiltonian contains both the quadratic hopping term $\hat{T}$ and the quartic interaction $\hat{V}$, and these two operators do not generally commute,

$$[\hat{T}, \hat{V}] \neq 0. \quad (2.10)$$

Consequently, the operator $e^{-\beta \hat{H}}$ cannot be factorized exactly as $e^{-\beta \hat{H}} = e^{-\beta \hat{T}} e^{-\beta \hat{V}}$. This difficulty is dealt with by discretizing the imaginary-time interval $[0, \beta]$ into L time slices of width

$$\Delta\tau = \frac{\beta}{L}. \quad (2.11)$$

The partition function may then be written exactly as

$$Z = \text{Tr} \left[\prod_{\ell=1}^L e^{-\Delta\tau \hat{H}} \right], \quad (2.12)$$

where, throughout this work, products over imaginary-time slices are ordered such that

$$\prod_{\ell=1}^L A_{\ell} \equiv A_L A_{L-1} \cdots A_1. \quad (2.13)$$

For sufficiently small $\Delta\tau$, the Trotter-Suzuki decomposition allows us to factorize the kinetic term from the interaction term, thus approximating each short-time propagator as [39, 40]

$$e^{-\Delta\tau(\hat{T}+\hat{V})} = e^{-\Delta\tau\hat{V}} e^{-\Delta\tau\hat{T}} + \mathcal{O}(\Delta\tau^2). \quad (2.14)$$

The error displayed in Eq. (2.14) is the error associated with one short-time factor. Since the full imaginary-time evolution contains $L = \beta/\Delta\tau$ such factors, the local errors accumulate into a systematic discretization error (the *Trotter error*) in the simulated observables. In the limit $\Delta\tau \rightarrow 0$, the decomposition becomes exact. The ordering adopted in Eq. (2.14) also fixes the convention for the single-particle propagators introduced below.

Applying the decomposition to every imaginary-time slice gives

$$Z \simeq \text{Tr} \left[\prod_{\ell=1}^L e^{-\Delta\tau \hat{V}} e^{-\Delta\tau \hat{T}} \right]. \quad (2.15)$$

The kinetic exponential is generated by an operator quadratic in the fermionic creation and annihilation operators. It can therefore be represented by the exponential of the $N \times N$ single-particle matrix K defined in Eq. (2.8). The interaction exponential, on the other hand, remains quartic in the fermionic operators. The Trotter decomposition has separated the two contributions, but an additional transformation is required before the interaction can be treated analytically.

2.2.2 Hubbard-Stratonovich transformation

The purpose of the HS transformation is to replace the direct interaction between the fermions by a coupling between the fermions and an auxiliary field. The interaction is not neglected or treated perturbatively. Instead, the exponential of the quartic interaction is written exactly, for each discrete imaginary-time slice, as a sum of exponentials which are only quadratic in the fermionic operators.

For the repulsive Hubbard model in the particle-hole-symmetric form, the interaction at a single site is

$$\hat{V}_i = U \left(\hat{n}_{i\uparrow} - \frac{1}{2} \right) \left(\hat{n}_{i\downarrow} - \frac{1}{2} \right). \quad (2.16)$$

Using the fermionic identity $\hat{n}_{i\sigma}^2 = \hat{n}_{i\sigma}$, the product of number operators can be expressed in terms of the square of the local magnetization,

$$\left(\hat{n}_{i\uparrow} - \frac{1}{2} \right) \left(\hat{n}_{i\downarrow} - \frac{1}{2} \right) = -\frac{1}{2} (\hat{n}_{i\uparrow} - \hat{n}_{i\downarrow})^2 + \frac{1}{4}. \quad (2.17)$$

Equation (2.17) suggests a decoupling in the spin channel, in which an auxiliary variable couples linearly, and with opposite signs, to the two spin species.

For the discrete imaginary-time formulation, we use Hirsch's discrete HS transformation [2],

$$e^{-\Delta\tau U \left(\hat{n}_{i\uparrow} - \frac{1}{2}\right) \left(\hat{n}_{i\downarrow} - \frac{1}{2}\right)} = C \sum_{h_i = \pm 1} e^{\nu h_i (\hat{n}_{i\uparrow} - \hat{n}_{i\downarrow})}, \quad (2.18)$$

where

$$C = \frac{1}{2} e^{-\Delta\tau U/4}, \quad \cosh \nu = e^{\Delta\tau U/2}. \quad (2.19)$$

The auxiliary variable h_i takes the two values $+1$ and -1 and is therefore Ising-like. The transformation in Eq. (2.18) can be verified by considering the four possible local occupation states,

$$|0\rangle, \quad |\uparrow\rangle, \quad |\downarrow\rangle, \quad |\uparrow\downarrow\rangle. \quad (2.20)$$

By acting on them with both sides of Eq. (2.18), one obtains the system of equations

$$\begin{cases} 2C = e^{-\Delta\tau U/4}, \\ 2C \cosh \nu = e^{\Delta\tau U/4}, \end{cases} \quad (2.21)$$

of which Eq. (2.19) is the solution. Thus, with that choice of parameters, both sides of Eq. (2.18) agree on a complete basis of the local Hilbert space.

The choice in Eq. (2.18) couples the auxiliary field to the local spin density. However, it should be noted that other HS decompositions are possible, including transformations which preserve different symmetries or couple to different fermion bilinears [41, 42]. Although these decompositions represent the same interaction after the auxiliary field is summed over, they may lead to different numerical properties. Throughout this thesis, we use the discrete spin-channel transformation in Eq. (2.18).

The interaction operator is a sum of mutually commuting on-site terms,

$$\hat{V} = \sum_i \hat{V}_i, \quad (2.22)$$

so Eq. (2.18) may be applied independently at every lattice site. For a single imaginary-time slice, this gives

$$e^{-\Delta\tau \hat{V}} = C^N \sum_{\{h_i\}} \exp \left[\nu \sum_i h_i (\hat{n}_{i\uparrow} - \hat{n}_{i\downarrow}) \right], \quad (2.23)$$

where $\{h_i\}$ denotes one of the 2^N possible auxiliary-field configurations on that slice.

A separate auxiliary variable must be introduced at every site and at every imaginary-time slice. We therefore denote the complete auxiliary-field configuration by

$$h \equiv \{h_{\ell,i}\}, \quad h_{\ell,i} \in \{-1, +1\}, \quad (2.24)$$

where $i = 1, \dots, N$ labels the lattice sites and $\ell = 1, \dots, L$ labels the imaginary-time slices. The auxiliary field may thus be viewed as a classical Ising field defined on an $N \times L$ space-time lattice. The complete field space contains 2^{NL} different configurations.

Substituting Eq. (2.23) into the partition function in Eq. (2.15), we obtain

$$Z \simeq C^{NL} \sum_h \text{Tr} \left[\prod_{\ell=1}^L \exp \left[\nu \sum_i h_{\ell,i} (\hat{n}_{i\uparrow} - \hat{n}_{i\downarrow}) \right] e^{-\Delta\tau \hat{T}} \right]. \quad (2.25)$$

The quartic fermion interaction has now disappeared from the operator appearing inside the trace. For any fixed configuration h , all remaining fermionic operators are quadratic. The price paid for this simplification is the introduction of the exponentially large sum over the auxiliary-field configurations. Unlike the original Fock-space trace, however, every term in this sum can now be evaluated using operations in the single-particle Hilbert space.

2.2.3 Fermionic trace and determinantal weights

Because the kinetic Hamiltonian is spin independent and does not mix the two spin species, it can be decomposed as

$$\hat{T} = \hat{T}_{\uparrow} + \hat{T}_{\downarrow}, \quad \hat{T}_{\sigma} = \sum_{i,j} \hat{c}_{i\sigma}^{\dagger} K_{ij} \hat{c}_{j\sigma}. \quad (2.26)$$

Similarly, the auxiliary-field term in Eq. (2.25) is diagonal in spin. Introducing the numerical convention

$$\sigma = \begin{cases} +1, & \sigma = \uparrow, \\ -1, & \sigma = \downarrow, \end{cases} \quad (2.27)$$

the short-time many-body propagator for spin sector σ and time slice ℓ is defined as

$$\hat{B}_{\ell,\sigma} = e^{\sigma\nu \sum_i h_{\ell,i} \hat{n}_{i\sigma}} e^{-\Delta\tau \hat{T}_{\sigma}}. \quad (2.28)$$

This operator propagates the spin- σ fermions across one imaginary-time interval in the background of the auxiliary field on slice ℓ . Since the two spin sectors act on independent

components of the Fock space, the trace factorizes,

$$Z \simeq C^{NL} \sum_h \prod_{\sigma=\pm 1} \text{Tr} \left[\prod_{\ell=1}^L \hat{B}_{\ell,\sigma} \right]. \quad (2.29)$$

The operators $\hat{B}_{\ell,\sigma}$ still act on the many-body Fock space. Their usefulness follows from the fact that they are exponentials of quadratic fermionic operators. Each of them therefore has a corresponding representation in the N -dimensional single-particle space. We denote this matrix by $B_{\ell,\sigma}$, without a hat. With the Trotter ordering adopted in Eq. (2.14), it is given by

$$B_{\ell,\sigma} = e^{V_{\ell,\sigma}} e^{-\Delta\tau K}, \quad (2.30)$$

where the auxiliary-field matrix $V_{\ell,\sigma}$ is diagonal,

$$[V_{\ell,\sigma}]_{ij} = \sigma \nu h_{\ell,i} \delta_{ij}. \quad (2.31)$$

The matrix $e^{-\Delta\tau K}$ describes the kinetic propagation during one time interval, while $e^{V_{\ell,\sigma}}$ contains the coupling to the auxiliary field on slice ℓ . Although the auxiliary field is common to the two spin sectors, its contribution enters with opposite signs for spin up and spin down.

The remaining trace over the fermionic Fock space can be evaluated using the trace-determinant identity [13, 43]. For any collection of quadratic fermionic operators,

$$\hat{A}_\ell = \sum_{i,j} \hat{c}_i^\dagger [A_\ell]_{ij} \hat{c}_j, \quad (2.32)$$

one has

$$\text{Tr} \left[\prod_{\ell=1}^L e^{\hat{A}_\ell} \right] = \det \left(I + \prod_{\ell=1}^L e^{A_\ell} \right). \quad (2.33)$$

The trace on the left-hand side is taken over the fermionic Fock space, whereas the determinant on the right-hand side is taken only over the single-particle space. A derivation of Eq. (2.33) is presented in Appendix A.1.

Applying this identity independently to the two spin sectors in Eq. (2.29), we arrive at a representation of the partition function in terms of determinants of one-particle operators:

$$Z \simeq C^{NL} \sum_h \det M_\uparrow(h) \det M_\downarrow(h), \quad (2.34)$$

where

$$M_\sigma(h) = I + \prod_{\ell=1}^L B_{\ell,\sigma}. \quad (2.35)$$

We may therefore assign to each auxiliary-field configuration h the weight

$$W(h) = C^{NL} \det M_\uparrow(h) \det M_\downarrow(h). \quad (2.36)$$

The factor C^{NL} is positive and independent of the auxiliary-field configuration. It therefore cancels from normalized expectation values and from the Metropolis ratios used to evolve the Markov chain. We thus absorb this factor into the overall normalization and define the DQMC configuration weight as

$$W(h) \equiv \det M_\uparrow(h) \det M_\downarrow(h). \quad (2.37)$$

Equation (2.34) is the central reduction underlying DQMC. The original partition function involves a trace over a many-body Fock space whose dimension grows exponentially with N . After introducing the auxiliary field, the fermionic degrees of freedom are integrated out exactly for every fixed field configuration, and their complete contribution is encoded in the determinants of two $N \times N$ matrices.

2.2.4 Checkerboard decomposition

The single-particle propagator in Eq. (2.30) contains the kinetic factor $e^{-\Delta\tau K}$. Although this matrix can be computed directly, its multiplication by another dense $N \times N$ matrix requires $\mathcal{O}(N^3)$ operations. Since the hopping term is local, this cost can be reduced by decomposing the kinetic matrix into groups of non-overlapping bonds, whose exponentials can be applied as a sequence of sparse two-site operations. This procedure is commonly known as the checkerboard decomposition [13, 17].

We first separate the hopping and chemical-potential contributions,

$$K = K_{\text{hop}} - \mu I. \quad (2.38)$$

The chemical-potential contribution is proportional to the identity and therefore commutes with all remaining factors, so that

$$e^{-\Delta\tau K} = e^{\Delta\tau\mu} e^{-\Delta\tau K_{\text{hop}}}. \quad (2.39)$$

The hopping matrix is then divided into a set of sparse matrices,

$$K_{\text{hop}} = \sum_{p=1}^{N_{\text{cb}}} K_p, \quad (2.40)$$

where each K_p contains a collection $\mathcal{B}_p$ of non-overlapping bonds,

$$K_p = - \sum_{\langle i,j \rangle \in \mathcal{B}_p} t_{ij} (|i\rangle\langle j| + |j\rangle\langle i|). \quad (2.41)$$

No lattice site belongs to more than one bond within the same set $\mathcal{B}_p$. Consequently, all bond terms contained in a given K_p commute.

For the nearest-neighbor square lattice considered here, the bonds can be separated into four groups: even and odd bonds along the x direction, and even and odd bonds along the y direction. Thus, $N_{\text{cb}} = 4$, provided that the linear lattice dimensions are compatible with this partition under PBC, which in this case corresponds to even linear lattice dimensions.

Applying an additional Trotter decomposition to Eq. (2.40), the kinetic propagator is approximated by

$$e^{-\Delta\tau K_{\text{hop}}} \simeq \prod_{p=1}^{N_{\text{cb}}} e^{-\Delta\tau K_p} + \mathcal{O}(\Delta\tau^2). \quad (2.42)$$

Since the bonds within each group are mutually disjoint, their exponential factorizes exactly,

$$e^{-\Delta\tau K_p} = \prod_{\langle i,j \rangle \in \mathcal{B}_p} e^{-\Delta\tau K^{(ij)}}, \quad (2.43)$$

where $K^{(ij)}$ contains only the hopping between sites i and j .

In the two-dimensional single-particle subspace spanned by $|i\rangle$ and $|j\rangle$, the bond matrix and its exponential are

$$K^{(ij)} = \begin{pmatrix} 0 & -t_{ij} \\ -t_{ij} & 0 \end{pmatrix}, \quad (2.44)$$

$$e^{-\Delta\tau K^{(ij)}} = \begin{pmatrix} \cosh(\Delta\tau t_{ij}) & \sinh(\Delta\tau t_{ij}) \\ \sinh(\Delta\tau t_{ij}) & \cosh(\Delta\tau t_{ij}) \end{pmatrix}. \quad (2.45)$$

All remaining diagonal elements of the full $N \times N$ bond propagator are equal to unity. Each factor therefore differs from the identity only in a 2×2 block.

The checkerboard decomposition has the advantage of reducing the complexity of the

operation of exponentiating the matrix K and applying the matrix $e^{-\Delta\tau K}$ from $\mathcal{O}(N^3)$ to $\mathcal{O}(N_b N^2)$, where N_b is the number of sites connected to a given site. In most models, $N_b \ll N$ so the checkerboard decomposition greatly reduces the computational cost to $\mathcal{O}(N^2)$ rather than the $\mathcal{O}(N^3)$ cost of dense matrix multiplication.

The checkerboard breakup introduces an additional discretization error, which vanishes as $\Delta\tau \rightarrow 0$, because the matrices K_p associated with different bond groups do not generally commute. Combining Eqs. (2.39) and (2.42), the kinetic factor used in the implementation is

$$e^{-\Delta\tau K} \simeq e^{\Delta\tau\mu} \prod_{p=1}^{N_{\text{cb}}} e^{-\Delta\tau K_p}. \quad (2.46)$$

2.3. Monte Carlo sampling and local updates

The auxiliary-field formulation derived in the previous section expresses the partition function as a sum over configurations of a classical Ising-like field,

$$Z \simeq \sum_h W(h), \quad (2.47)$$

where the weight $W(h)$ can be evaluated from determinants of matrices acting on the single-particle space. This expression resembles a classical partition function, in which h plays the role of a classical configuration and $W(h)$ plays the role of its statistical weight, and so it makes it possible to apply the methods of classical statistical mechanics to the interacting fermion problem. The number of auxiliary-field configurations is nevertheless 2^{NL} (much larger than the original 4^N -dimensional Hilbert space for $L > 2$), so an explicit evaluation of the sum remains impossible for the system sizes of interest. Instead, DQMC uses a Markov chain to sample the configurations that contribute most strongly to the partition function.

We first describe this procedure under the assumption that $W(h)$ is non-negative for every auxiliary-field configuration. In that case, the normalized probability distribution is

$$P(h) = \frac{W(h)}{Z}. \quad (2.48)$$

The modification required when the determinantal weights fluctuate in sign is discussed in Sec. 2.4.

2.3.1 Markov-chain Monte Carlo sampling

Rather than drawing statistically independent configurations directly from Eq. (2.48), which is generally not possible, we construct a Markov chain

$$h^{(0)} \longrightarrow h^{(1)} \longrightarrow h^{(2)} \longrightarrow \dots \quad (2.49)$$

whose asymptotic distribution is $P(h)$. A transition from a configuration h to another configuration h' is described by a transition probability

$$T(h \rightarrow h') = q(h \rightarrow h')A(h \rightarrow h'), \quad (2.50)$$

where $q(h \rightarrow h')$ is the probability of proposing the new configuration and $A(h \rightarrow h')$ is its acceptance probability. To ensure that the desired distribution is stationary, it is sufficient to choose transitions that obey detailed balance [17],

$$P(h)T(h \rightarrow h') = P(h')T(h' \rightarrow h). \quad (2.51)$$

In the Metropolis-Hastings algorithm, the acceptance probability is chosen as [44, 45]

$$A(h \rightarrow h') = \min \left\{ 1, \frac{P(h')q(h' \rightarrow h)}{P(h)q(h \rightarrow h')} \right\}. \quad (2.52)$$

The elementary updates used here consist of flipping one auxiliary-field variable,

$$h_{\ell,i} \longmapsto h'_{\ell,i} = -h_{\ell,i}, \quad (2.53)$$

while leaving all other variables unchanged. The same local flip transforms h' back into h , so the proposal probabilities are symmetric,

$$q(h \rightarrow h') = q(h' \rightarrow h). \quad (2.54)$$

Equation (2.52) therefore reduces to

$$A(h \rightarrow h') = \min \left\{ 1, \frac{W(h')}{W(h)} \right\}. \quad (2.55)$$

A proposed flip is accepted by drawing a uniformly distributed random number $r \in [0, 1]$ and comparing it with the probability in Eq. (2.55). If $r \leq A(h \rightarrow h')$, the Markov chain moves to h' . Otherwise, the original configuration is retained as the next element of the chain. Rejected proposals must therefore also be counted when constructing MC averages.

After an initial thermalization period, the distribution of configurations generated by the

chain approaches Eq. (2.48). An observable with auxiliary-field estimator $O(h)$ can then be evaluated as

$$\langle O \rangle = \sum_h P(h) O(h) \simeq \frac{1}{M} \sum_{m=1}^M O(h_m). \quad (2.56)$$

However, one needs to take into account the fact that successive configurations of a Markov chain are correlated, so the number M of measurements is generally larger than the effective number of statistically independent samples. The associated autocorrelations and the statistical procedure used in this work are discussed in Sec. 2.6.

Although Eq. (2.55) defines the sampling procedure, computing the ratio $W(h')/W(h)$ directly would be expensive. Each local proposal changes the matrices entering the determinant weight, and a naive implementation would require rebuilding the complete imaginary-time product and evaluating new determinants. The quadratic form of the fermionic problem allows these ratios to be obtained much more efficiently from equal-time Green functions.

2.3.2 Equal-time Green functions

Besides being used in the computation of the ratio $W(h')/W(h)$, as will be shown in Sec. 2.3.3, the equal-time Green functions are also useful in computing observables. In practice, observables are computed as averages over the sampled configurations,

$$\langle \mathcal{O} \rangle = \frac{1}{N_{\text{samples}}} \sum_{i=1}^{N_{\text{samples}}} \langle \mathcal{O} \rangle_{h_i}, \quad (2.57)$$

where the subscript h_i means that the expectation value is computed for a fixed configuration h_i . Now, for every fixed auxiliary-field configuration h , the fermionic problem is quadratic and its observables can be reduced to one-particle Green functions through Wick's theorem [5]. As such, we define the equal-time Green function at imaginary-time slice ℓ by

$$G_{\sigma,ij}(\ell) = \left\langle \hat{c}_{i\sigma}(\ell\Delta\tau) \hat{c}_{j\sigma}^\dagger(\ell\Delta\tau) \right\rangle_h, \quad (2.58)$$

where $\hat{A}(\tau) = e^{H\tau} \hat{A} e^{-H\tau}$ is an imaginary-time Heisenberg operator. This expectation value may be evaluated in terms of the short-time propagators $B_{\ell,\sigma}$ as [13, 43]

$$G_\sigma(\ell) = \left[I + B_{\ell,\sigma} \cdots B_{1,\sigma} B_{L,\sigma} \cdots B_{\ell+1,\sigma} \right]^{-1}, \quad (2.59)$$

a proof of which is shown in Appendix A.2. The product in Eq. (2.59) is obtained by cutting the periodic imaginary-time sequence immediately after slice ℓ . For example, at $\ell = L$ one recovers

$$G_\sigma(L) = [I + B_{L,\sigma} \cdots B_{1,\sigma}]^{-1} = M_\sigma^{-1}. \quad (2.60)$$

Moving the cut cyclically changes the matrix representation of the equal-time Green function, while leaving the determinant weight invariant. It is relevant to note that, for a fixed auxiliary-field configuration, $G_\sigma(\ell)$ generally depends on the position of the imaginary-time cut, although its determinant does not.

The convention adopted in Eq. (2.58) implies

$$\langle \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} \rangle_h = \delta_{ij} - G_{\sigma,ji}(\ell) \quad (2.61)$$

due to the anticommutation relation $\{\hat{c}_{i\sigma}, \hat{c}_{j\sigma}^\dagger\} = \delta_{ij}$. In particular, the local density is

$$\langle \hat{n}_{i\sigma} \rangle_h = 1 - G_{\sigma,ii}(\ell). \quad (2.62)$$

Since the two spin sectors are independent for a fixed auxiliary-field configuration, the double occupation factorizes as

$$\langle \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} \rangle_h = [1 - G_{\uparrow,ii}(\ell)][1 - G_{\downarrow,ii}(\ell)], \quad (2.63)$$

from which one can, for instance, compute the spin-spin correlation function

$$\langle S_i^z S_j^z \rangle = \frac{1}{4} \langle (n_{i\uparrow} - n_{i\downarrow})(n_{j\uparrow} - n_{j\downarrow}) \rangle. \quad (2.64)$$

Since the effective Hamiltonian is quadratic for a fixed auxiliary-field configuration, then higher-order equal-time correlation functions can be obtained by applying Wick's theorem. For instance, for two distinct sites and equal spins,

$$\langle \hat{n}_{i\sigma} \hat{n}_{j\sigma} \rangle_h = \langle \hat{n}_{i\sigma} \rangle_h \langle \hat{n}_{j\sigma} \rangle_h - \langle \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} \rangle_h \langle \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma} \rangle_h, \quad i \neq j, \quad (2.65)$$

which follows directly from Wick's theorem and the anticommutation relation $\hat{c}_{i\sigma} \hat{c}_{j\sigma}^\dagger = \delta_{ij} - \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}$.

Other observables, such as spin susceptibilities, can also be computed, but they require imaginary-time-dependent correlation functions and therefore the evaluation of time-displaced Green functions [13]. These are defined analogously to Eq. (2.58), but with the

fermionic operators evaluated at distinct time slices. Since such quantities are not necessary in this work, we restrict the discussion to equal-time Green functions.

2.3.3 Fast local updates

Consider the local proposal in Eq. (2.53). Since the auxiliary field enters only through the diagonal matrix $V_{\ell,\sigma}$, flipping $h_{\ell,i}$ changes only one diagonal element of the corresponding potential factor. We may write

$$\begin{aligned} e^{V'_{\ell,\sigma}} &= \text{diag}\left(e^{\nu\sigma h_{\ell,1}}, \dots, e^{\nu\sigma h'_{\ell,i}}, \dots, e^{\nu\sigma h_{\ell,N}}\right) \\ &= \text{diag}\left(e^{\nu\sigma h_{\ell,1}}, \dots, e^{\nu\sigma(h'_{\ell,i}-h_{\ell,i})} e^{\nu\sigma h_{\ell,i}}, \dots, e^{\nu\sigma h_{\ell,N}}\right) \\ &= \text{diag}\left(1, \dots, e^{\nu\sigma(h'_{\ell,i}-h_{\ell,i})}, \dots, 1\right) e^{V_{\ell,\sigma}} \\ &= [I + \Delta_{\ell,\sigma}(i)] e^{V_{\ell,\sigma}}, \end{aligned} \quad (2.66)$$

where $\Delta_{\ell,\sigma}(i)$ has a single non-zero entry,

$$[\Delta_{\ell,\sigma}(i)]_{jk} = \left(e^{\nu\sigma(h'_{\ell,i}-h_{\ell,i})} - 1\right) \delta_{ji} \delta_{ki}. \quad (2.67)$$

Under the convention adopted in Eq. (2.30), the updated one-slice propagator is therefore

$$B'_{\ell,\sigma} = [I + \Delta_{\ell,\sigma}(i)] B_{\ell,\sigma}. \quad (2.68)$$

Since the proposed flip produces a rank-one change in the propagator, it is possible to rewrite the determinant ratio in a simpler form without recomputing the determinant,

$$R_{\sigma} \equiv \frac{\det M'_{\sigma}}{\det M_{\sigma}} = 1 + [\Delta_{\ell,\sigma}(i)]_{ii} [1 - G_{\sigma,ii}(\ell)], \quad (2.69)$$

a derivation of which is given in Appendix A.3. Thus, the determinant ratio for each spin sector requires only one diagonal element of the equal-time Green function. The ratio of the complete configuration weights is

$$R \equiv R_{\uparrow} R_{\downarrow} = \frac{W(h')}{W(h)}. \quad (2.70)$$

Under the present assumption that the determinantal weights are non-negative, the local Metropolis probability becomes

$$A(h \rightarrow h') = \min\{1, R\}. \quad (2.71)$$

If the proposal is accepted, the equal-time Green function must be updated to remain consistent with the new auxiliary-field configuration. Similarly to Eq. (2.69), we may derive an update formula for the Green function without explicitly recomputing the long product of short-time propagators and doing the matrix inversion in Eq. (2.59). A precise derivation is shown in Appendix A.4 and it yields the expression

$$G'_\sigma(\ell) = G_\sigma(\ell) - \frac{1}{R_\sigma} G_\sigma(\ell) \Delta_{\ell,\sigma}(i) [I - G_\sigma(\ell)], \quad (2.72)$$

or, in matrix element form,

$$G'_{\sigma,jk}(\ell) = G_{\sigma,jk}(\ell) - \frac{1}{R_\sigma} G_{\sigma,ji}(\ell) [\Delta_{\ell,\sigma}(i)]_{ii} (\delta_{ik} - G_{\sigma,ik}(\ell)). \quad (2.73)$$

Once the initial Green function has been constructed, each accepted local update requires $\mathcal{O}(N^2)$ operations, rather than the $\mathcal{O}(N^3)$ cost of recomputing it through a new matrix inversion.

After all sites on slice ℓ have been visited, the equal-time Green function is propagated to the next time slice according to

$$G_\sigma(\ell + 1) = B_{\ell+1,\sigma} G_\sigma(\ell) B_{\ell+1,\sigma}^{-1}, \quad (2.74)$$

or, if one wishes to propagate in the opposite direction,

$$G_\sigma(\ell) = B_{\ell+1,\sigma}^{-1} G_\sigma(\ell + 1) B_{\ell+1,\sigma}. \quad (2.75)$$

The same procedure is repeated as the simulation moves through the imaginary-time direction. In exact arithmetic, the determinant ratios, rank-one updates, and propagation formula would be sufficient to evolve the chain indefinitely. In floating-point arithmetic, however, repeated multiplication by the non-unitary propagators $B_{\ell,\sigma}$ and $B_{\ell,\sigma}^{-1}$ gradually introduces numerical errors. The stabilization required to control this loss of accuracy is discussed in Sec. 2.5.

The construction above provides an efficient Markov-chain algorithm whenever the determinantal weights form a non-negative distribution. In a generic fermionic system, however, $W(h)$ may take both positive and negative values. The modifications required in this case lead to the fermion sign problem.

2.4. The fermion sign problem

The MC construction in the previous section assumed that the determinantal weight could be normalized into a probability distribution. In general, however, the product

$$W(h) = \det M_{\uparrow}(h) \det M_{\downarrow}(h) \quad (2.76)$$

is not guaranteed to be non-negative. A signed quantity cannot itself be interpreted as a probability, so the distribution in Eq. (2.48) cannot be used directly for sampling configurations through MC methods.

We define the sign of an auxiliary-field configuration by

$$s(h) = \frac{W(h)}{|W(h)|} \in \{-1, +1\}. \quad (2.77)$$

Throughout this thesis, the configuration sign refers to the sign of the complete product $\det M_{\uparrow} \det M_{\downarrow}$, rather than to the sign of either determinant separately.

The standard procedure is to sample the absolute value of the determinantal weight and include the sign in the measurements [3, 43]. We introduce the absolute-weight partition function

$$Z_{|W|} = \sum_h |W(h)| \quad (2.78)$$

and the corresponding normalized distribution

$$P_{|W|}(h) = \frac{|W(h)|}{Z_{|W|}}. \quad (2.79)$$

Since this distribution is positive, it can be sampled using the Markov-chain procedure described above. The local acceptance probability in Eq. (2.71) must now be replaced by

$$A(h \rightarrow h') = \min \left\{ 1, \left| \frac{W(h')}{W(h)} \right| \right\} = \min \{1, |R|\}. \quad (2.80)$$

The determinant-ratio and Green-function formulas themselves remain unchanged. Only the absolute value of the total ratio enters the acceptance probability, while its sign records whether the update moves between the two sign sectors.

For an observable represented by the auxiliary-field estimator $O(h)$, its physical expectation value is recovered through reweighting,

$$\begin{aligned}\langle O \rangle_W &= \frac{\sum_h O(h)W(h)}{\sum_h W(h)} \\ &= \frac{\sum_h O(h)s(h)|W(h)|}{\sum_h s(h)|W(h)|} \\ &= \frac{\langle sO \rangle_{|W|}}{\langle s \rangle_{|W|}}.\end{aligned}\tag{2.81}$$

The sign has therefore not been eliminated. It has been transferred from the sampling distribution to the estimator, so every physical measurement becomes a ratio whose denominator is the average sign,

$$\langle s \rangle_{|W|} = \frac{\sum_h W(h)}{\sum_h |W(h)|} = \frac{Z}{Z_{|W|}}.\tag{2.82}$$

When one sign sector dominates, $\langle s \rangle_{|W|}$ is close to unity and the reweighting procedure introduces little additional uncertainty. When positive and negative contributions occur with nearly equal total magnitude, they cancel strongly in the numerator of Eq. (2.82). The average sign then becomes small and must be estimated as the difference between much larger positive contributions.

The resulting loss of precision can be understood by treating Eq. (2.81) as a ratio estimator. For M statistically independent configurations, the uncertainty typically contains a factor of the form

$$\delta O \sim \frac{1}{\sqrt{M} \langle s \rangle_{|W|}},\tag{2.83}$$

up to an observable-dependent combination of variances and covariance. For Markov-chain data, M must be replaced by the effective number of independent samples, so autocorrelation further increases the computational effort. As such, it follows that for reduced average sign one must increase the number of samples accordingly to achieve a fixed target error bar. This effect may be worsened whenever the sign variable itself has a long autocorrelation time, corresponding to a slow mixing of the sign sectors under the Markov-chain dynamics; consequently the estimator for the average sign $\langle s \rangle$ itself acquires larger error bars and converges more slowly to its equilibrium value. Since reweighting involves the ratio of two sample means, $\overline{sO}/\overline{s}$, such ratio estimators can exhibit non-Gaussian, heavy-tailed sampling distributions when fluctuations of the estimated denominator bring it close to zero [46]. Autocorrelation therefore compounds the statistical

difficulty of reweighting by reducing the effective number of independent measurements entering both sample means.

In that case, the estimator of the average sign $\langle s \rangle$ has an increased sampling variance and converges more slowly to its equilibrium value.

The average sign also has a thermodynamic interpretation. Defining the free-energy densities of the original and absolute-weight ensembles by

$$f = -\frac{1}{\beta N} \log Z, \quad f_{|W|} = -\frac{1}{\beta N} \log Z_{|W|}, \quad (2.84)$$

one obtains

$$\langle s \rangle_{|W|} = e^{-\beta N \Delta f}, \quad \Delta f = f - f_{|W|} \geq 0. \quad (2.85)$$

When Δf remains finite and positive in the thermodynamic limit, the average sign decreases exponentially with inverse temperature and system size [3, 4]. Since the number of independent samples required to maintain a fixed relative uncertainty scales approximately as the inverse square of the average sign, the computational effort grows exponentially. This exponential loss of statistical precision is the practical manifestation of the fermion sign problem.

Whether a sign problem is present is not determined solely by the interacting Hamiltonian. It also depends on the single-particle basis, the chosen HS decomposition, and the symmetries of the quadratic fermionic problem produced for each auxiliary-field configuration [20, 26, 42]. An important example is the repulsive Hubbard model at half filling on a bipartite lattice. At $\mu = 0$, a particle-hole transformation on one spin species relates the two spin determinants so that their product is non-negative for every auxiliary-field configuration [15, 37, 43]. The model is therefore sign-problem-free at this point and

$$\langle s \rangle_{|W|} = 1. \quad (2.86)$$

A derivation of this result is given in Appendix A.5. Away from half filling, the chemical-potential term breaks the particle-hole relation and negative configuration weights are generally present [21, 47]. The simulations discussed in this thesis are performed in this regime.

2.5. Numerical stabilization

The matrices $B_{\ell,\sigma}$ represent imaginary-time evolution under a one-body Hamiltonian. When we multiply many such matrices, the resulting propagator develops singular values which grow or decay exponentially with the total imaginary-time interval. This becomes an issue at low temperatures, as the full product $B_{L,\sigma} \cdots B_{1,\sigma}$ contains scales that span many orders of magnitude and so its evaluation becomes severely ill-conditioned [48, 49]. This is especially true in the evaluation of the fermionic matrix in Eq. (2.35), as the addition of the identity causes a loss of the smallest scales. A direct evaluation of the Green function through the inversion of this matrix may then return a matrix that no longer represents the correct physical propagator.

This problem is separate from the fermion sign problem. Stabilization controls deterministic floating-point errors and allows the determinants and Green functions to be evaluated accurately at large β . It does not change the statistical cancellation between positive and negative configurations.

2.5.1 Stabilized matrix products

To preserve the different numerical scales in a long matrix product, we factorize it periodically in the form

$$B_{m,\sigma} \cdots B_{1,\sigma} = U_{m,\sigma} D_{m,\sigma} V_{m,\sigma}, \quad (2.87)$$

where D is diagonal and contains the large and small scales of the product, while U and V are well-conditioned matrices. In the implementation used here, this decomposition is constructed from a column-pivoted QR decomposition, with the diagonal magnitudes of the triangular factor being extracted into D [49, 50].

The factorization can then be updated recursively as one extends the product in imaginary time. For notational simplicity, we suppress the spin index in the following. Suppose then that a partial product has already been written as

$$P_{m-1} = U_{m-1} D_{m-1} V_{m-1}. \quad (2.88)$$

After multiplying by a new block $\mathcal{B}_m$, one has

$$\begin{aligned} P_m &= \mathcal{B}_m P_{m-1} \\ &= \mathcal{B}_m U_{m-1} D_{m-1} V_{m-1}. \end{aligned} \quad (2.89)$$

The product containing this newly introduced block is then factorized again,

$$\mathcal{B}_m U_{m-1} D_{m-1} = U_m D_m \tilde{V}_m, \quad (2.90)$$

and the remaining factor is absorbed according to

$$V_m = \tilde{V}_m V_{m-1}. \quad (2.91)$$

The updated partial product therefore retains the form

$$P_m = U_m D_m V_m. \quad (2.92)$$

By proceeding in this manner, the exponentially separated scales remain isolated in the diagonal factor D_m , and the full ill-conditioned product never needs to be formed explicitly as a single dense matrix.

The inverse in Eq. (2.59) must be constructed with equal care. In our algorithmic procedure, at every instant we store the UDV representation of each of the two long matrix products appearing in the cyclic product. Representing these as $U_1 D_1 V_1$ and $U_2 D_2 V_2$, the inverse procedure is applied as

$$\begin{aligned} G &= [I + U_1 D_1 V_1 U_2 D_2 V_2]^{-1} \\ &= [U_1 (U_1^{-1} V_2^{-1} + D_1 V_1 U_2 D_2) V_2]^{-1} \\ &= [U_1 (U' D' V') V_2]^{-1} \\ &= (V' V_2)^{-1} D'^{-1} (U_1 U')^{-1}, \end{aligned} \quad (2.93)$$

where an intermediate UDV decomposition was evaluated. The important point is that the scale information stored in D is retained throughout the inversion.

Throughout this thesis, a central focus is the analysis of the magnitude of the determinantal weights and their corresponding signs. Directly reconstructing the matrices M_σ and evaluating their determinants after every local update would, however, be both computationally costly and numerically unsafe. Besides requiring $\mathcal{O}(N^3)$ operations, the determinants may span many orders of magnitude and therefore be susceptible to overflow, underflow, and loss of precision. The stabilized factorizations used to reconstruct the equal-time Green functions also provide a reliable way of evaluating these quantities. In particular, once M_σ has been expressed in the form

$$M_\sigma = U_\sigma D_\sigma V_\sigma, \quad (2.94)$$

where U_σ and V_σ are well conditioned and D_σ contains the separated numerical scales, the logarithm of the determinant magnitude is evaluated as

$$\log |\det M_\sigma| = \log |\det U_\sigma| + \sum_i \log |[D_\sigma]_{ii}| + \log |\det V_\sigma|. \quad (2.95)$$

The corresponding determinant sign is

$$s_\sigma = \text{sgn}(\det U_\sigma) \text{sgn}(\det V_\sigma) \prod_i \text{sgn}([D_\sigma]_{ii}). \quad (2.96)$$

The remaining determinants involve only well-conditioned matrices and may be evaluated using a pivoted QR decomposition. The weight and configuration sign are then

$$\log |W(h)| = \sum_\sigma \log |\det M_\sigma|, \quad (2.97)$$

$$s(h) = \prod_{\sigma=\pm 1} s_\sigma. \quad (2.98)$$

Between stabilized reconstructions, these quantities may be updated after an accepted local proposal using

$$\log |W(h')| = \log |W(h)| + \log |R|, \quad s(h') = s(h) \text{sgn}(R). \quad (2.99)$$

They are periodically recomputed from the stabilized factors in order to prevent the accumulation of round-off errors in the successive local updates.

2.5.2 Stable propagation through imaginary time

Reconstructing a stabilized product after every local proposal would remove the efficiency gained from the rank-one update. Instead, we follow the stabilized propagation scheme described in [51], in which the imaginary-time axis is divided into $L_s = L/\Delta_s$ blocks containing Δ_s slices. Within each block we use Eqs. (2.72) and (2.74), while at the block boundaries the Green function is recomputed from stabilized partial products. The stabilization interval Δ_s is chosen sufficiently small that the error accumulated by direct propagation remains below the desired numerical tolerance.

Let

$$\bar{B}_{s,\sigma} \equiv B_{s\Delta_s,\sigma} \cdots B_{(s-1)\Delta_s+1,\sigma} \quad (2.100)$$

denote a block propagator, and let

$$\begin{aligned}\bar{B}_{s,s',\sigma} &\equiv \bar{B}_{s,\sigma} \bar{B}_{s-1,\sigma} \cdots \bar{B}_{s'+1,\sigma} \\ &= \underbrace{B_{s\Delta_s,\sigma} B_{s\Delta_s-1,\sigma} \cdots B_{s'\Delta_s+1,\sigma}}_{(s-s')\Delta_s}\end{aligned}\quad (2.101)$$

be the product of such block propagators. Under this notation, the equal-time Green function at time slices on the boundaries of these blocks, i.e., $\ell = s\Delta_s$, is written as

$$G_\sigma(\ell) = \left[I + \bar{B}_{s,0,\sigma} \bar{B}_{L_s,s,\sigma} \right]^{-1}. \quad (2.102)$$

The two products appearing in this expression correspond to the two parts of the cyclic imaginary-time product entering the equal-time Green function. The factor $\bar{B}_{s,0,\sigma}$ contains the propagation from $\tau = 0$ up to the current block boundary $\ell = s\Delta_s$, whereas $\bar{B}_{L_s,s,\sigma}$ contains the complementary propagation from that boundary to $\tau = \beta$. Both of these products are stored in stabilized factorized form. During a forward propagation in imaginary time, $\bar{B}_{s,0,\sigma}$ is built recursively as successive blocks are crossed, while the complementary factors $\bar{B}_{L_s,s,\sigma}$ are taken from products stored during the preceding reverse propagation. At each block boundary, the two stabilized products are combined to reconstruct the cyclic product entering the Green function in Eq. (2.102). In this way, the Green function only needs to be propagated directly over the Δ_s slices contained within each block, limiting the accumulation of round-off errors while preserving the $\mathcal{O}(N^2)$ cost of the local rank-one updates between stabilizations.

The procedure is then repeated in the opposite direction. After reaching $\tau = \beta$, the simulation traverses the imaginary-time axis in decreasing slice order. During this reverse pass, the products required for the next forward traversal are rebuilt analogously, with the roles of the two partial products interchanged. The stabilized factors generated during one traversal are therefore reused during the next, avoiding the need to reconstruct all partial products independently at every block boundary [51].

2.6. Simulation procedure and validation

A DQMC calculation begins by choosing the lattice geometry and the parameters U , μ , β , L , and Δ_s . The auxiliary field is initialized, the matrices $B_{\ell,\sigma}$ are constructed, and the stabilized block products are built. The initial Green functions are then reconstructed from these products. Starting from this configuration, the simulation performs a sequence of local Metropolis proposals. For each space-time point, the two determinant ratios are

evaluated from Eq. (2.69); the flip is then accepted with the probability in Eq. (2.80), and, if accepted, both Green functions are updated using Eq. (2.72). After all sites of a time slice have been visited, the Green functions are propagated to the next slice. At each block boundary, they are reconstructed from stabilized factors. After the simulation reaches the final time slice, the propagation procedure is applied in the opposite time direction.

The first part of the Markov chain is discarded as thermalization. During this stage, the field evolves from the arbitrary initial configuration toward the absolute-weight distribution in Eq. (2.79). Measurements are retained only after the monitored quantities fluctuate around stationary values.

We use the following units of MC time. An elementary proposal is one attempted flip of a single field variable $h_{\ell,i}$. A time-slice pass is one traversal of the N spatial sites at fixed ℓ . A forward pass visits all L slices in increasing order and a reverse pass visits them in decreasing order. One complete MC sweep is defined as one forward pass followed by one reverse pass and therefore contains $2NL$ elementary proposals. Quantities used to compare simulations with different numbers of time slices are naturally expressed in sweeps.

Successive measurements obtained from the same Markov chain are correlated. For an observable O_t sampled at equally spaced times, we define the normalized autocorrelation function

$$\rho_O(k) = \frac{\langle O_t O_{t+k} \rangle - \langle O \rangle^2}{\langle O^2 \rangle - \langle O \rangle^2}. \quad (2.103)$$

The corresponding integrated autocorrelation time is

$$\tau_{\text{int},O} = 1 + 2 \sum_{k=1}^{\infty} \rho_O(k), \quad (2.104)$$

where the sum is truncated using a finite window in practical estimates. If M measurements are stored, the effective number of independent samples is approximately

$$M_{\text{eff}} \simeq \frac{M}{\tau_{\text{int},O}}, \quad (2.105)$$

that is, it is reduced by a factor equal to $\tau_{\text{int},O}$. To get around this issue, statistical errors are not estimated from the full series of measurements. Instead, one divides the time series data into bins of size larger than the autocorrelation time, and the variance of the bin averages is used to estimate the uncertainty. For a reweighted observable, the numerator and denominator of Eq. (2.81) are blocked together. Independent simulation threads are

thermalized and analyzed separately before their estimates are combined.

The implementation was validated at several levels. The local determinant ratio and the rank-one Green-function update were first compared against direct determinant evaluations and matrix inversions on small systems. The propagated Green functions were then compared with stabilized reconstructions at the block boundaries. For a four-site 1D Hubbard system, the equal-time spin correlation functions obtained from DQMC were also found to agree with the results obtained from exact diagonalization within the MC uncertainties, as can be seen in Fig. 2.1. With the goal of identifying a parameter regime in which the sign problem is severe, we also validated the dependence of the average sign on the band filling by comparing our results for the Hubbard model on a square lattice with those obtained using the QUEST DQMC code [21, 52], as shown in Fig. 2.2.




Figure 2.1: Comparison between the equal-time spin correlations obtained with DQMC and ED for a four-site system. The left panel shows the spin correlations as a function of the distance from a reference site, with the exact result compared with the S^z and S^x correlations measured in DQMC. The right panel shows the corresponding relative deviations from the exact values, with the error bars representing the statistical uncertainties of the DQMC estimates.




Figure 2.2: Average sign $\langle s \rangle$ as a function of the average band filling $\langle n \rangle$, comparing the present DQMC implementation with the QUEST code. The two calculations show excellent agreement over the full filling range. The average sign is equal to unity close to the empty- and half-filled limits and decreases at intermediate fillings, reaching its minimum around $\langle n \rangle \simeq 0.85$. This comparison both validates the treatment of the determinant sign in the present implementation and identifies a parameter region in which the sign problem is sufficiently pronounced for the analysis developed in the following chapters.

3. Effective-energy and sign-resolved structure of configuration space

The average sign introduced in Chapter 2 quantifies the total cancellation between positive and negative determinantal weights, but compresses the entire auxiliary-field configuration space into a single number. It does not reveal whether the two sign sectors occupy different ranges of determinant magnitude, whether the configurations dominating the absolute-weight ensemble are especially balanced between the two signs, or how the severity of the cancellation changes across configuration space. In this chapter, we begin to resolve this internal structure by using the logarithm of the determinant magnitude as an effective-energy coordinate. We first study its canonical distribution and establish its dependence on the imaginary-time discretization, temperature, and system size. We then introduce sign-resolved densities of configurations and their associated entropies. Since ordinary DQMC samples only the narrow range selected by the absolute weight, we use generalized-ensemble sampling and the Wang-Landau method to determine these entropies over a broader interval. Finally, we analyze their thermodynamic scaling and show how the local balance between positive and negative configurations determines the global average sign.

3.1. The effective-energy variable

The determinantal formulation introduced in the previous chapter assigns to each auxiliary-field configuration h a weight

$$W(h) = \det M_{\uparrow}(h) \det M_{\downarrow}(h), \quad (3.1)$$

which may be either positive or negative. Before introducing a statistical description of these weights, it is useful to examine how their magnitude and sign evolve during a conventional DQMC simulation. Figure 3.1 shows the logarithm of the weight magnitude during simulations performed at two inverse temperatures. Positive- and negative-weight configurations are distinguished by color.

The magnitude of the determinantal weight fluctuates over a broad range along both trajectories in Fig. 3.1, with the typical scale increasing substantially as the inverse temperature is increased. This makes a direct description in terms of $|W|$ inconvenient both numerically and graphically. This broad variation motivates the introduction of the logarithmic




Figure 3.1: Representative Monte Carlo time series of the logarithm of the determinantal weight magnitude at $\beta = 4$ and $\beta = 8$. Blue and red segments correspond to configurations with positive and negative determinantal weights, respectively. The average signs measured during the simulations are displayed above each panel. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

coordinate

$$x(h) = -\log |W(h)|, \quad (3.2)$$

which will be referred to as the *effective energy* of the auxiliary-field configuration. With this definition,

$$W(h) = s(h)e^{-x(h)}, \quad (3.3)$$

so the magnitude of the determinantal weight takes the same exponential form as a classical Boltzmann factor. The terminology “effective energy” does not imply that $x(h)$ is the physical energy of the Hubbard model. It is instead a scalar coordinate assigned to each auxiliary-field configuration through the magnitude of its fermionic determinant. Unlike the physical energy, which is obtained as an expectation value of the Hamiltonian, $x(h)$ characterizes the statistical importance of an individual configuration in the DQMC ensemble. It is nevertheless a natural variable for describing configuration space because it converts weights spanning many orders of magnitude into an additive quantity.

At the same time, the trajectory repeatedly moves between the positive and negative sign sectors. The sign changes appear irregular and display no obvious long-time pattern by visual inspection, particularly at $\beta = 8$, where the average sign is smaller. This observation is only qualitative and does not imply that successive signs are statistically independent. For the purposes of the present work, however, the relevant information is not the temporal ordering of these sign changes, but how positive- and negative-weight configurations are distributed over effective-energy space. We therefore retain the sign as an additional

discrete label and organize the configuration space through the projection

$$h \mapsto (x(h), s(h)). \quad (3.4)$$

This mapping discards the microscopic arrangement of the NL auxiliary variables, but retains the two properties that determine the magnitude and sign of the contribution of the configuration to the partition function. It therefore provides a convenient starting point for studying how the positive and negative sectors are distributed across configuration space.

The time series in Fig. 3.1 show how the determinantal weight evolves along a single Monte Carlo trajectory. A more statistical view is obtained by considering the distribution of its magnitude over the sampled configurations. Figure 3.2 shows representative normalized histograms of $x = -\log |W|$, resolved according to the sign of the determinantal weight, for the same model parameters at two different inverse temperatures. The histograms are normalized such that

$$\int dx [P_+(x) + P_-(x)] = 1, \quad (3.5)$$

where $P_\sigma(x)$ denotes the probability of the simulation sampling a weight with effective energy x and sign σ^* .




Figure 3.2: Representative normalized distributions of the effective energy for configurations with positive and negative sign. The left and right panels correspond to $\beta = 4$ and $\beta = 8$, respectively. Blue and red histograms denote configurations with $W > 0$ and $W < 0$, while the dashed lines show Gaussian fits to the corresponding distributions. The simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

*Here and throughout the remainder of this dissertation, the variable σ refers to the determinant sign sector and should not be confused with the physical spin index of Chapter 2.

At $\beta = 4$, the positive- and negative-sign sectors occupy broadly the same range of effective energies, but the positive sector is still visibly more populated, consistent with the finite average sign shown in the figure. Furthermore, their distributions are approximately Gaussian, with similar centers and widths for the two signs. At $\beta = 8$, where the average sign is much smaller, the two distributions become nearly indistinguishable both in shape and in integrated probability. This last point implies that the average sign is indeed quite close to zero, since

$$\langle s \rangle = \int dx [P_+(x) - P_-(x)]. \quad (3.6)$$

These two observations already suggest that the sign problem is associated not with a small isolated set of negative-weight configurations, but rather with an almost complete overlap between the sampled positive and negative sectors in the space of weight magnitudes, leading to a near-complete cancellation between them.

3.1.1 Scaling properties of the effective energy

The representative trajectories and distributions above show that the typical value and width of the effective-energy variable change substantially with the simulation parameters. To compare different temperatures, system sizes, and imaginary-time discretizations, it is therefore necessary to determine how x scales with the corresponding space-time volume. We begin by examining its dependence on the imaginary-time step $\Delta\tau$ and then study its scaling with β and N .

Before making these comparisons, it is important to recall that the determinantal weight is defined only up to a positive configuration-independent normalization. Under the transformation

$$W(h) \mapsto A W(h), \quad (3.7)$$

where A does not depend on h , the effective energy is shifted according to

$$x(h) \mapsto x(h) - \log A. \quad (3.8)$$

Such a transformation does not affect the Markov-chain dynamics, normalized expectation values, the configuration sign, or the width of the effective-energy distribution. It does, however, shift its mean. Consequently, comparisons of the absolute value of x between different simulations are meaningful only after fixing a common normalization convention. Throughout this work, we use the determinant weight defined in Eq. (2.37), in which the

positive configuration-independent factor C^{NL} introduced by the HS transformation has been omitted.

The expected scaling of x can be understood directly from the determinantal representation. For a fixed auxiliary-field configuration, the matrices appearing in the determinantal weight are given by

$$M_\sigma(h) = I + B_{L,\sigma} \cdots B_{1,\sigma}. \quad (3.9)$$

Denoting the singular values of M_σ by $d_{\sigma,a}$, with $a = 1, \dots, N$, the determinant magnitude satisfies

$$\log |\det M_\sigma(h)| = \sum_{a=1}^N \log d_{\sigma,a}, \quad (3.10)$$

and the effective energy can therefore be written as

$$x(h) = - \sum_{\sigma} \sum_{a=1}^N \log d_{\sigma,a}. \quad (3.11)$$

Each product $B_{L,\sigma} \cdots B_{1,\sigma}$ describes propagation through a total imaginary-time interval of length β . Products of non-unitary imaginary-time propagators generally develop singular values which grow or decay exponentially with this interval [49],

$$d_{\sigma,a} \sim e^{\beta \gamma_{\sigma,a}}, \quad (3.12)$$

up to finite- β corrections, where $\gamma_{\sigma,a}$ denotes the corresponding asymptotic growth rate. Consequently, $\log d_{\sigma,a}$ is generically of order β . Since there are N single-particle modes in each spin sector, Eq. (3.11) suggests the extensive scaling

$$\langle x \rangle = \beta N x^{(\infty)} + o(\beta N), \quad (3.13)$$

where the coefficient $x^{(\infty)}$ depends on the interaction strength, chemical potential, lattice geometry, and normalization convention. This motivates defining the intensive effective-energy variable

$$\bar{x}(h) = \frac{x(h)}{\beta N}, \quad (3.14)$$

which will be used throughout the remainder of this chapter. If x is extensive, the distribution of $\bar{x}$ should remain centered at a finite value as either β or N is increased. On the other hand, reducing $\Delta\tau$ at fixed β and N increases the number of time slices according to $L = \beta/\Delta\tau$, but it does not change the total imaginary-time interval represented by the

product of propagators. In the limit $\Delta\tau \rightarrow 0$, this product converges to the corresponding continuous imaginary-time evolution operator. Once $\Delta\tau$ is sufficiently small, the distribution of $\bar{x}$ should therefore approach a finite limit which is independent of the discretization, apart from the residual Trotter error,

$$\langle \bar{x} \rangle(\Delta\tau) = \bar{x}^{(0)} + \mathcal{O}(\Delta\tau). \quad (3.15)$$

This convergence depends on the normalization convention adopted above. Had the HS prefactor C^{NL} been retained in the configuration weight, it would contribute

$$-\log C^{NL} = -NL \log C \quad (3.16)$$

to the effective energy. Since $L = \beta/\Delta\tau$, this term would contain an explicit and physically irrelevant dependence on the imaginary-time discretization. By omitting this positive configuration-independent prefactor, the remaining $\Delta\tau$ dependence reflects only the approximation of the imaginary-time evolution and should vanish in the continuous-time limit.

Figure 3.3 shows the distributions of $\bar{x}$ obtained for several imaginary-time steps, separately for positive- and negative-weight configurations.




Figure 3.3: Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different imaginary-time steps. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. The remaining simulation parameters are $N = 16$, $\beta = 4$, $U = 6$, and $\mu = -1.5$.

For the coarsest discretizations, both the position and width of the distributions depend visibly on $\Delta\tau$. As the imaginary-time step is reduced, this systematic displacement becomes progressively smaller and the distributions approach a common limiting region. This behavior is consistent with the interpretation that the remaining dependence is a discretization effect rather than an intrinsic scaling with the number of time slices.

The convergence of the mean is displayed more directly in Fig. 3.4. For each sign sector, the central point represents the sign-resolved mean of $\bar{x}$, while the vertical bar represents

the standard deviation of the corresponding distribution. These bars therefore quantify the physical width of the sampled distribution and should not be interpreted as statistical uncertainties on its mean.




Figure 3.4: Sign-resolved mean of the intensive effective energy as a function of the inverse imaginary-time step. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively. The vertical bars show one standard deviation of each effective-energy distribution rather than the statistical uncertainty of the estimated mean.

The mean initially shifts as the discretization is refined but ceases to display a significant systematic drift for the smallest values of $\Delta\tau$. These results are consistent with convergence toward a finite and well-defined continuous-time limit for $\bar{\epsilon}$.

Having established that the distributions converge as the imaginary-time discretization is refined, we now examine their dependence on the total imaginary-time extent, that is, on the inverse temperature β . Figure 3.5 shows the sign-resolved distributions of the intensive effective energy for several inverse temperatures at fixed lattice size and imaginary-time step.

At the smallest values of β , the distributions are comparatively broad and their centers depend appreciably on the inverse temperature. As β is increased, they become increasingly concentrated and their centers approach a common finite region of $\bar{\epsilon}$. This behavior is consistent with an extensive leading contribution to the effective energy, supplemented by finite-temperature corrections, as in Eq. (3.13). Such a scaling behavior predicts that their centers should approach a finite asymptotic value as the inverse temperature becomes large. A direct test is obtained by multiplying the mean of $\bar{\epsilon}$ by β and checking whether it depends linearly on β .




Figure 3.5: Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different inverse temperatures. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. We note the absence of distributions of the negative weights for $\beta = 1$ and 2 , meaning that no negative-signed configurations were visited during the simulation. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

As shown in Fig. 3.6, the quantity $\beta\langle\bar{x}\rangle$ is approximately linear in β over the range considered. Equivalently, the mean of x grows approximately linearly with inverse temperature at fixed N . The results for the two sign sectors are nearly indistinguishable on the scale of the figure, showing that positive- and negative-weight configurations have very similar typical determinant magnitudes even when their contributions to the partition function have opposite signs.




Figure 3.6: Scaling of the mean effective energy with inverse temperature. The plotted quantity satisfies $\beta\langle\bar{x}\rangle = \langle x \rangle/N$. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively, while the dashed line shows a common linear fit. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

We next study the spatial scaling by varying the number of lattice sites at fixed inverse temperature and imaginary-time interval. The resulting distributions are shown in Fig. 3.7.




Figure 3.7: Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different lattice sizes. The left and right panels correspond to configurations with $W > 0$ and $W < 0$, respectively. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

The distributions remain centered within a finite range of $\bar{x}$ as the lattice size is increased, while their widths decrease substantially. The smallest lattice displays the largest finite-size corrections, but the centers of the distributions rapidly approach a common region for larger systems. As before, a direct test is done by verifying the linear dependence of the quantity $N\langle\bar{x}\rangle = \langle x \rangle / \beta$.




Figure 3.8: Scaling of the mean effective energy with the number of lattice sites. The plotted quantity satisfies $N\langle\bar{x}\rangle = \langle x \rangle / \beta$. Blue and red points correspond to configurations with positive and negative determinantal weights, respectively, while the dashed line shows a linear fit. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

The linear behavior in Fig. 3.8 confirms that the mean effective energy is proportional to the number of lattice sites at fixed β . Together with the corresponding temperature scaling, this supports the leading form

$$\langle x \rangle \sim \beta N. \quad (3.17)$$

This is the same leading scaling obeyed by the thermodynamic energy of a classical extensive system. For the analogy between x and a classical energy to extend beyond the scaling of the mean, its fluctuations should also exhibit the usual thermodynamic behavior. In particular, at fixed β , the variance of an extensive energy is expected to grow proportionally to N , so that the variance of the corresponding intensive quantity scales as

$$\sigma_{\bar{x}}^2 \equiv \langle \bar{x}^2 \rangle - \langle \bar{x} \rangle^2 \sim \frac{1}{N}. \quad (3.18)$$

Testing this relation therefore provides the remaining scaling check required for $\bar{x}$ to behave as an intensive energy-like variable in the usual statistical-mechanical sense.




Figure 3.9: Variance of the intensive effective energy as a function of the inverse lattice size. Blue and red points correspond to positive- and negative-weight configurations, respectively. The dashed lines show fits of the form $\sigma_{\bar{x}}^2 \propto 1/N$ excluding the smallest lattice, for which finite-size effects are most pronounced. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

For $N > 4$, the data in Fig. 3.9 are consistent with an approximately linear dependence on $1/N$. The smallest system deviates from this trend, particularly in the negative-sign sector. We attribute this deviation to finite-size effects and therefore exclude the smallest system from the fits. The observed scaling implies

$$\sigma_{\bar{x}} \sim \frac{1}{\sqrt{N}}, \quad (3.19)$$

so that $\bar{x}$ is consistent with self-averaging in the thermodynamic limit. Although the absolute fluctuations of x increase with the system size, they grow more slowly than its mean, and the relative width of the effective-energy distribution consequently decreases.

3.2. Sign-resolved density of states and entropy

The effective-energy variable introduced above provides a one-dimensional coordinate with which the auxiliary-field configuration space can be organized, as configurations with very different microscopic arrangements of the HS field may have the same value of x , and may therefore be grouped into effective-energy shells. Since configurations within the same shell may nevertheless have opposite determinantal signs, the density of states must be resolved according to the sign sector. We therefore define the sign-resolved density of states as

$$N_\sigma(x) = \sum_h \delta_{s(h),\sigma} \delta(x - x(h)), \quad (3.20)$$

where the sum runs over all auxiliary-field configurations. Thus, $N_+(x) dx$ and $N_-(x) dx$ count the numbers of positive- and negative-weight configurations, respectively, whose effective energies lie in the interval $[x, x + dx]$. The total density of configurations is then

$$N(x) = N_+(x) + N_-(x). \quad (3.21)$$

In analogy with the microcanonical entropy of classical statistical mechanics [53], we define the sign-resolved entropies

$$S_\sigma(x) = \log N_\sigma(x). \quad (3.22)$$

These functions quantify the number of auxiliary-field configurations available at a given effective energy in each sign sector, and they should not be confused with the thermodynamic entropy of the original Hubbard model. Rather, they describe the organization of the auxiliary-field configuration space after it has been projected onto the variables (x, σ) .

Using $W(h) = s(h)e^{-x(h)}$, the partition function may be written in terms of the sign-resolved densities as

$$Z = \sum_h s(h)e^{-x(h)} = \int dx [N_+(x) - N_-(x)] e^{-x}, \quad (3.23)$$

while the corresponding partition function of the absolute-weight ensemble is

$$Z_{|W|} = \sum_h |W(h)| = \int dx [N_+(x) + N_-(x)] e^{-x}. \quad (3.24)$$

Since the probability of visiting a given configuration h in the absolute-weight ensemble is given by

$$P(h) = \frac{|W(h)|}{Z_{|W|}} = \frac{e^{-x(h)}}{Z_{|W|}}, \quad (3.25)$$

then the sign-resolved probability densities measured in the absolute-weight ensemble are therefore

$$P_\sigma(x) \equiv \sum_h \delta_{s(h),\sigma} \delta(x(h) - x) P(h) = \frac{N_\sigma(x) e^{-x}}{Z_{|W|}} = \frac{e^{S_\sigma(x) - x}}{Z_{|W|}}. \quad (3.26)$$

The histograms discussed in Fig. 3.2 are precisely canonical distributions of this form. The density of states and the canonical distribution contain different information, since $N_\sigma(x)$ counts just the number of configurations in each shell, whereas the factor e^{-x} determines the statistical importance assigned to that shell in the absolute-weight ensemble.

The competition between these two factors is analogous to the usual competition between energy and entropy in classical statistical mechanics. In each sign sector, the probability of observing a given effective energy is controlled by

$$P_\sigma(x) \propto \exp[S_\sigma(x) - x]. \quad (3.27)$$

A region of configuration space may therefore contain a very large number of states but still have a small canonical probability if their effective energies are too large. Conversely, configurations with small x receive large individual weights, although the corresponding shell may contain comparatively few states.

The sign structure may also be characterized locally in effective-energy space. This is done by introducing the average sign restricted to one effective-energy shell,

$$\sigma(x) = \frac{N_+(x) - N_-(x)}{N_+(x) + N_-(x)}. \quad (3.28)$$

We shall refer to $\sigma(x)$ as the *local average sign*, where *local* refers to localization in effective-energy space rather than in the physical lattice. Its magnitude measures the imbalance between the positive and negative sectors within a single effective-energy shell. It satisfies $-1 \leq \sigma(x) \leq 1$, with $\sigma(x) = 1$ when the shell contains only positive configurations, $\sigma(x) = -1$ when it contains only negative configurations, and $\sigma(x) = 0$ when both signs occur equally often.

In terms of the sign-resolved entropies, Eq. (3.28) may be rewritten as

$$\sigma(x) = \frac{e^{S_+(x)} - e^{S_-(x)}}{e^{S_+(x)} + e^{S_-(x)}} = \tanh\left(\frac{S_+(x) - S_-(x)}{2}\right). \quad (3.29)$$

The fixed- x sign is therefore determined entirely by the entropy difference

$$\Delta S(x) = S_+(x) - S_-(x). \quad (3.30)$$

Finally, the usual average sign can be expressed as an average of $\sigma(x)$ over the canonical distribution of the effective energy,

$$\begin{aligned} \langle s \rangle_{|W|} &= \frac{\int dx (N_+(x) - N_-(x)) e^{-x}}{\int dx (N_+(x) + N_-(x)) e^{-x}} \\ &= \frac{\int dx \sigma(x) (N_+(x) + N_-(x)) e^{-x}}{\int dx (N_+(x) + N_-(x)) e^{-x}} \\ &= \langle \sigma(x) \rangle_{P(x)}, \end{aligned} \quad (3.31)$$

where $\langle \sigma(x) \rangle_{P(x)}$ denotes the average of $\sigma(x)$ with respect to the distribution

$$P(x) \equiv P_+(x) + P_-(x). \quad (3.32)$$

The global average sign is thus determined jointly by the sign imbalance within each shell and by the probability with which the different effective energies are sampled.

The functions $S_{\pm}(x)$ therefore provide a complete sign-resolved description of configuration space at the level of the effective-energy coordinate. The next step is to determine these functions numerically over the relevant range of x . For that purpose, it is relevant to note that, in a numerical calculation, the effective-energy axis is divided into finite bins rather than treated as a continuous variable. The densities in Eq. (3.20) should therefore be understood as the continuum representation of the corresponding histogram counts. Changing the bin width multiplies both densities by a common constant and consequently shifts their logarithms by an additive constant. Such a shift does not affect normalized probabilities, entropy differences, or any of the quantities considered above, provided that a single binning convention is used consistently.

3.3. Wang-Landau determination of the entropy

The sign-resolved canonical distributions introduced in the previous section contain, in principle, enough information to reconstruct the entropy. From Eq. (3.26), one has

$$S_{\sigma}(x) = \log P_{\sigma}(x) + x + \text{const.} \quad (3.33)$$

In practice, however, this relation is useful only over the range of effective energies that is visited with appreciable probability. The factor e^{-x} concentrates conventional absolute-weight sampling around a comparatively narrow interval of x , while configurations outside this region occur too rarely for their density to be estimated accurately. The narrowing of the intensive effective-energy distributions with increasing system size makes this limitation progressively more severe, and so a different sampling weight is required to determine $S_\sigma(x)$ over an extended range.

To overcome this limitation, we replace the absolute-weight ensemble by a generalized ensemble in which the probability assigned to a configuration is controlled by an adjustable function of its effective energy and sign [54]. We introduce the positive sampling weight

$$P_\omega(h) \propto e^{-\omega_{s(h)}(x(h))}, \quad (3.34)$$

where $\omega_\sigma(x)$ is an arbitrary biasing function, defined separately for the two sign sectors. It then follows that the probability of observing an effective energy x in sign sector σ is given by

$$P_{\omega,\sigma}(x) \propto e^{S_\sigma(x) - \omega_\sigma(x)}. \quad (3.35)$$

The sampled distribution is therefore determined by the difference between the entropy of the shell and the imposed bias. Ordinary absolute-weight sampling is recovered by choosing

$$\omega_\sigma(x) = x, \quad (3.36)$$

independently of the sign sector, since in that case

$$P_\omega(h) \propto e^{-x(h)} = |W(h)|. \quad (3.37)$$

More generally, if a simulation performed with a known bias $\omega_\sigma(x)$ produces a sign-resolved histogram $H_{\omega,\sigma}(x)$, then

$$H_{\omega,\sigma}(x) \propto P_{\omega,\sigma}(x), \quad (3.38)$$

and the entropy may consequently be reconstructed as

$$S_\sigma(x) = \log H_{\omega,\sigma}(x) + \omega_\sigma(x) + \text{const.} \quad (3.39)$$

over the range in which the histogram has sufficient statistics. The additive constant is common to the two sign sectors when they are sampled within the same generalized

ensemble and does not affect entropy differences.

For the local auxiliary-field updates used here, the proposal probability is symmetric. A proposed transition from a configuration h to a configuration h' is therefore accepted with probability

$$A(h \rightarrow h') = \min \left\{ 1, \exp \left[-\omega_{s(h')}(x(h')) + \omega_{s(h)}(x(h)) \right] \right\}. \quad (3.40)$$

The microscopic proposal mechanism is thus unchanged from conventional DQMC, as only the probability with which the proposed configuration is accepted is modified.

A simple way of modifying the sampled range is to change the strength with which the effective energy enters the configuration weight. We therefore first consider the one-parameter family of positive sampling weights

$$P_\lambda(h) \propto e^{-\lambda x(h)} = |W(h)|^\lambda, \quad (3.41)$$

obtained by setting $\omega_\sigma(x) = \lambda x$, where λ is a real parameter. We note that this one-parameter family is equivalent to the weight rescaling introduced in Ref. [29]. Conventional absolute-weight sampling is recovered for $\lambda = 1$, whereas varying λ changes the relative importance assigned to different effective-energy shells. Reducing λ weakens the suppression of configurations with larger x , while increasing λ favors configurations with smaller x . At $\lambda = 0$, all microscopic auxiliary-field configurations are assigned the same sampling weight. This does not produce a flat distribution in x , since effective-energy shells generally contain different numbers of configurations. In fact, using the sign-resolved entropy, the probability density of observing an effective energy x in sign sector σ becomes proportional to

$$P_{\lambda,\sigma}(x) \propto e^{S_\sigma(x) - \lambda x}. \quad (3.42)$$

The parameter λ therefore controls the competition between the entropy of an effective-energy shell and the sampling penalty assigned to its effective energy. In this sense, λ plays a role analogous to an inverse temperature conjugate to x . Furthermore, when the distribution is sufficiently sharply peaked, its most probable effective energy $x_{\lambda,\sigma}^\star$ approximately satisfies

$$\left. \frac{dS_\sigma}{dx} \right|_{x=x_{\lambda,\sigma}^\star} = \lambda. \quad (3.43)$$

Different values of λ therefore select regions of the entropy curve with different local slopes.

A collection of simulations at different λ values can consequently be used to explore a broader range of effective energies than the conventional $\lambda = 1$ ensemble.

The effect of varying λ is shown in Fig. 3.10. Each value of λ produces a distribution concentrated around a different region of effective-energy space. Increasing λ shifts the sampled configurations toward smaller values of $\bar{x}$, since these receive a larger relative weight. Conversely, decreasing λ weakens this preference and shifts the distribution toward larger effective energies. Negative values of λ reverse the usual bias and strongly suppress the configurations favored by conventional absolute-weight sampling.




Figure 3.10: Sign-resolved probability distributions of the intensive effective energy $\bar{x} = x/(\beta N)$ obtained using the generalized sampling weight $P_\lambda(h) = e^{-\lambda x(h)}$. The upper and lower panels correspond to configurations with positive and negative determinantal weights, respectively. The color indicates the value of λ . Note that the two panels use different vertical scales. The remaining simulation parameters are $N = 16$, $\beta = 8$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

The family of distributions therefore spans a much broader interval of $\bar{x}$ than the conventional $\lambda = 1$ ensemble alone, which confirms that the restricted range observed under ordinary DQMC sampling is not the complete range of accessible effective energies, but results from the particular exponential bias e^{-x} . The λ scans also provide a practical estimate of the interval over which the sign-resolved entropy should subsequently be determined.

Although varying λ allows different regions of effective-energy space to be explored, this procedure is not particularly practical for determining the entropy over a broad interval. For a fixed value of λ , the relation

$$S_\sigma(x) = \log P_{\lambda,\sigma}(x) + \lambda x + \text{const.} \quad (3.44)$$

provides a reliable estimate only within the limited region in which the corresponding distribution has sufficient statistics. Reconstructing the full entropy would therefore require a large number of simulations at different values of λ , with sufficiently overlapping distributions to determine the relative additive constants and combine the resulting estimates. As the distributions become narrower with increasing system size, the number of required values of λ would also increase, while the poorly sampled tails would remain difficult to resolve accurately.

A more efficient strategy is to replace the linear bias λx by a general sign-dependent function $\omega_\sigma(x)$ chosen to compensate the variation of the density of states across the entire interval of interest. More concretely, the “perfect” choice of such a function is the entropy itself, $\omega_\sigma(x) = S_\sigma(x)$, since in that case the probability of visiting a given x -shell with sign σ is

$$P_\sigma(x) = \text{const.} \quad (3.45)$$

across the entire effective-energy range. The corresponding generalized ensemble can then be designed to sample all effective-energy bins with approximately equal probability, rather than moving a narrow sampling window through configuration space one value of λ at a time. The required entropy is, however, precisely the unknown quantity to be determined.

The Wang-Landau procedure addresses this circularity by constructing the bias function $\omega_\sigma(x)$ iteratively during the simulation without any prior knowledge of the entropy [55, 56]. The Wang-Landau algorithm has been previously combined with DQMC to improve the exploration of a larger range of weight magnitudes and reduce autocorrelation times [34]. Here, we extend the same general strategy by resolving the sampling coordinate by the determinant sign, allowing the separate determination of the sign-resolved entropies.

In the present application, the effective-energy interval of interest, $[x_{\min}, x_{\max}]$, is first divided into a finite number of bins. With each bin and sign sector we associate an estimate

$\omega_\sigma(x)$ of the entropy and a visitation histogram $H_\sigma(x)$, with both being initially set to zero,

$$\omega_\sigma(x) = 0, \quad (3.46)$$

$$H_\sigma(x) = 0. \quad (3.47)$$

The simulation then performs the same local auxiliary-field proposals as conventional DQMC, but accepts them according to the current estimate of the generalized sampling weight through Eq. (3.40). The only exception is for proposals whose effective energy lies outside the predefined interval, in which case all such moves are rejected, thereby imposing reflecting boundaries at $x_{\min}$ and $x_{\max}$.

After every attempted update, regardless of whether the proposed move is accepted or rejected, the bin corresponding to the configuration occupied after the attempt is modified as follows. If the current state belongs to the bin (x, σ) , the updates are

$$\omega_\sigma(x) \mapsto \omega_\sigma(x) + f, \quad (3.48)$$

$$H_\sigma(x) \mapsto H_\sigma(x) + 1, \quad (3.49)$$

where $f > 0$ is the current modification factor. Since the sampling weight is proportional to $e^{-\omega_\sigma(x)}$, increasing $\omega_\sigma(x)$ decreases the probability of immediately returning to a bin that has already been visited, and so the random walk is consequently driven away from frequently sampled regions and toward bins that have been visited less often.

At the beginning of the simulation, the additive increment f is chosen to be unity, so it is sufficiently large that the bias can adapt rapidly. Typically, the modification factor is reduced as

$$f \mapsto \frac{f}{2} \quad (3.50)$$

once the visitation histogram is judged to be sufficiently flat. In our implementation, however, we chose the number of MC sweeps in each iteration to be proportional to the number of x bins and increased it exponentially from one iteration to the next. This choice thereby increases the likelihood of each x bin being visited and the simulation time spent on refining the increasingly smaller scales, which has been shown to improve the convergence of the algorithm [57]. After the modification factor is reduced, the histogram $H_\sigma(x)$ is reset to zero and the procedure is then repeated with the more slowly varying bias.

As f is successively reduced, the changes made to $\omega_\sigma(x)$ become smaller and the estimate converges toward the entropy,

$$\omega_\sigma(x) \longrightarrow S_\sigma(x) + C, \quad (3.51)$$

where C is an undetermined additive constant.

Because both sign sectors are explored by a single random walk within the same generalized ensemble and transitions between them are accepted using Eq. (3.40), their relative normalization is determined by the simulation. The remaining constant in Eq. (3.51) is therefore common to both $S_+(x)$ and $S_-(x)$ and it cancels from the entropy difference and from all normalized quantities derived from the two densities of states, but for graphical purposes we always subtract $C = \max_x S_+(x)$ from the sign-resolved entropies.

The iterative stage is terminated once the modification factor falls below a prescribed final value $f_{\min}$. At that point, $\omega_\sigma(x)$ provides an approximation to the bias required to produce a flat random walk over the joint (x, σ) space.

Figure 3.11 displays the evolution of the bias during the later Wang-Landau iterations, in which we see that the largest changes occur during the earliest iterations displayed. As the Wang-Landau modification factor is reduced, the successive curves approach one another over the complete sampled interval and only small corrections remain during the final iterations. Since the ideal flat-histogram bias satisfies

$$\omega_\sigma(x) = S_\sigma(x) + \text{const.}, \quad (3.52)$$

this progressive stabilization of $\omega_\sigma(x)$ provides a direct visual indication of the convergence of the entropy estimate. Because the entropy is defined only up to an additive constant, the curves are shifted so that the maximum value of $\omega_+(x)$ is zero.

To quantify this convergence, we compare the bias obtained at successive Wang-Landau iterations. Defining

$$\Delta\omega_\sigma^{(n)}(x) = \omega_\sigma^{(n)}(x) - \omega_\sigma^{(n-1)}(x), \quad (3.53)$$

we monitor the maximum change over the sampled effective-energy interval,

$$\max_x |\Delta\omega_\sigma^{(n)}(x)|. \quad (3.54)$$

This quantity measures the largest correction made to the intensive bias between two consecutive iterations, and is shown in Fig. 3.12.




Figure 3.11: Evolution of the sign-resolved Wang-Landau bias during the later iterations of the algorithm. The left and right panels show $\omega_+(x)/(\beta N)$ and $\omega_-(x)/(\beta N)$, respectively, as functions of the intensive effective energy $\bar{x} = x/(\beta N)$. The color indicates the Wang-Landau iteration.




Figure 3.12: Convergence of the sign-resolved Wang-Landau bias. The plotted quantity is the maximum absolute change in the intensive bias between consecutive Wang-Landau iterations, as defined in Eq. (3.53). Blue and red points correspond to the positive- and negative-sign sectors, respectively. The logarithmic vertical scale shows the progressive reduction of the corrections as the modification factor is decreased.

As shown in Fig. 3.12, the maximum change decreases by several orders of magnitude over the course of the simulation. The decrease is not strictly monotonic because the location of the largest correction may change between iterations and because the two sign sectors are visited with different frequencies. Nevertheless, both sectors display the same overall downward trend, reaching changes of order 10^{-6} in the intensive bias during the final iterations.

The agreement between the visual stabilization in Fig. 3.11 and the systematic decrease of $\max_x |\Delta \omega_\sigma^{(n)}(x)|$ indicates that the iterative bias has reached a stable shape over the chosen effective-energy interval. This convergence diagnostic alone does not determine the statistical uncertainty of the entropy estimate, nor does it guarantee that every bin has

been sampled equally well, but it does show that further Wang-Landau iterations produce progressively smaller modifications to the estimated bias.

The adaptive Wang-Landau stage produces a bias close to the entropy but not an exact estimate of it. We therefore hold the final bias fixed and perform a long production run. This separates the adaptive determination of the bias from the measurement stage and allows residual deviations from a perfectly flat histogram to be used to refine the entropy estimate, which is essential in obtaining more precise curves for the entropy difference $\Delta S(x)$. Denoting the production histogram by $H_\sigma^{\text{prod}}(x)$, the refined entropy is obtained from

$$S_\sigma(x) = \omega_\sigma(x) + \log H_\sigma^{\text{prod}}(x) + C. \quad (3.55)$$

As was the case with the iterated bias function $\omega_\sigma(x)$, the constant C is common to both sign sectors.

3.4. Entropy results and thermodynamic scaling

The Wang-Landau procedure described in the previous section provides the sign-resolved entropies over an interval of effective energies much broader than the region explored by conventional absolute-weight sampling. We now use these results to examine how the positive- and negative-sign sectors evolve with inverse temperature and system size.

All results are expressed in terms of the intensive effective energy

$$\bar{x} = \frac{x}{\beta N}, \quad (3.56)$$

defined in Sec. 3.1.1. The entropies are divided by the corresponding space-time volume βN and an additive constant common to both sign sectors is removed such that $\max_x S_+(x) = 0$ for each individual set of parameters. Consequently, the absolute vertical positions of entropy curves belonging to different simulations should not be compared, but their shapes, their scaling with βN , and the difference $S_+(x) - S_-(x)$ remain meaningful.

The results are presented in the order established by the formal relations in Sec. 3.2. We first examine the two extensive entropy curves, then isolate their difference and the associated local average sign, and finally restore the factor e^{-x} in order to recover the canonical distribution $P_\sigma(x) \propto e^{S_\sigma(x)-x}$, thus determining which part of effective-energy space is relevant to conventional DQMC sampling.

3.4.1 Dependence on inverse temperature

We first keep the lattice size and imaginary-time step fixed and vary the inverse temperature. Figure 3.13 shows the positive- and negative-sign entropies over the interval explored by the Wang-Landau simulations.




Figure 3.13: Sign-resolved entropies as functions of the intensive effective energy $\bar{x} = x/(\beta N)$ for several inverse temperatures. The left and right panels show $S_+(x)/(\beta N)$ and $S_-(x)/(\beta N)$, respectively. For each inverse temperature, the same additive constant is subtracted from both sign sectors so that the largest entropy is zero. No negative-sign entropy is shown for $\beta = 1$, since no negative-weight configurations were observed in the interval studied. The remaining simulation parameters are $N = 16$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

Both entropies are smooth, concave functions of the intensive effective energy and have almost identical overall shapes. The largest deviations from a common scaling form occur at the smallest inverse temperatures, particularly in the positive-sign sector at $\beta = 1$. For the remaining values of β , the curves increasingly approach a common form when plotted as $S_\sigma(x)/(\beta N)$ against $\bar{x}$.

This collapse is consistent with an exponential scaling of the sign-resolved density of states with the space-time volume,

$$N_\sigma(x) \sim e^{\beta N s_\sigma(\bar{x})}, \quad (3.57)$$

up to corrections that are subexponential in βN . Equivalently, the sign-resolved entropy is extensive,

$$S_\sigma(x) = \beta N s_\sigma(\bar{x}) + o(\beta N), \quad (3.58)$$

where $s_\sigma(\bar{x})$ is the entropy per unit space-time volume. The results therefore indicate that

the number of configurations in each sign sector grows exponentially with βN , while its dependence on the intensive effective energy is described by an approximately temperature-independent entropy density. We note that, since $\Delta\tau$ is held fixed throughout these comparisons, βN is proportional to the total number NL of auxiliary-field spins. Extensivity in βN is therefore equivalent here to extensivity in the size of the auxiliary-field space-time lattice.

The similarity between the two panels of Fig. 3.13 also shows that the leading extensive contributions to the positive- and negative-sign entropies are nearly identical. Their difference is too small to be resolved on the scale of the individual entropy curves and is shown directly in Fig. 3.14.




Figure 3.14: Difference between the positive- and negative-sign entropies, divided by the space-time volume, as a function of the intensive effective energy for several inverse temperatures. The dashed horizontal line marks equal sign-resolved entropies.

The entropy difference is positive over essentially the complete resolved interval, implying that the positive-sign sector contains more configurations than the negative-sign sector at the same effective energy. Its magnitude, however, varies strongly across configuration space. The difference is largest toward the boundaries of the interval and reaches a broad minimum in the central region around $\bar{x} \simeq -3$. Positive configurations therefore dominate strongly near the boundaries, while the multiplicities of the two signs become much more similar in the center.

The scaled difference

$$\frac{\Delta S(x)}{\beta N} = \frac{S_+(x) - S_-(x)}{\beta N} \quad (3.59)$$

decreases systematically as β is increased and approaches zero over an increasingly broad central interval. This shows that the positive- and negative-sign sectors approach the same entropy density in the low-temperature regime and that their leading exponential growth with the space-time volume is identical within this region.

The comparison between Figs. 3.13 and 3.14 reveals a clear separation of scales. The individual entropies are extensive in βN , whereas their difference is much smaller. The sign structure is consequently controlled by the residual imbalance between two sectors whose numbers of configurations are both exponentially large and whose leading entropy densities are nearly identical.

The vanishing of $\Delta S(x)/(\beta N)$ does not, by itself, imply that the unscaled difference $\Delta S(x)$ vanishes. Since the local average sign depends on the unscaled entropy difference,

$$\sigma(x) = \tanh\left(\frac{\Delta S(x)}{2}\right), \quad (3.60)$$

its behavior must be examined separately. The resulting curves are shown in Fig. 3.15.




Figure 3.15: Local average sign as a function of the intensive effective energy for several inverse temperatures. The curves are reconstructed from the sign-resolved entropy difference using Eq. (3.29).

At $\beta = 2$, the local average sign remains essentially equal to unity throughout the displayed interval, indicating that the negative-sign sector is still strongly suppressed. At $\beta = 4$, a broad minimum develops around $\bar{x} \simeq -3$. As the inverse temperature is increased further, the minimum rapidly approaches zero and the interval over which the two sign sectors occur with comparable multiplicities becomes broader.

For the largest inverse temperatures, $\sigma(x)$ is nearly zero over a substantial central region.

Within these effective-energy shells, positive- and negative-weight configurations therefore occur in almost equal numbers. In contrast, the local sign returns to unity toward both boundaries of the sampled interval. Negative-weight configurations are thus not present in a fixed proportion throughout configuration space. They become competitive mainly within an intermediate range of effective energies, while the outer shells remain dominated by positive configurations.

The local average sign describes the balance between the two sectors within a fixed effective-energy shell, but it does not determine how strongly that shell contributes to the absolute-weight ensemble. The latter is controlled by the sign-resolved canonical distributions

$$P_{\sigma}(x) \propto e^{S_{\sigma}(x)-x}. \quad (3.61)$$

The corresponding profiles reconstructed from the Wang-Landau entropies are shown in Fig. 3.16. Since the purpose of this figure is to compare their positions and widths, each curve is divided by its own maximum. Their relative heights and integrated areas should therefore not be interpreted as the relative probabilities of the two sign sectors.




Figure 3.16: Sign-resolved canonical profiles reconstructed from $e^{S_{\sigma}(x)-x}$ for several inverse temperatures. Solid and dashed lines correspond to the positive- and negative-sign sectors, respectively. Each curve is divided by its own maximum, so the figure compares only the positions and shapes of the distributions.

The canonical profiles occupy only a narrow portion of the full effective-energy interval over which the entropy was determined. At the smallest inverse temperatures, the distributions are comparatively broad and their maxima depend appreciably on β . As the temperature is lowered, they move toward a common region near $\bar{x} \simeq -3.2$ and become progressively

narrower. The positive- and negative-sign profiles also become nearly indistinguishable in both position and shape.

Comparison with Fig. 3.15 shows that the region selected by the canonical weight coincides with the minimum of the local average sign. At large β , conventional absolute-weight sampling is therefore concentrated precisely in the effective-energy shells where positive and negative configurations occur with almost equal multiplicity.

The deterioration of the global average sign with decreasing temperature can consequently be understood as the combination of the development of an extended central region in which $\sigma(x)$ is close to zero and the increasing concentration of the canonical distribution within that same region. The low-temperature sign problem is therefore associated not merely with the existence of negative-weight configurations, but with the fact that the configurations carrying most of the absolute statistical weight belong to shells in which the two signs are almost perfectly balanced.

3.4.2 Dependence on system size

We next keep the inverse temperature and imaginary-time step fixed and vary the number of lattice sites. Since β and $\Delta\tau$ are fixed, the number of time slices remains constant and increasing N directly enlarges the spatial volume and the total number NL of auxiliary-field variables. The resulting sign-resolved entropy densities for systems ranging from $N = 16$ to $N = 144$ are shown in Fig. 3.17.




Figure 3.17: Sign-resolved entropies as functions of the intensive effective energy $\bar{x} = x/(\beta N)$ for different lattice sizes. The left and right panels show $S_+(x)/(\beta N)$ and $S_-(x)/(\beta N)$, respectively. For each lattice size, the same additive constant is subtracted from both sign sectors so that the largest entropy is zero. The remaining simulation parameters are $\beta = 4$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$.

The shifted entropy densities display a clear collapse as the spatial volume is increased. The positive-sign curves are almost indistinguishable over the complete displayed interval, while the negative-sign sector exhibits somewhat larger finite-size corrections near the lower- $\bar{x}$ boundary. These deviations decrease with increasing N , and both sectors approach smooth limiting profiles.

At fixed β , this collapse establishes the extensivity of the shape of the entropy. It is consistent with the scaling form in Eq. (3.58), according to which the variation of $S_\sigma(x)$ across effective-energy space grows proportionally to N and is described by an approximately size-independent function of $\bar{x}$.

The positive- and negative-sign profiles remain too similar for their difference to be resolved on the scale of Fig. 3.17. Unlike the individual entropies, however, their difference is independent of the common additive normalization and can therefore be compared directly between lattice sizes. The results are shown in Fig. 3.18.




Figure 3.18: Difference between the positive- and negative-sign entropies, divided by the space-time volume, as a function of the intensive effective energy for different lattice sizes. The dashed horizontal line marks equal sign-resolved entropies.

The entropy difference retains a similar dependence on effective energy for all lattice sizes. It is largest near the boundaries of the sampled interval and reaches a broad minimum around $\bar{x} \simeq -3$. The positive-sign sector is therefore strongly favored in the outer shells, whereas the relative multiplicities of the two signs become much more similar in the central region.

The scaled difference $\Delta S(x)/(\beta N)$ decreases rapidly as N is increased. This behavior is consistent with the positive- and negative-sign sectors approaching the same leading

entropy density, while the distinction between them is carried by a much smaller residual contribution.

The effect of the unscaled entropy difference on the relative populations of the two sectors is seen more directly through the local average sign, shown in Fig. 3.19.




Figure 3.19: Local average sign as a function of the intensive effective energy for different lattice sizes. The curves are reconstructed from the sign-resolved entropy difference using Eq. (3.29).

For $N = 16$, the local average sign remains appreciably above zero throughout the sampled interval. As the lattice size is increased, the central minimum becomes deeper and extends over a broader range of effective energies. For the largest systems studied, the reconstructed local sign is close to zero around $\bar{x} \simeq -3$, consistent with positive- and negative-weight configurations occurring in nearly equal numbers in those shells.

Near the boundaries, $\sigma(x)$ remains close to unity for every lattice size and a local minimum is again formed in an intermediate effective-energy interval. The average sign around this minimum worsens as the lattice size is increased. The roughness visible in the $N = 144$ curve near its minimum reflects the sensitivity of $\sigma(x)$ to small fluctuations in $\Delta S(x)$ when the latter is close to zero.

To determine the statistical relevance of this central low-sign region, we reconstruct the canonical profiles proportional to $e^{S_\sigma(x)-x}$. The results are shown in Fig. 3.20. As in Fig. 3.16, each curve is divided by its own maximum, so the plotted profiles retain information about their positions and widths but not about the relative total weights of the two sign sectors.

The canonical profiles become progressively narrower as the lattice size is increased.




Figure 3.20: Peak-normalized sign-resolved canonical profiles reconstructed from $e^{S_\sigma(x)-x}$ for different lattice sizes. Solid and dashed lines correspond to the positive- and negative-sign sectors, respectively. Each curve is divided by its own maximum, so the figure compares only the positions and shapes of the profiles.

This behavior is consistent with the self-averaging of the intensive effective energy established in Sec. 3.1.1, where its variance was found to decrease approximately as $1/N$. The maxima also approach a common location near the center of the sampled interval.

The peak-normalized positive- and negative-sign profiles become increasingly similar in both position and width. Thus, within the canonically relevant region, the two sign sectors are associated with very similar distributions of determinant magnitude. The normalization used in the figure removes their relative integrated probabilities, so it does not by itself quantify the total population imbalance between the two sectors.

Comparison with Fig. 3.19 shows that the narrowing profiles are centered inside the region where the local average sign is smallest. Increasing the lattice size therefore does more than reduce the width of the sampled effective-energy distribution as it also concentrates the absolute-weight ensemble within shells whose positive- and negative-sign multiplicities are already nearly balanced.

The system-size dependence consequently complements the temperature dependence discussed above. Increasing N suppresses the entropy difference per unit volume, reduces the local sign in the canonically relevant region, and simultaneously narrows the absolute-weight distribution around that region. Through Eq. (3.31), these effects reinforce one another, since the effective-energy shells receiving most of the statistical weight are also those in which the cancellation between the two sign sectors is strongest.

4. Local geometry under single-spin excitations

The sign-resolved entropy analysis of the previous chapter determines the relative numbers of configurations in both sign sectors at a given effective energy. It does not, however, reveal how the determinant sign responds to small modifications of configurations drawn from that effective-energy shell. Two shells may have the same value of $\sigma(x)$ while differing substantially in the fraction of single-spin updates that change the sign.

To investigate this structure, we introduce a notion of local geometry or neighborhood based on the elementary updates used in DQMC. An auxiliary-field configuration $h = \{h_{\ell,i}\}$ can be viewed as a vertex of an NL -dimensional hypercube, since each of its NL components takes one of the values ± 1 . The complete configuration space therefore contains 2^{NL} vertices, with two vertices being connected whenever their configurations differ at exactly one space-time point, so the edges of the hypercube correspond precisely to the single-spin flips used in the standard local-update algorithm. Here, “single-spin” refers to one HS spin rather than to the physical spin of a fermion.

This graph structure allows us to ask how frequently an elementary displacement from a configuration crosses from one sign sector to the other. We first define quantities that characterize this geometric connectivity and then include the corresponding change in determinant magnitude in order to determine whether the sign-changing directions are dynamically accessible under the usual Metropolis rule.

4.1. Geometric and dynamical diagnostics

The elementary excitation considered in this chapter is the inversion of the auxiliary field at one space-time point (ℓ, i) ,

$$h_{\ell,i} \mapsto -h_{\ell,i}. \quad (4.1)$$

For a system with N spatial sites and L imaginary-time slices, each configuration has

$$n_{\text{exc}} = NL \quad (4.2)$$

such neighboring configurations. We label these excitations by $a = 1, \dots, n_{\text{exc}}$ and denote by h_a the configuration obtained by applying excitation a to h . The corresponding ratio of

determinantal weights is

$$R_a(h) = \frac{W(h_a)}{W(h)}. \quad (4.3)$$

Its magnitude determines the change in effective energy, while its sign determines whether the excitation changes the determinantal sign,

$$s(h_a) = s(h) \operatorname{sgn} R_a(h). \quad (4.4)$$

For a fixed configuration, we define

$$p(h) = \frac{1}{n_{\text{exc}}} \sum_{a=1}^{n_{\text{exc}}} \delta_{\operatorname{sgn} R_a(h), -1}. \quad (4.5)$$

The quantity $p(h)$ is the fraction of the edges incident on h which connect it to the opposite sign sector. Equivalently, it is the probability that an excitation selected uniformly from the n_{exc} possible single-spin flips changes the sign, before any Metropolis acceptance criterion is applied. In the hypercube representation, $n_{\text{exc}} p(h)$ is therefore the number of opposite-sign nearest neighbors of h .

The value of $p(h)$ generally varies between configurations belonging to the same sign-resolved effective-energy shell. We characterize these fluctuations through the distribution

$$\rho(p \mid x, \sigma) = \frac{1}{N_\sigma(x)} \sum_h \delta(x - x(h)) \delta_{\sigma, s(h)} \delta(p - p(h)), \quad (4.6)$$

where $N_\sigma(x)$ is the sign-resolved density of configurations introduced in the previous chapter. The distribution is normalized according to

$$\int dp \rho(p \mid x, \sigma) = 1. \quad (4.7)$$

It is important to distinguish $\rho(p \mid x, \sigma)$ from a sign-changing probability. The distribution describes how the fraction of sign-changing flips varies between configurations in the same shell. A narrow distribution indicates that most configurations have similar values of $p(h)$, while a broad distribution instead shows that configurations with the same effective energy and sign can respond very differently to a single-spin excitation, as some have few or no sign-changing flips, whereas others have many more. The shell is therefore heterogeneous, with some configurations lying much closer to the sign-sector boundary than others, even though all of its configurations share approximately the same effective energy and determinant sign.

The first moment of the $\rho(p \mid x, \sigma)$ distribution,

$$\langle p \rangle_{x,\sigma} = \int dp p \rho(p \mid x, \sigma), \quad (4.8)$$

is the probability that a uniformly chosen single-spin excitation changes the sign when its initial configuration is selected uniformly from the shell (x, σ) . This quantity is purely geometric as it counts every sign-changing direction independently of whether the corresponding proposal would be accepted in a DQMC simulation.

To include the change in determinant magnitude, we introduce the effective-energy displacement

$$\Delta x_a(h) \equiv x(h_a) - x(h) = -\log \left| \frac{W(h_a)}{W(h)} \right| = -\log |R_a(h)|. \quad (4.9)$$

An excitation with $\Delta x_a < 0$ increases the magnitude of the determinantal weight and is always accepted under absolute-weight sampling, whereas an excitation with $\Delta x_a > 0$ decreases the weight magnitude and is accepted with probability $e^{-\Delta x_a}$.

The displacement and final sign are resolved simultaneously through the joint distribution

$$\rho(\Delta x, \sigma' \mid x, \sigma) = \frac{1}{N_\sigma(x)} \sum_h \delta(x - x(h)) \delta_{\sigma, s(h)} \frac{1}{n_{\text{exc}}} \sum_{a=1}^{n_{\text{exc}}} \delta(\Delta x - \Delta x_a(h)) \delta_{\sigma', s(h_a)}. \quad (4.10)$$

This is the probability density that an excitation chosen uniformly from a configuration in the shell (x, σ) produces a displacement Δx and terminates in the sign sector σ' . Its normalization is joint over the displacement and final sign,

$$\sum_{\sigma'=\pm 1} \int_{-\infty}^{+\infty} d(\Delta x) \rho(\Delta x, \sigma' \mid x, \sigma) = 1. \quad (4.11)$$

The two final-sign sectors are therefore not normalized independently. Instead,

$$\int d(\Delta x) \rho(\Delta x, \sigma' \mid x, \sigma) \quad (4.12)$$

is the total probability that the proposal ends in sector σ' . The sign-changing component of the $\rho(\Delta x, \sigma' \mid x, \sigma)$ distribution therefore integrates to the mean sign-changing fraction

$$\langle p \rangle_{x,\sigma} = \int_{-\infty}^{+\infty} d(\Delta x) \rho(\Delta x, -\sigma \mid x, \sigma). \quad (4.13)$$

Thus, $\rho(\Delta x, \sigma' \mid x, \sigma)$ reproduces the mean sign-changing fraction after integration over the opposite-sign sector, while also determining whether the corresponding flips increase or decrease the determinant magnitude and how strongly they are suppressed by the Metropolis factor.

The probability that a uniformly selected single-spin excitation both changes the sign and is accepted by the absolute-weight Metropolis rule is then

$$P_{\sigma \rightarrow -\sigma}(x) = \int_{-\infty}^{+\infty} d(\Delta x) \rho(\Delta x, -\sigma | x, \sigma) \min\{1, e^{-\Delta x}\}. \quad (4.14)$$

We emphasize that this is an accepted-proposal probability conditioned on the initial shell, rather than a finite-time transition probability of the actual Markov chain. Comparing it with $\langle p \rangle_{x,\sigma}$ separates the geometric availability of sign-changing directions from their dynamical suppression by the Metropolis factor.

We also define the conditional acceptance probability

$$A_{\sigma}(x) = \frac{P_{\sigma \rightarrow -\sigma}(x)}{\langle p \rangle_{x,\sigma}}, \quad (4.15)$$

or equivalently,

$$A_{\sigma}(x) = \frac{\int d(\Delta x) \rho(\Delta x, -\sigma | x, \sigma) \min\{1, e^{-\Delta x}\}}{\int d(\Delta x) \rho(\Delta x, -\sigma | x, \sigma)}. \quad (4.16)$$

Whenever $\langle p \rangle_{x,\sigma} > 0$, this quantity answers the following question: given that the proposed excitation changes the sign, what is the probability that it is accepted? It isolates the dynamical penalty associated with the displacement Δx from the scarcity of sign-changing proposals themselves.

In practice, these distributions are measured using the converged entropic bias obtained in the previous chapter, with configurations being sampled from a frozen generalized ensemble of the form

$$\pi_{\omega}(h) \propto e^{-\omega_s(h)(x(h))}, \quad (4.17)$$

where $\omega_{\sigma}(x)$ is our previous approximation to the sign-resolved entropy. This bias permits a broad range of effective-energy shells to be explored uniformly.

For every stored configuration, all n_{exc} neighboring configurations are probed without modifying the Markov-chain state, i.e., without applying an accept-reject decision. Then, for each excitation, we calculate $R_a(h)$, the final sign, and $\Delta x_a(h)$. Each stored configuration contributes one value of $p(h)$ to $\rho(p | x, \sigma)$, whereas every individual excitation contributes to the joint histogram $\rho(\Delta x, \sigma' | x, \sigma)$. Recording all proposals, rather than only accepted moves, is essential, as restricting the analysis to accepted excitations would incorporate the Metropolis factor into the measured geometry and would bias both distributions.

4.2. Local sign-changing connectivity

For the purpose of probing the local geometry in a regime where the sign problem is severe, we consider the parameter set $N = 16$, $\beta = 8$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$, for which the local average sign $\sigma(x)$ becomes close to zero over a broad central interval of effective energy. The same parameter set is used throughout Chapters 4 and 5.

The distributions $\rho(p \mid x, \sigma)$ of the configuration-dependent sign-changing fraction are shown in Fig. 4.1. The left and right panels correspond to initially positive and initially negative configurations, respectively.

A pronounced asymmetry is visible outside the central low-sign region. At small values of $\bar{x} = x/(\beta N)$, where $\sigma(x)$ is close to unity, the distribution for positive configurations is concentrated extremely close to $p = 0$. Almost all nearest neighbors of a typical positive configuration therefore remain in the positive sector. Sign-changing directions from this sector are either absent within the available statistics or occur with an extremely small frequency.

Negative configurations in the same region, when they occur, display a broader distribution extending to larger values of p . They consequently possess more elementary directions leading back to the positive sector than positive configurations possess directions leading into the negative sector. Of course, this is consistent with negative configurations forming a comparatively small population within the good-sign shells and having a larger average number of opposite-sign nearest neighbors.

As the effective energy enters the region in which $\sigma(x)$ is small, both distributions broaden and extend toward larger values of p . Configurations of each sign therefore differ more strongly in the number of flips that change the determinant sign, while sign-changing flips become more common overall. At the same time, the positive- and negative-sector distributions become increasingly similar. Positive and negative configurations therefore have comparable probabilities of changing sign under a single-spin flip in this region, even though most individual flips still preserve the sign.

The broad distributions also show that the mean value $\langle p \rangle_{x, \sigma}$ does not fully describe the configurations within a shell. At fixed x and fixed sign, the distribution is concentrated near $p = 0$ but extends to substantially larger values. Thus, most configurations have very few single-spin flips that change the determinant sign, whereas a smaller subset has a noticeably larger number of such flips. The possibility of changing sign through a




Figure 4.1: Distribution $\rho(p | x, \sigma)$ of the configuration-dependent sign-changing fraction $p(h)$ for single-spin excitations. The left panel shows initially positive configurations and the right panel initially negative configurations. The dashed red curve shows the local average sign $\sigma(x)$ and is included to identify the central low-sign region. The insets show fixed- x cuts of the distributions at representative values of the intensive effective energy. The parameters used are $N = 16$, $\beta = 8$, $U = 6$, $\mu = -1.5$, and $\Delta\tau = 1/16$, corresponding to $L = 128$.

single-spin excitation is therefore not distributed uniformly among the configurations in the shell.

The corresponding first moments of $\rho(p | x, \sigma)$ are shown in Fig. 4.2.




Figure 4.2: Mean sign-changing fraction $\langle p \rangle_{x, \sigma}$ for single-spin excitations as a function of the intensive effective energy. The blue curve corresponds to positive-to-negative excitations and the red curve to negative-to-positive excitations. The dashed black curve shows the local average sign $\sigma(x)$.

The mean values summarize the asymmetry seen in the full distributions. In the positive-dominated region, the positive-to-negative fraction lies below the resolved range, while the negative-to-positive fraction remains finite. A negative configuration located in this region is therefore much more likely to possess a nearest neighbor in the positive sector than a positive configuration is to possess a nearest neighbor in the negative sector.

The two curves approach one another as $\sigma(x)$ decreases, reaching comparable values throughout much of the central low-sign interval. This approximate equality indicates a local symmetry between the two directions of sign change, which may be due to the balance between the sign-sector populations.

Nevertheless, sign-changing flips remain rare in absolute terms. Even around the minimum of $\sigma(x)$, only approximately 1% of the available single-spin excitations change the determinant sign. Thus, the presence of nearly equal numbers of positive and negative configurations does not imply that a typical local modification changes the sign very often. It means instead that the small sign-changing fractions are comparable for the two initial signs. The configuration space can therefore contain almost equal numbers of positive- and negative-weight configurations within the relevant sampled region while still having a sparse sign boundary under single-spin adjacency.

4.3. Effective-energy displacements and dynamical accessibility

The previous quantities describe the geometry of the sign sectors without considering whether the corresponding moves are favored when one applies the usual sampling procedure. The distributions resolved in both Δx and final sign add the additional information required to determine how accessible these directions are under the Metropolis rule. Figures 4.3 and 4.4 show the sign-conserving and sign-changing components, respectively.

Since the distributions in Figs. 4.3 and 4.4 are normalized jointly over the final sign, and since the sign-changing components have much smaller probability density, we verify again that most single-spin excitations preserve the determinant sign. Of course, this is in agreement with the small values of $\langle p \rangle_{x,\sigma}$. Furthermore, the positive-to-negative excitations are unresolved in most of the positive-dominated region, which is exactly the asymmetry found previously in $\rho(p | x, \sigma)$ and its first moment.

These distributions are also sharply concentrated around small values of Δx , although they possess extended tails toward both positive and negative displacements. A modification of one auxiliary-field spin therefore usually produces a small change in determinant magnitude, while rare excitations can modify the weight by several orders of magnitude. The distributions are also characterized by heavy tails toward positive effective-energy displacements, which correspond to moves that are rarely accepted.

It is relevant to note that the support of the sign-changing and sign-conserving components is quite similar, which indicates that sign-changing moves do not have a particularly large effective-energy displacement, and so they are not generally associated with prohibitively small acceptance probabilities. Instead, their rarity has its origin in the reduced number of possible moves that flip the determinant sign. The effect of the penalty introduced by the Metropolis factor is summarized by the accepted sign-changing probability shown in Fig. 4.5.

The accepted probabilities closely follow the geometric fractions in Fig. 4.2, with positive-to-negative moves remaining strongly suppressed outside the low-sign region, while negative-to-positive moves remain possible over a much broader range of effective energies. In the central region, the two accepted probabilities become nearly equal, which reflects the similar number of sign-changing directions and the similar Δx distributions of the two sectors.

It is important to note that the magnitude of the probabilities in Fig. 4.5 are smaller than




Figure 4.3: Sign-conserving component of the displacement distribution $\rho(\Delta x, \sigma' | x, \sigma)$ for single-spin excitations. The left panel shows positive-to-positive excitations and the right panel negative-to-negative excitations. The dashed red curve shows the local average sign $\sigma(x)$. The insets show fixed- x cuts at representative values of the intensive effective energy.




Figure 4.4: Sign-changing component of the displacement distribution $\rho(\Delta x, \sigma' | x, \sigma)$ for single-spin excitations. The left panel shows positive-to-negative excitations and the right panel negative-to-positive excitations. The dashed red curve shows the local average sign $\sigma(x)$. The insets show fixed- x cuts at representative values of the intensive effective energy.




Figure 4.5: Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for single-spin excitations. The blue and red curves correspond to positive-to-negative and negative-to-positive transitions, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

those in Fig. 4.2, which is to be expected due to the inclusion of the Metropolis factor. To better determine the magnitude of this penalty and confirm whether it is responsible for the low dynamical accessibility of sign-changing moves, we next analyze the conditional probability that such moves are accepted, represented in Fig. 4.6.




Figure 4.6: Conditional acceptance probability $A_{\sigma}(x) = P_{\sigma \rightarrow -\sigma}(x) / \langle p \rangle_{x,\sigma}$ for sign-changing single-spin excitations. The blue and red curves correspond to initially positive and initially negative configurations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

Figure 4.6 shows that, throughout most of the resolved interval, approximately 40–60% of the sign-changing proposals would be accepted by the absolute-weight Metropolis rule. This confirms our previous guess that sign-changing excitations are not generally subject to an overwhelming Metropolis suppression. Furthermore, the positive- and negative-sector curves are almost identical wherever both are resolved. This means that, once

a sign-changing excitation has been selected, its acceptance probability depends only weakly on the sign of the initial configuration.

Taken together, the results connect the sign-resolved entropy analysis of Chapter 3 with the response of individual configurations to local modifications. In the region where $\sigma(x)$ is small, positive- and negative-weight configurations occur in nearly equal numbers and respond similarly to single-spin excitations, that is, the fractions of sign-changing flips from the two sectors become comparable, and the corresponding distributions of Δx become similar. Nevertheless, fewer than approximately 1% of the available single-spin flips change the determinant sign. For the parameter set studied here, the accepted sign-changing probability is therefore small mainly because sign-changing proposals are uncommon, rather than because their acceptance probability is particularly low.

These conclusions concern single-spin excitations, and updates that modify several auxiliary-field variables simultaneously may change the determinant sign more frequently or produce different effective-energy displacements. This motivates the collective excitations studied in the following chapter, where their behavior is compared with the single-spin results established here.

5. Collective excitations

The previous chapter examined how the determinant sign responds to the smallest possible modification of the auxiliary field, that being the inversion of a single HS spin. In the central low-sign region, positive- and negative-weight configurations were found to respond similarly to these updates. Nevertheless, only of order 1% of the available single-spin flips changed the determinant sign. Once a sign-changing flip was selected, its acceptance probability was not particularly small, which indicates that the accepted sign-changing probability was limited mainly by the rarity of proposals that modify the sign.

This previous result motivates the study of moves that modify several auxiliary-field spins simultaneously, as such collective modifications may reach opposite-sign configurations with a higher probability than a single elementary flip. This response may also depend on how the modified spins are arranged over the space-time lattice. For example, flipping several consecutive spins along imaginary time need not have the same effect as flipping the same number of spins at unrelated time slices.

The diagnostics introduced in Chapter 4 are used throughout this chapter and they are evaluated using the same set of parameters. We begin by defining several families of collective excitations; we then determine the fraction of proposals that change the determinant sign, the probability that a proposal both changes the sign and is accepted, and the conditional acceptance probability once a sign-changing move has been selected. The single-spin results provide the reference against which the collective transformations are compared.

5.1. Collective excitation families

A general collective excitation is specified by a set $\mathcal{A}$ of points on the auxiliary-field space-time lattice. Its action on a configuration h is

$$h_{\ell,i} \mapsto -h_{\ell,i}, \quad (\ell, i) \in \mathcal{A}, \quad (5.1)$$

while all spins outside $\mathcal{A}$ remain unchanged. We denote the resulting configuration by $h_{\mathcal{A}}$. The ratio of determinantal weights is

$$R_{\mathcal{A}}(h) = \frac{W(h_{\mathcal{A}})}{W(h)}, \quad (5.2)$$

and the corresponding change in effective energy is

$$\Delta x_{\mathcal{A}}(h) = x(h_{\mathcal{A}}) - x(h) = -\log |R_{\mathcal{A}}(h)|. \quad (5.3)$$

The excitation changes the determinant sign whenever $R_{\mathcal{A}}(h) < 0$.

To characterize how accessible these collective moves would be under the usual DQMC dynamics, we associate each of them with the Metropolis acceptance probability it would have if proposed as a single move, that is, with all spins in $\mathcal{A}$ being inverted simultaneously (or one by one with no intermediate acceptance decision). The total ratio in Eq. (5.2) is then evaluated between the initial and final configurations, and, since these transformations are their own inverse and they are selected symmetrically within each excitation family, we obtain the acceptance probability

$$A(h \rightarrow h_{\mathcal{A}}) = \min\{1, e^{-\Delta x_{\mathcal{A}}}\}. \quad (5.4)$$

This is different from performing a sequence of ordinary single-spin updates. In the sequential process, every intermediate configuration is subjected to a separate acceptance decision and influences the following proposal. In a collective excitation, the intermediate configurations are not visited and only the total change between the two endpoints enters the acceptance probability.

For a given excitation family, let $\mathcal{A}_a$, with $a = 1, \dots, n_{\text{exc}}$, denote the possible transformations. The configuration-dependent sign-changing fraction is defined exactly as in Chapter 4,

$$p(h) = \frac{1}{n_{\text{exc}}} \sum_{a=1}^{n_{\text{exc}}} \delta_{\text{sgn } R_{\mathcal{A}_a}(h), -1}. \quad (5.5)$$

Only the set of allowed excitations and its size n_{exc} change between the different cases. The other quantities introduced in Chapter 4 are also defined exactly as before.

The first excitation considered is the inversion of a complete fixed-time spatial slice, i.e., for all spatial sites at a chosen imaginary-time slice ℓ_0 , we apply the map

$$h_{\ell_0, i} \mapsto -h_{\ell_0, i}, \quad i = 1, \dots, N. \quad (5.6)$$

Since any of the L slices may be selected, the number of distinct full-slice excitations is

$$n_{\text{exc}}^{\text{slice}} = L. \quad (5.7)$$

The complementary transformation is the inversion of a complete fixed-site worldline. Global updates of this form, involving a line of HS spins at a single spatial site extending over all imaginary-time slices, were previously introduced in DQMC to alleviate ergodicity problems at strong coupling [36]. Here, rather than using this transformation as part of the sampling dynamics, we employ it as a probe of the connectivity between the two sign sectors. This collective move then corresponds to flipping all auxiliary-field spins along the imaginary-time axis for a chosen spatial site i_0 ,

$$h_{\ell, i_0} \longrightarrow -h_{\ell, i_0}, \quad \ell = 1, \dots, L. \quad (5.8)$$

There is one such transformation for each lattice site, so

$$n_{\text{exc}}^{\text{worldline}} = N. \quad (5.9)$$

The full-slice and full-worldline excitations probe two qualitatively different organizations of the auxiliary-field modification. The former is extended in space and localized in imaginary time, whereas the latter is localized in space and extended in imaginary time. Their direct comparison does not, however, isolate the effect of this organization alone, since a full slice flips N spins, while a complete worldline flips L spins, and these numbers are generally different. Any difference between their sign-changing probabilities may therefore reflect both their orientation on the space-time lattice and the total number of modified spins.

To vary the temporal extent continuously while keeping the excitation localized at one spatial site, we introduce a contiguous fixed-site segment. Choosing a site i_0 , an initial time slice ℓ_0 , and a segment length M , we flip

$$h_{\ell, i_0} \longmapsto -h_{\ell, i_0}, \quad \ell = \ell_0, \dots, \ell_0 + M - 1 \pmod{L}, \quad (5.10)$$

where the periodic identification reflects the periodicity of the auxiliary field in imaginary time. For $1 \leq M < L$, there are N possible spatial sites and L possible starting slices, giving

$$n_{\text{exc}}^{\text{segment}}(M) = NL. \quad (5.11)$$

For $M = L$, the choice of starting slice becomes redundant and the excitation reduces to the complete fixed-site worldline, leaving only N distinct transformations.

The segment family interpolates between the two limits

$$M = 1 \implies \text{single-spin excitation,} \quad (5.12)$$

$$M = L \implies \text{complete fixed-site worldline.} \quad (5.13)$$

It therefore makes it possible to increase the number of consecutive time slices involved in the excitation while keeping it confined to a single spatial site.

Increasing M changes both the number of inverted spins and the length of the interval over which they are distributed. To determine whether temporal contiguity has an effect beyond the number of modified spins, we compare the contiguous segment with a random fixed-site excitation containing the same number M of flips. For a chosen site i_0 , let $\mathcal{L}_M$ be a set of M distinct imaginary-time slices,

$$\mathcal{L}_M \subset \{1, \dots, L\}, \quad |\mathcal{L}_M| = M. \quad (5.14)$$

The corresponding excitation is

$$h_{\ell, i_0} \longrightarrow -h_{\ell, i_0}, \quad \ell \in \mathcal{L}_M. \quad (5.15)$$

The contiguous and random excitations therefore modify the same number of auxiliary-field spins and remain confined to one spatial site. They differ only in their arrangement along imaginary time, as the contiguous excitation acts on M consecutive slices, whereas the random excitation acts on M slices selected without imposing any temporal ordering.

The complete family of random fixed-site excitations contains

$$n_{\text{exc}}^{\text{random}}(M) = N \binom{L}{M} \quad (5.16)$$

distinct transformations. This number grows rapidly with M , making it impractical to evaluate every possible random excitation for each stored configuration. As such, direct estimates must be obtained from a finite sample of the possible choices and consequently have substantially poorer statistics than the corresponding contiguous-segment results.

The excitation families introduced above are designed to separate the role of the number of flipped auxiliary-field spins from the role of their organization in imaginary time. By varying the segment length M , we can determine how the sign-changing properties evolve as the temporal extent of the excitation is increased, while the comparison with random fixed-site excitations at the same value of M keeps both the number of flips and their spatial support unchanged. Any remaining difference between the two types of excitations

can therefore be attributed to the temporal contiguity of the flipped spins.

As a final note, the collective determinant ratios entering our quantities of interest could be naively computed by just considering a series of single-spin flips, that is, after each such flip, one would compute the corresponding single-spin ratio, then update the Green function for the new configuration and so on. Of course, there are no free lunches, and this simple procedure has the disadvantage of updating the Green function once for each of the flipped spins, $N_{\mathcal{A}}$, which comes at a computational and memory cost of order $\mathcal{O}(N_{\mathcal{A}}N^2)$. If we compound this with the fact that we must test all n_{exc} excitations, then this inefficiency becomes clear. Instead, we made use of more efficient algorithmic approaches to computing these determinant ratios, which are present in Appendix B.

5.2. Spatial and temporal collective excitations

Having defined the different excitation families, we first compare the fixed-time spatial slice and fixed-site worldline transformations. As in Chapter 4, the configurations are sampled via the usual DQMC dynamics and the possible excitations are probed without modifying these configurations. For each stored configuration, all L time-slice excitations and all N worldline excitations are evaluated. Details of the efficient calculation of the corresponding determinant ratios are given in Appendix B.

In this chapter, we focus on the integrated quantities $\langle p \rangle_{x,\sigma}$, $P_{\sigma \rightarrow -\sigma}(x)$, and $A_{\sigma}(x)$, which are sufficient to compare how frequently the two excitations change the sign and how strongly the resulting proposals are suppressed by the Metropolis rule.

Figure 5.1 shows the mean fraction of excitations that change the determinant sign. The positive-to-negative and negative-to-positive fractions become increasingly similar as the local average sign approaches zero, following the behavior already observed for single-spin excitations. However, the two excitation families differ in the magnitude of these fractions, with the worldline excitations changing the determinant sign considerably more often. The difference is particularly clear in the low-sign region, where the worldline fraction reaches values of order $1/2$, while the spatial-slice fraction remains of order 10^{-1} . Outside the low-sign region, the probabilities depend strongly on the initial sign, with positive-to-negative excitations becoming quite suppressed, which reflects the dominance of positive configurations in those effective-energy shells. Thus, extending an excitation through the complete imaginary-time direction at one site produces, on average, many more sign

changes than reversing all spatial spins at one time slice. We emphasize that this comparison does not yet determine whether this difference originates from their organization in space-time.




Figure 5.1: Mean sign-changing fraction $\langle p \rangle_{x,\sigma}$ for fixed-time spatial-slice and fixed-site worldline excitations. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

The sign-changing fraction does not account for the change in determinant magnitude produced by the excitation. Figure 5.2 therefore shows the corresponding accepted sign-changing probabilities after inclusion of the Metropolis factor. In it we see that the acceptance penalty reduces the sign-changing probabilities substantially, which is due to the fact that these collective sign-changing moves lead to configurations with a larger effective energy. Nevertheless, the ordering observed in Fig. 5.1 is preserved, with complete worldlines retaining a larger accepted sign-changing probability than spatial slices, with a difference of a factor of several.




Figure 5.2: Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for fixed-time spatial-slice and fixed-site worldline excitations. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

The larger sign-changing fraction of the worldline excitation is therefore not canceled by the larger determinant-magnitude changes associated with it. Complete worldlines not only produce sign changes more frequently, but enough of these proposals remain acceptable for their overall accepted probability to remain substantially larger.

To properly determine the origin of this difference, Fig. 5.3 shows the conditional acceptance probability $A_\sigma(x)$. This figure shows that, whenever any of these excitation families produces a sign-changing move, the corresponding acceptance probability is quite similar, with the fixed-time spatial slice having a slight advantage. The larger value of $P_{\sigma \rightarrow -\sigma}(x)$ for worldline excitations therefore does not result from their sign-changing proposals being more easily accepted, but instead arises primarily because they change the determinant sign much more frequently.




Figure 5.3: Conditional acceptance probability $A_\sigma(x)$ for sign-changing fixed-time spatial-slice and fixed-site worldline excitations. The left and right panels correspond to initially positive and initially negative configurations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

The conditional acceptance is also substantially smaller than for the single-spin excitations studied in Chapter 4. While approximately 40–50% of the single-spin sign-changing proposals were accepted through most of the resolved interval, the collective excitations typically have conditional acceptances of approximately 10% across much of the most sampled region. Modifying a larger part of the auxiliary field therefore produces a trade-off, as sign changes become more frequent, but a smaller fraction of them survive the Metropolis acceptance step. The accepted probabilities in Fig. 5.2, however, continue to depend on both this acceptance and the number of available sign-changing excitations. The three figures must therefore be considered together, with $\langle p \rangle_{x,\sigma}$ determining how frequently the sign changes, while $A_\sigma(x)$ determines what fraction of those changes is accepted.

For the parameters studied here, complete fixed-site worldlines are consequently more

effective than fixed-time spatial slices at producing accepted sign changes. As noted previously, this comparison alone cannot determine whether this enhancement results from the organization of the excitation itself, due to the difference in the number of spins flipped. We next address this distinction using contiguous fixed-site segments of variable length M , as these moves allow the response to be followed continuously between the single-spin and complete-worldline limits.

5.3. Dependence on the segment length

The comparison in the previous section showed that complete fixed-site worldlines produce sign changes much more often than fixed-time spatial slices. Since these transformations modify different numbers of auxiliary-field spins, that comparison does not really reveal how the sign-change frequency develops as the temporal extent of the fixed-site excitation is increased. We therefore consider contiguous fixed-site segments of variable length M , which allow us to interpolate between the single-spin limit at $M = 1$ and the complete-worldline limit at $M = L$.

As before, we focus on the integrated quantities $\langle p \rangle_{x,\sigma}$, $P_{\sigma \rightarrow -\sigma}(x)$, and $A_\sigma(x)$, which directly show how the sign-changing and acceptance probabilities evolve with the segment length.

Figure 5.4 shows the mean sign-changing fraction for segment lengths ranging from $M = 1$ to $M = L = 128$, in which the shortest segment reproduces the single-spin result of Chapter 4, while the longest reproduces the fixed-site worldline result of Sec. 5.2.




Figure 5.4: Mean sign-changing fraction $\langle p \rangle_{x,\sigma}$ for contiguous fixed-site segments of different lengths M . The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$. The cases $M = 1$ and $M = 128$ correspond to the single-spin and complete-worldline limits, respectively.

Increasing the segment length produces a strong enhancement in the average probability of the excitation changing the sign. In the relevant low-sign region, the single-spin value is

of order 10^{-2} , whereas the corresponding probability for segments extending over several consecutive time slices eventually approaches one half. In that limit, sign-changing and sign-preserving segment excitations become equally frequent.

Moreover, just as before, the asymmetry between the two initial signs remains visible outside the low-sign region, with negative-to-positive excitations being more common. In either case, increasing M typically extends the range over which positive-to-negative sign changes are observed and, within the low-sign region, the two directions become similar independently of the segment length.

It is also relevant to note the non-monotonicity in the segment length as one approaches the full worldline limit, which is particularly visible outside the low-sign region. This could already indicate an advantage of intermediate-length segments before even accounting for the larger Metropolis penalty naturally associated with larger excitations. To better analyze this we move on to the accepted sign-changing probabilities of Fig. 5.5.




Figure 5.5: Accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$ for contiguous fixed-site segments of different lengths M . The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

For the shortest segments, the accepted sign-changing probability increases rapidly with M , which is particularly clear when comparing the single-spin result with segments of length $M = 8, 16$, or 32 . Longer segments are therefore able to produce accepted sign changes substantially more frequently than the elementary update but, unlike the mean sign-changing fraction, the accepted probability does not continue to increase systematically up to the complete-worldline limit. Within much of the low-sign region, the largest values are obtained for intermediate segments, and so the effectiveness of an excitation cannot be inferred from its sign-changing fraction alone.

To better analyze this dependence on the segment length, we now fix the initial effective energy near the peak of the canonical distribution, which for this set of parameters takes the value $\bar{x}_{\text{peak}} = -3.23$, and evaluate the three diagnostics as functions of M . The resulting behavior is shown in Fig. 5.6.




Figure 5.6: Dependence of the sign-changing response on the length M of a contiguous fixed-site segment at $\bar{x}_{\text{peak}} = -3.23$. The upper panel shows the mean sign-changing fraction $\langle p \rangle_{x, \sigma}$, the middle panel the accepted sign-changing probability $P_{\sigma \rightarrow -\sigma}(x)$, and the lower panel the conditional acceptance probability $A_{\sigma}(x)$. Blue and red curves correspond to initially positive and initially negative configurations, respectively. The horizontal dashed lines indicate the corresponding values for the fixed-time spatial-slice excitation.

The three quantities display clearly the competition already suggested by Figs. 5.4 and 5.5. The mean sign-changing fraction increases strongly with the segment length and approaches values of order one half for the longest segments, meaning that these collective modifications encounter the opposite sign sector increasingly often. At the same time, the conditional acceptance probability decreases steadily with M , falling from the single-spin value of roughly 40–50% to only a few percent for the longer segments. The accepted sign-changing probability therefore results from a balance between these two effects and reaches its largest values at an intermediate segment length, around $M \simeq 32$ for the parameters considered here, before decreasing again.

A particularly useful comparison is obtained at $M = 16 = N$. A contiguous segment of this length and a complete fixed-time spatial slice both modify exactly 16 auxiliary-field spins for our parameter choice, yet their sign-changing properties are markedly different, with the contiguous segment producing sign changes more easily. The number of flipped spins alone is therefore clearly insufficient to characterize the effect that a collective transformation has on the determinant sign, as their arrangement in space and imaginary time also plays an important role. To further investigate this dependence, we next compare contiguous segments with random collective excitations containing the same number of flipped spins.

5.4. Contiguous and random fixed-site excitations

The results of the previous section show that the sign-changing properties of collective excitations are not determined solely by the number of auxiliary-field spins that are modified. In particular, we saw that a contiguous fixed-site segment produces sign changes more easily than a complete fixed-time spatial slice of the same size. To better isolate the role that temporal contiguity might have, we now compare the contiguous fixed-site segments with the random fixed-site excitations introduced in Sec. 5.1. Both collective moves share the same spatial support and the number of flipped spins but differ in their arrangement along imaginary time.

The comparison is carried out in three stages. We first examine direct measurements obtained by explicitly sampling random fixed-site excitations. Because the number of possible random choices grows combinatorially with M , these measurements become statistically noisy when only a finite subset can be evaluated. We then construct a smoother approximation from the accurately measured single-spin displacement distributions, which is verified to somewhat follow the noisier curves of the random fixed-site excitations, particularly for smaller excitation sizes, and to follow quite closely the curves for fully random excitations. Finally, this reconstructed random baseline is compared with the contiguous-segment results in order to determine whether temporal contiguity produces an additional enhancement beyond the effect of increasing the number of flipped spins.

5.4.1 Direct random fixed-site excitations

For a fixed segment length M , the complete random excitation family contains

$$n_{\text{exc}}^{\text{random}}(M) = N \binom{L}{M} \quad (5.17)$$

possible transformations. This number becomes too large for all of them to be evaluated for every stored configuration, even for moderate values of M . We therefore sample a finite number n_{samp} of random excitations uniformly from the complete family.

For a configuration h , the directly sampled estimate of the sign-changing fraction is

$$\hat{p}_M(h) = \frac{1}{n_{\text{samp}}} \sum_{a=1}^{n_{\text{samp}}} \delta_{\text{sgn } R_{\mathcal{A}_a}(h), -1}, \quad (5.18)$$

where each $\mathcal{A}_a$ contains M distinct time slices at a fixed spatial site. When the excitations are sampled uniformly, Eq. (5.18) provides an unbiased estimate of the fraction over the complete random family. Its statistical fluctuations nevertheless increase when sign-changing excitations are rare or when the number of sampled transformations is small compared with the total number of possible choices.

Each set of M flips is treated as one collective proposal. The determinant ratio is evaluated between the initial configuration and the configuration obtained after all M spins have been inverted through successive single-spin flips, without applying acceptance decisions to the intermediate configurations. Thus, no efficient single determinant calculation is implemented in this case, in contrast to the other excitation types, which further complicates the measurements of the quantities of interest.

Figure 5.7 shows the resulting mean sign-changing fractions for the random fixed-site excitations. The case $M = 1$ reproduces the single-spin result, for which every possible excitation can be evaluated. For $M > 1$, only a finite sample of the much larger random family is used.




Figure 5.7: Mean sign-changing fraction $\langle p \rangle_{x,\sigma}$ obtained by directly sampling random fixed-site excitations containing M flipped auxiliary-field spins. The left and right panels show positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$. The case $M = 1$ corresponds to the single-spin limit.

The direct measurements show that the probability of changing the determinant sign increases strongly with the number of randomly selected time slices. In the low-sign region, the single-spin probability is of order 10^{-2} , while the largest values of M considered produce sign-changing fractions of order 10^{-1} and approach one half, which is a behavior similar to that seen for the contiguous segments. Randomly modifying a larger portion of a fixed-site worldline therefore makes a sign reversal progressively more likely.

The curves for $M > 1$ display visible bin-to-bin fluctuations, particularly in regions where the underlying sign-changing probability is small. These fluctuations arise both from the rarity of sign-changing moves and from the finite sampling of the random excitation family.

Despite this noise, the direct results establish the main qualitative behavior. Increasing M also increases the sign-changing fraction when the flipped spins are distributed randomly along the same worldline. The growth observed for contiguous segments therefore cannot by itself be attributed to temporal contiguity, since part of it follows simply from modifying a larger number of auxiliary-field spins.

5.4.2 Fixed-shell reconstruction from single-spin excitations

The direct random-excitation measurements provide the desired sign-changing probabilities without additional approximations, but their statistical precision is limited by the number of random subsets that can be sampled for each configuration. A smoother reference can be constructed from the much more accurately determined single-spin distributions of Chapter 4. The basic idea is to approximate a random M -spin excitation as the combination of M independent single-spin contributions.

Let

$$\rho_1(\Delta x, \sigma' \mid x, \sigma) \tag{5.19}$$

denote the joint distribution of the effective-energy displacement and final sign produced by a single-spin excitation. Starting from a configuration in the shell (x, σ_0) , we combine successive draws from this distribution while updating the sign after every contribution. The effective-energy shell used to select the elementary distribution is, however, kept fixed at its initial value x . For this reason, we refer to the construction as a fixed-shell convolution.

If Δx_j denotes the displacement associated with the j th spin flip, the total displacement after M draws is

$$\Delta X_M = \sum_{j=1}^M \Delta x_j. \quad (5.20)$$

We denote the resulting distribution by

$$\rho_M^{\text{conv}}(\Delta X, \sigma_M | x, \sigma_0). \quad (5.21)$$

Starting from the measured single-spin distribution,

$$\rho_1^{\text{conv}}(\Delta X, \sigma_1 | x, \sigma_0) = \rho_1(\Delta X, \sigma_1 | x, \sigma_0), \quad (5.22)$$

the distribution for $m + 1$ contributions can be obtained recursively as

$$\begin{aligned} \rho_{m+1}^{\text{conv}}(\Delta X, \sigma_{m+1} | x, \sigma_0) &= \sum_{\sigma_m = \pm 1} \int d(\Delta Y) \rho_m^{\text{conv}}(\Delta Y, \sigma_m | x, \sigma_0) \\ &\quad \times \rho_1(\Delta X - \Delta Y, \sigma_{m+1} | x, \sigma_m). \end{aligned} \quad (5.23)$$

More concretely, the full distribution $\rho_M^{\text{conv}}(\Delta X, \sigma_M | x, \sigma_0)$ can be written as a convolution of the single-spin distributions,

$$\rho_M^{\text{conv}}(\Delta X, \sigma_M | x, \sigma_0) = \sum_{\sigma_1, \dots, \sigma_{M-1}} \int \prod_{j=1}^M d(\Delta x_j) \prod_{j=1}^M \rho_1(\Delta x_j, \sigma_j | x, \sigma_{j-1}) \delta \left(\Delta X - \sum_{j=1}^M \Delta x_j \right). \quad (5.24)$$

The reconstructed sign-changing fraction is then obtained by integrating over the final sector with the opposite sign,

$$\langle p \rangle_{x, \sigma; M}^{\text{conv}} = \int d(\Delta X) \rho_M^{\text{conv}}(\Delta X, -\sigma | x, \sigma). \quad (5.25)$$

No Metropolis decision is applied between the individual contributions in this construction and so the M draws are combined into a single excitation, so that the intermediate signs only determine which single-spin distribution is used for the following contribution. This distinguishes the reconstruction from an actual sequence of M ordinary DQMC updates, which will be considered separately in Sec. 5.5.

The construction makes two important approximations. First, the elementary contributions are treated as independent, as the virtual intermediate configurations are uncorrelated. Second, their displacement distributions are always evaluated at the initial effective energy x , rather than at the value reached after the preceding contributions, which is justified

by the relatively small support of the $\rho(\Delta x, \sigma' | x, \sigma)$ distributions of Chapter 4. In addition, the construction contains no constraint forcing the M contributions to belong to the same spatial site or to form any particular arrangement in imaginary time. It should therefore be regarded as an independent-spin baseline whose relation to the different random excitation families must be tested numerically rather than as an exact reconstruction.

Figure 5.8 compares this approximation with the direct random-excitation measurements of the previous subsection. Solid curves show the directly sampled results, while colored dashed curves show the fixed-shell reconstruction.




Figure 5.8: Comparison between the directly sampled and reconstructed mean sign-changing fractions for random fixed-site excitations. Solid curves show direct measurements obtained by sampling random sets of M distinct time slices at one spatial site, while colored dashed curves show the fixed-shell reconstruction obtained from the single-spin displacement distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

For the shortest excitations, the reconstructed curves closely follow the directly sampled results. The agreement remains good throughout much of the low-sign region for intermediate values of M , reproducing both the increase of the sign-changing fraction with M and the tendency of the positive-to-negative and negative-to-positive probabilities to become comparable. The reconstruction also removes the strong bin-to-bin fluctuations present in the direct measurements where sign-changing events are rare.

The agreement becomes progressively less accurate for the longest excitations. The largest deviations are visible for $M = 64$, for which the directly sampled sign-changing fractions can be substantially larger than the reconstructed values, particularly away from the center of the low-sign region. Smaller differences are already visible for $M = 16$. The fixed-shell construction therefore captures the general dependence on x and M , but does not reproduce all of the structure that develops when a large fraction of a worldline is modified simultaneously.

A plausible interpretation of this discrepancy is that it already reflects the effect of temporal contiguity within the direct random excitations. The reconstruction combines independent single-spin contributions and therefore resembles a set of flips that are distributed without any temporal structure. The directly sampled random excitations, however, are restricted to a single spatial site. When M becomes a sizable fraction of L , a random subset of time slices on one worldline inevitably contains many neighboring slices and short contiguous segments. In particular, for $M = 64$ and $L = 128$, half of the slices of the worldline are flipped, so substantial local contiguity becomes unavoidable even though no consecutive block is imposed explicitly. The direct random excitations may therefore already contain some of the temporal structure that is absent from the fixed-shell reconstruction. To verify this hypothesis, we repeat the direct measurement using fully random excitations, in which the M flipped spins are selected uniformly from the complete space-time lattice.

For a fully random excitation, the set of flipped variables is chosen from all NL space-time points,

$$\mathcal{L}_M \subset \{(\ell, i) : \ell = 1, \dots, L, i = 1, \dots, N\}, \quad |\mathcal{L}_M| = M. \quad (5.26)$$

The complete family therefore contains

$$n_{\text{exc}}^{\text{full rand}}(M) = \binom{NL}{M} \quad (5.27)$$

possible transformations. As for the random fixed-site excitations, only a finite subset is sampled directly for each stored configuration.

Figure 5.9 compares these direct measurements with the same fixed-shell reconstruction used above. Solid curves show the finitely sampled fully random excitations, while dashed curves show the reconstruction from the single-spin distributions.

In contrast with the random fixed-site excitations, the fully random measurements follow the reconstructed curves closely over the complete range of M considered. The agreement remains very good even for $M = 32$ and $M = 64$, for which clear systematic differences were observed when the flipped spins were restricted to a single worldline. Instead, the comparison indicates that the fixed-shell construction provides a good approximation to excitations composed of spins distributed randomly over the complete space-time lattice. The discrepancy observed for random fixed-site excitations at large M is consequently associated with the restriction of all flipped variables to a common worldline. When M/L becomes appreciable, this restriction necessarily produces neighboring flipped time




Figure 5.9: Comparison between the directly sampled fully random excitations and the fixed-shell reconstruction. Solid curves show direct measurements obtained by selecting M distinct auxiliary-field spins uniformly from the complete space-time lattice, while dashed curves show the reconstruction obtained from the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. Colors indicate the number M of flipped spins.

slices and short contiguous groups even though no contiguous segment is imposed explicitly.

The two direct comparisons therefore reveal a progression in the organization of the excitation. Fully random spins are well described by the independent-spin reconstruction, while restricting the same number of flips to one worldline produces systematic deviations as temporal contiguity becomes important. The contiguous segments studied in Sec. 5.3 represent the limiting case in which this temporal organization is imposed explicitly. Nevertheless, this fixed-shell convolution remains a good approximation for the random fixed-site excitations.

5.4.3 Effect of temporal contiguity

The previous comparisons provide a reference against which the effect of temporal contiguity can be examined, as the fixed-shell convolution closely reproduces excitations in which the M flipped spins are distributed randomly over the complete space-time lattice. Random fixed-site excitations also remain close to this baseline for small and intermediate values of M , particularly from $M = 1$ up to $M = 16$, indicating that restricting the flipped spins to a common worldline has only a minor effect while the selected time slices remain sparse. For this reason, we now restrict ourselves to this range of segment lengths and compare the contiguous segments of Sec. 5.3 with the fixed-shell convolution to determine if the differences between the contiguous and reconstructed curves can be attributed primarily to the imposed temporal organization of the former.

We first consider the mean sign-changing fraction. Figure 5.10 compares the contiguous-segment results with the fixed-shell reconstruction for equal values of M . Solid curves show the contiguous excitations, while dashed curves show the reconstructed random baseline.




Figure 5.10: Comparison between the mean sign-changing fractions of contiguous fixed-site segments and the fixed-shell random baseline. Solid curves show contiguous segments, while dashed curves show the reconstructed random fixed-site excitations obtained from the fixed-shell convolution of the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

The contiguous segments produce a systematically larger sign-changing fraction than the random baseline over most of the low-sign region. The difference is relatively small for $M = 2$ and $M = 4$, but becomes increasingly pronounced for $M = 8$ and especially $M = 16$. This behavior is observed for both sign-changing directions. Thus, the advantage of the contiguous excitation develops progressively as a longer interval of imaginary time is modified coherently and this enhancement is therefore not simply a consequence of modifying more auxiliary-field variables, since M is held fixed in each comparison.

We note that the effect is not uniform over the complete effective-energy interval. Toward the upper boundary, the contiguous and reconstructed curves approach one another and can change their relative ordering. Nevertheless, the clearest enhancement is found in the central low-sign region, which is also the region sampled in a usual DQMC simulation.

A relative enhancement of the same magnitude is observed even when accounting for the acceptance-decision penalty, as can be seen in the accepted sign-changing probabilities of Fig. 5.11 for the same set of segment lengths.

The comparison of Figs. 5.10 and 5.11 also shows that the advantage of temporal contiguity grows progressively with the length of the excitation. For the shortest segments, imposing contiguity produces only a modest change relative to randomly distributed flips.




Figure 5.11: Comparison between the accepted sign-changing probabilities of contiguous fixed-site segments and the fixed-shell random baseline. Solid curves show contiguous segments, while dashed curves show the reconstructed random fixed-site excitations obtained from the fixed-shell convolution of the single-spin distributions. The left and right panels correspond to positive-to-negative and negative-to-positive excitations, respectively. The dashed black curve shows the local average sign $\sigma(x)$.

For longer segments, the difference becomes much more pronounced, indicating that the determinant responds to modifications extending over a finite interval of imaginary time rather than only to the number of auxiliary-field spins that are inverted.

The results of this section therefore show that collective excitations are sensitive not only to the number of modified auxiliary-field variables but also to their arrangement along imaginary time. Contiguous fixed-site segments provide a more effective route between the two sign sectors than a comparable set of randomly distributed flips, and this advantage remains after the Metropolis acceptance criterion is applied.

5.5. Collective excitations and local-update dynamics

The comparisons above concern different collective proposals evaluated from the same initial configuration. They do not yet determine whether a collective segment is more effective at changing the determinant sign than the standard DQMC dynamics itself. In the usual Markov chain, the same number of single-spin proposals is performed sequentially, with an acceptance decision after every step and with each subsequent proposal acting on the resulting intermediate configuration. As such, the auxiliary field evolves through a correlated sequence of single-spin proposals and independent acceptance decisions. The configuration reached after several such updates therefore depends on all accepted and rejected proposals along the intermediate trajectory and intermediate configurations may also contribute to the sampling.

To characterize this process directly, we measure sign-transition probabilities at a finite lag in the ordinary DQMC Markov chain. The idea is to measure the probability that a

configuration which, at some MC time, belongs to a given sign-resolved x -shell (x, σ) has sign $-\sigma$ after a fixed number M of elementary proposals. With that in mind, we let h_t denote the configuration after the t th elementary Metropolis proposal, including rejected proposals, and define

$$x_t = x(h_t), \quad \sigma_t = s(h_t). \quad (5.28)$$

One unit of MC time in this section therefore corresponds to one single-spin proposal rather than one complete sweep. For instance, if $M = 2$, the corresponding transition contains all possible two-proposal histories, including accept-accept, accept-reject, reject-accept, and reject-reject events. This measured quantity is therefore a property of the Markov chain, not an artificial process in which every proposed spin flip is forced to be accepted.

For a lag M , we denote this probability of inverting the sign of a configuration initially in the shell (x, σ) after M Markov-chain steps by

$$P_M(-\sigma | x, \sigma) = \Pr[\sigma_{t+M} = -\sigma | x_t \in B_x, \sigma_t = \sigma], \quad (5.29)$$

where B_x denotes the effective-energy bin containing x . We note that the conditioning refers only to the configuration at the initial time t . During the following M proposals, both the sign and the effective energy are free to evolve according to the usual DQMC dynamics. In practice, this probability can be estimated directly from the Markov-chain time series. For every configuration h_t whose effective energy lies in B_x and whose sign is σ , we inspect the configuration M elementary proposals later and record whether its sign is $-\sigma$. We therefore define the number of eligible initial configurations in a Markov chain with T proposals as

$$H_M(x, \sigma) = \sum_{t=M}^T \mathbf{1}[x_{t-M} \in B_x, \sigma_{t-M} = \sigma], \quad (5.30)$$

where $\mathbf{1}[\dots]$ is the characteristic function, and the number of those configurations that have the opposite sign after M proposals as

$$H_M(-\sigma | x, \sigma) = \sum_{t=M}^T \mathbf{1}[x_{t-M} \in B_x, \sigma_{t-M} = \sigma, \sigma_t = -\sigma]. \quad (5.31)$$

The measured transition probability is then

$$\hat{P}_M(-\sigma | x, \sigma) = \frac{H_M(-\sigma | x, \sigma)}{H_M(x, \sigma)}. \quad (5.32)$$

This procedure uses the actual sequence generated by the DQMC simulation. In particular, rejected proposals are retained in the time series and therefore contribute naturally to the estimator. No assumptions are made about the intermediate configurations between $t - M$ and t , and as such the chain may change sign several times, move between different effective-energy shells, or remain unchanged during rejected proposals. Only the sign of the configuration after exactly M elementary proposals is recorded. Naturally, in the infinite-sample limit, the measured probability converges to the desired probability,

$$\lim_{T \rightarrow \infty} \hat{P}_M(-\sigma | x, \sigma) = P_M(-\sigma | x, \sigma). \quad (5.33)$$

It is relevant to note that, since the conditioning in Eq. (5.32) applies only to the initial configuration, as M increases, the chain can move progressively farther from its initial state and eventually lose memory of both x_t and σ_t . Assuming an equilibrated ergodic chain,

$$\lim_{M \rightarrow \infty} P_M(-\sigma | x, \sigma) = P_{-\sigma}, \quad (5.34)$$

where P_σ is the global probability of a configuration having sign σ under absolute-weight sampling. Since

$$\langle s \rangle_{|W|} = P_+ - P_-, \quad (5.35)$$

the probability P_σ is given by

$$P_\sigma = \frac{1 + \sigma \langle s \rangle_{|W|}}{2}. \quad (5.36)$$

Consequently,

$$\lim_{M \rightarrow \infty} P_M(-\sigma | x, \sigma) = \frac{1 - \sigma \langle s \rangle_{|W|}}{2}. \quad (5.37)$$

This limit is independent of the initial effective energy and approaches 1/2 if the average sign is reduced, which is the case for the present set of parameters.

Figure 5.12 shows the measured transition probabilities for lags ranging from $M = 1$ to $M = 2^{16}$, the latter corresponding to 16 MC sweeps.

The probability of finding the opposite sign increases systematically with the lag. For $M = 1$, sign changes remain rare, in agreement with the single-spin results of Chapter 4. As more local proposals are allowed, the chain has a greater probability of reaching the




Figure 5.12: Lagged sign-transition probability $P_M(-\sigma | x, \sigma)$ measured from the ordinary single-spin DQMC Markov chain. The left and right panels correspond to positive-to-negative and negative-to-positive transitions, respectively. The colorbar indicates the number M of elementary Metropolis proposals separating the initial and final configurations. The dashed red curve shows the large- M limit of Eq. (5.37) and the dashed black curve shows the canonical distribution of Eq. (3.32).

opposite sign sector, and the transition probability rises throughout the sampled effective-energy interval. Eventually, the curves flatten and converge to a value near $1/2$, in agreement with the prediction given by (5.37), since the parameter choice here is such that $\langle s \rangle_{|W|}$ is very close to zero.

We now use these probabilities to compare the ordinary DQMC dynamics with the contiguous collective excitations. For a segment of length M , the collective process inverts M consecutive auxiliary-field spins and applies a single Metropolis decision to the complete transformation, while the usual absolute-weight dynamics instead performs M sequential single-spin proposals, each with its own acceptance decision. A natural comparison is therefore

$$P_M(-\sigma | x, \sigma) \quad \text{and} \quad P_{\sigma \rightarrow -\sigma}^{\text{seg}}(x; M), \quad (5.38)$$

where the two processes involve the same number M of proposal attempts and elementary spin flips, respectively.

Figure 5.13 shows this comparison for $M = 1$ through $M = 128$ for the effective-energy shell containing the peak of the canonical distribution. The blue curve corresponds to the lagged probabilities of the ordinary Markov chain, while the red curve shows the accepted sign-changing probabilities of the contiguous segments.

For $M = 1$, the two curves coincide, as expected since they correspond to the same single spin-flip process, but they immediately deviate as one increases the segment length. The contiguous segments produce opposite-sign configurations more frequently than the




Figure 5.13: Comparison between sign changes generated by the ordinary DQMC dynamics and by contiguous collective excitations at the canonically most relevant effective energy $\bar{x}_{\text{peak}} = -3.23$. The blue curve shows the lagged probability $P_M(-\sigma | x, \sigma)$ of finding the opposite sign after M sequential single-spin Metropolis proposals, while the red curve shows the accepted sign-changing probability for a single contiguous segment containing M flipped auxiliary-field spins. Since the transition probabilities $P_{+\rightarrow-}$ and $P_{-\rightarrow+}$ were closely matched, only their average is shown. The black dashed line shows the reference linear scaling $MP_{\sigma\rightarrow-\sigma}(x; M=1)$.

absolute-weight dynamics and the corresponding sign-changing probabilities follow an initially linear scaling, but they lose that relative advantage for longer segment lengths, particularly past $M = 32$. That loss of advantage is due to the already described decrease in the conditional Metropolis acceptance in Sec. 5.3, which is caused by the larger effective-energy displacements of the longer collective moves.

The ordinary dynamics instead displays a monotonic increase with M , following a sublinear scaling in the shown range of M . Successive single-spin proposals allow the Markov chain to progressively lose memory of its initial sign, so its probability of being found in the opposite sign sector continues to grow even after the collective probability has begun to decrease. The local-update dynamics therefore overtakes the contiguous excitation between $M = 32$ and $M = 64$ for the parameters considered here. The enhancement associated with a single contiguous excitation is thus restricted to an intermediate range of M and does not persist for the larger segment lengths.

Based on these results, an additional hypothesis may be made on the reason that longer segment lengths show a decrease in the sign-changing acceptance probabilities. Large collective moves, just like the ordinary DQMC dynamics at longer lags, may facilitate the loss of memory of the initial sign. For a low-sign regime, this may mean that the probability of a proposed large collective update causing a sign inversion may be close to $1/2$,

analogous to the prediction in Eq. (5.37). Indeed, this saturation near $1/2$ was observed in Fig. 5.4 for the probability of proposing a sign-changing move. However, it is not expected that such a saturation happens at the level of the corresponding effective-energy displacement, which is exactly the cause of the decrease of the $P_{\sigma \rightarrow -\sigma}(x)$ probability in Fig. 5.5.

Finally, this comparison complements the fixed-shell convolution comparison of the previous section because $P_M(-\sigma | x, \sigma)$ is measured directly from the ordinary DQMC trajectory. It therefore includes the correlations between the successive proposals, rejected updates, drifts in the effective energy, and the dependence of each proposal on the previous configurations of the Markov chain. At intermediate M , the larger probability obtained from a contiguous excitation provides further evidence that such temporally organized collective moves may be more efficient at affecting the sign than the sequential single-spin flip dynamics.

6. Conclusions and outlook

The sign problem in DQMC is characterized by the reduced value of the average sign of the determinantal weights, which directly decreases the statistical accuracy of reweighted MC estimates at a fixed number of samples. This global quantity, however, does not describe how positive- and negative-weight configurations are organized within the space of auxiliary-field configurations, nor how easily a DQMC simulation can move between them. The purpose of this dissertation was to investigate this internal structure in DQMC, using the magnitude of the configuration weight to organize the configuration space and different classes of moves in this space to probe the connectivity between the two sign sectors. This type of characterization is an essential first step towards the goal of, faced with a fermionic model of interest, identifying the most efficient MC sampling prescription.

A central step in this program was the recognition that the quantity $x = -\log |W|$, which we designated by effective energy, is extensive in both inverse temperature and system size, which motivated introducing the intensive variable $\bar{x} = x/(\beta N)$ and made x a natural coordinate for describing the structure of the auxiliary-field configuration space. By resolving the number of configurations at fixed x according to their determinant sign we were led to the sign-resolved entropies $S_+(x)$ and $S_-(x)$. These were found to have very similar extensive behavior, while their relatively small difference $\Delta S(x)$ quantifies the local relative number of positive and negative configurations. In the regime where the sign problem becomes severe, a broad region in this effective-energy space develops in which both signs occur with nearly equal multiplicity. At the same time, the canonical distribution of the determinantal weights becomes increasingly concentrated in this same region. We therefore established that the small average sign can be understood as arising from the sampling of effective-energy shells in which two exponentially large populations of configurations almost cancel.

This near equality of the populations belonging to either sign sector does not imply that the corresponding configurations are locally well connected. By exhaustively probing neighbors connected via single-spin moves to configurations at fixed effective energy, we found that sign-changing moves remain rare even inside the region where positive and negative configurations occur with almost equal frequency. This establishes an important distinction between the relative abundance of the two signs and their connectivity under the elementary single-spin move employed by conventional DQMC simulations.

Including the Metropolis acceptance factor further showed that, for the parameter regime studied here, the rarity of sign-changing proposals is more restrictive than their subsequent acceptance. This motivated the study of collective auxiliary-field moves, to determine whether stronger connections between the two sign sectors emerge under transformations beyond the elementary single-spin update.

The collective moves studied reveal that this is indeed the case. Inverting a complete fixed-site worldline produces sign changes considerably more frequently than inverting a complete spatial slice. Contiguous fixed-site segments allowed us to interpolate continuously between a single-spin excitation and a complete worldline. As the segment length M increases, the probability of changing the determinant sign rises strongly, while the conditional Metropolis acceptance decreases, producing an intermediate range of segment lengths for which accepted sign changes are most frequent. By comparing excitations involving the same number of flipped spins, we established that this improvement cannot be attributed solely to the number of modified auxiliary-field spins, since their arrangement in space and imaginary time also has a significant effect.

The importance of this arrangement was confirmed by comparing contiguous segments with fully random excitations, the latter of which were accurately reproduced via a fixed-shell convolution. The former displayed an enhanced sign-changing probability over the random spin baseline, which remained substantial even after accounting for the Metropolis factor. Since these excitation types differ primarily through the organization of the modified spins, this improvement indicates that the determinant sign is sensitive to correlations along imaginary time and identifies temporal contiguity as a particularly effective way of connecting the two sign sectors.

The comparison with the actual local-update dynamics addresses a complementary question, since, by measuring the determinant sign after exactly M proposals of an ordinary DQMC simulation, we incorporate rejected moves, correlations between successive updates, and the drift of the configuration through effective-energy space. A single contiguous collective proposal at the same update scale M was found to be more effective at producing a final configuration with the opposite sign than the standard DQMC dynamics for intermediate segment lengths. Because these two processes differ both in how the Metropolis criterion is applied and in the number of auxiliary-field spins that are ultimately flipped, this comparison does not isolate the effect of temporal contiguity. However, together with the comparison with the random excitations, it shows that such organized

collective moves can explore the two sign sectors more effectively than the standard local-update dynamics over an intermediate range of M .

Taken together, these results distinguish between the statistical balance of the two sign sectors, their connectivity under different classes of auxiliary-field excitations, and how easily the standard DQMC dynamics moves between the two sign sectors. While positive- and negative-weight configurations can occur with nearly equal frequency in parameter regions where the average sign is very small, they remain only sparsely connected through single-spin moves, whereas substantially stronger connections appear under collective moves, particularly those organized along imaginary time. The severity of the sign problem and the ability of the Markov chain to explore the two sign sectors are therefore related, but distinct, problems.

This distinction may also have practical consequences in the quest to efficiently improve the accuracy of DQMC simulations when the physical parameters of interest fall in a range associated with a severe sign problem. In this regard, it is important to emphasize that the small average sign itself cannot be improved by changing the MC updates, since it is a property of the equilibrium ensemble. Nevertheless, if the Markov chain moves slowly between the two sign sectors, a long autocorrelation time of the sign may cause the simulation to spend a disproportionate amount of time in one sign sector, leading to an increased error bar for the estimated average sign. Since reweighted observables depend on both $\langle sO \rangle$ and $\langle s \rangle$, this effect may introduce an additional inefficiency in the sampling procedure, which adds to the intrinsic statistical cost associated with the sign problem. Collective moves that enhance mixing between the two sign sectors could help reduce this additional dynamical source of inefficiency.

Materializing such type of improvement requires going beyond the diagnostic calculations performed in this work, since the collective excitations were not incorporated into the Markov chain. Furthermore, we note that a larger probability of reaching the opposite sign does not necessarily imply a shorter autocorrelation time or smaller error bars for a physical observable, although fixed-site worldline updates have previously been used to improve DQMC ergodicity [36]. A direct test would therefore consist of introducing collective moves in an otherwise conventional DQMC simulation and measuring their effect on the autocorrelation times of the sign and of physical observables.

Another natural extension is to determine how general the configuration-space structure

found here is. The sign-resolved entropy analysis was performed at fixed chemical potential, interaction strength and lattice geometry, while the excitation analysis was restricted to a single point in the parameter space of the Hubbard model. Extending these calculations across other parameter ranges and lattice geometries would reveal whether the relative advantage of collective moves remains more broadly as the physical parameters and severity of the sign problem are varied. In particular, the sign-resolved entropy analysis could provide a more detailed characterization of the distinction identified in Ref. [29], between a reduced average sign intrinsic to the configuration space and those enhanced by absolute-weight sampling. By resolving the relative abundance of positive- and negative-sign configurations as a function of x , and combining this with the distribution of x selected by the sampling procedure for the models studied in said reference, one could better determine whether the small average sign arises predominantly from a local sign balance of typical regions of configuration space, or instead arises because absolute-weight sampling is concentrated in regions which have a similar number of positive- and negative-sign configurations.

References

- [1] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, “Monte Carlo calculations of coupled boson-fermion systems. I,” *Physical Review D*, vol. 24, no. 8, pp. 2278–2286, Oct. 1981. [Cited on pages iv, v, 1, and 8.]
- [2] J. E. Hirsch, “Discrete Hubbard-Stratonovich transformation for fermion lattice models,” *Physical Review B*, vol. 28, no. 7, pp. 4059–4061, Oct. 1983. [Cited on pages iv, v, 1, and 11.]
- [3] 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,” *Physical Review B*, vol. 41, no. 13, pp. 9301–9307, May 1990. [Cited on pages iv, v, 2, 22, and 24.]
- [4] M. Troyer and U.-J. Wiese, “Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations,” *Physical Review Letters*, vol. 94, no. 17, p. 170201, May 2005. [Cited on pages iv, v, 3, and 24.]
- [5] P. Coleman, *Introduction to Many-Body Physics*, 1st ed. Cambridge University Press, Nov. 2015. [Cited on pages 1, 6, and 18.]
- [6] J. Hubbard, “Electron correlations in narrow energy bands,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 276, no. 1365, pp. 238–257, Nov. 1963. [Cited on pages 1 and 6.]
- [7] V. J. Emery, “Theory of the quasi-one-dimensional electron gas with strong ‘on-site’ interactions,” *Physical Review B*, vol. 14, no. 7, pp. 2989–2994, Oct. 1976. [Cited on pages 1, 6, and 7.]
- [8] N. D. Mermin and H. Wagner, “Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models,” *Physical Review Letters*, vol. 17, no. 22, pp. 1133–1136, Nov. 1966. [Cited on page 1.]
- [9] K. G. Wilson and M. E. Fisher, “Critical Exponents in 3.99 Dimensions,” *Physical Review Letters*, vol. 28, no. 4, pp. 240–243, Jan. 1972. [Cited on page 1.]

- [10] T. Siro and A. Harju, “Exact diagonalization of the Hubbard model on graphics processing units,” *Computer Physics Communications*, vol. 183, no. 9, pp. 1884–1889, Sep. 2012. [Cited on page 1.]
- [11] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof’ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, E. Gull, and Simons Collaboration on the Many-Electron Problem, “Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms,” *Physical Review X*, vol. 5, no. 4, p. 041041, Dec. 2015. [Cited on pages 1 and 6.]
- [12] E. Gull, P. Werner, O. Parcollet, and M. Troyer, “Continuous-time auxiliary-field Monte Carlo for quantum impurity models,” *EPL (Europhysics Letters)*, vol. 82, no. 5, p. 57003, Jun. 2008. [Cited on page 1.]
- [13] R. R. D. Santos, “Introduction to quantum Monte Carlo simulations for fermionic systems,” *Brazilian Journal of Physics*, vol. 33, no. 1, pp. 36–54, Mar. 2003. [Cited on pages 1, 13, 14, 18, 19, and 122.]
- [14] D. J. Scalapino and R. L. Sugar, “Monte Carlo calculations of coupled boson-fermion systems. II,” *Physical Review B*, vol. 24, no. 8, pp. 4295–4308, Oct. 1981. [Cited on pages 1 and 8.]
- [15] J. E. Hirsch, “Two-dimensional Hubbard model: Numerical simulation study,” *Physical Review B*, vol. 31, no. 7, pp. 4403–4419, Apr. 1985. [Cited on pages 1, 7, and 24.]
- [16] —, “Simulations of the three-dimensional Hubbard model: Half-filled band sector,” *Physical Review B*, vol. 35, no. 4, pp. 1851–1859, Feb. 1987. [Cited on pages 1 and 7.]
- [17] J. E. Gubernatis, N. Kawashima, and P. Werner, *Quantum Monte Carlo Methods: Algorithms for Lattice Models*. New York: Cambridge University Press, 2016. [Cited on pages 2, 8, 14, and 17.]
- [18] W. Janke, “Monte Carlo Methods in Classical Statistical Physics,” in *Computational Many-Particle Physics*, H. Fehske, R. Schneider, and A. Weiße, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, vol. 739, pp. 79–140. [Cited on page 2.]

- [19] A. W. Sandvik, “Stochastic series expansion method with operator-loop update,” *Physical Review B*, vol. 59, no. 22, pp. R14 157–R14 160, Jun. 1999. [Cited on page 2.]
- [20] C. Wu and S.-C. Zhang, “Sufficient condition for absence of the sign problem in the fermionic quantum Monte Carlo algorithm,” *Physical Review B*, vol. 71, no. 15, p. 155115, Apr. 2005. [Cited on pages 2, 3, and 24.]
- [21] V. I. Iglovikov, E. Khatami, and R. T. Scalettar, “Geometry dependence of the sign problem in quantum Monte Carlo simulations,” *Physical Review B*, vol. 92, no. 4, p. 045110, Jul. 2015. [Cited on pages 2, 24, and 30.]
- [22] E. Dagotto, “Correlated electrons in high-temperature superconductors,” *Reviews of Modern Physics*, vol. 66, no. 3, pp. 763–840, Jul. 1994. [Cited on pages 2 and 6.]
- [23] A. I. Lichtenstein and M. I. Katsnelson, “Antiferromagnetism and d -wave superconductivity in cuprates: A cluster dynamical mean-field theory,” *Physical Review B*, vol. 62, no. 14, pp. R9283–R9286, Oct. 2000. [Cited on page 2.]
- [24] Z.-X. Li, Y.-F. Jiang, and H. Yao, “Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation,” *Physical Review B*, vol. 91, no. 24, p. 241117, Jun. 2015. [Cited on page 3.]
- [25] ———, “Majorana-Time-Reversal Symmetries: A Fundamental Principle for Sign-Problem-Free Quantum Monte Carlo Simulations,” *Physical Review Letters*, vol. 117, no. 26, p. 267002, Dec. 2016. [Cited on page 3.]
- [26] R. Levy and B. K. Clark, “Mitigating the Sign Problem through Basis Rotations,” *Physical Review Letters*, vol. 126, no. 21, p. 216401, May 2021. [Cited on pages 3 and 24.]
- [27] S. Chandrasekharan and U.-J. Wiese, “Meron-Cluster Solution of Fermion Sign Problems,” *Physical Review Letters*, vol. 83, no. 16, pp. 3116–3119, Oct. 1999. [Cited on page 3.]
- [28] M. Iazzi, A. A. Soluyanov, and M. Troyer, “Topological origin of the fermion sign problem,” *Physical Review B*, vol. 93, no. 11, p. 115102, Mar. 2016. [Cited on page 3.]
- [29] S. Tarat, B. Xiao, R. Mondaini, and R. T. Scalettar, “Deconvolving the components of the sign problem,” *Physical Review B*, vol. 105, no. 4, p. 045107, Jan. 2022. [Cited on pages 3, 46, and 100.]

- [30] T.-C. Yi, R. T. Scalettar, and R. Mondaini, “Hamming distance and the onset of quantum criticality,” *Physical Review B*, vol. 106, no. 20, p. 205113, Nov. 2022. [Cited on page 3.]
- [31] K. Ch’ng, N. Vazquez, and E. Khatami, “Unsupervised machine learning account of magnetic transitions in the Hubbard model,” *Physical Review E*, vol. 97, no. 1, p. 013306, Jan. 2018. [Cited on page 3.]
- [32] R. Larson, R. Mondaini, and R. T. Scalettar, “Sign-resolved statistics and the origin of bias in quantum Monte Carlo,” *Physical Review B*, vol. 113, no. 19, p. 195124, May 2026. [Cited on page 4.]
- [33] W. Chen and R. Mondaini, “Sign-problem-resilient singular-value probe in determinant quantum Monte Carlo,” arXiv:2608.19028, Aug. 2026. [Cited on page 4.]
- [34] M. Yao, D. Wang, and Q.-H. Wang, “Reducing autocorrelation time in determinant quantum Monte Carlo using the Wang-Landau algorithm: Application to the Holstein model,” *Physical Review E*, vol. 104, no. 2, p. 025305, Aug. 2021. [Cited on pages 4 and 48.]
- [35] S. Beyl, F. Goth, and F. F. Assaad, “Revisiting the hybrid quantum Monte Carlo method for Hubbard and electron-phonon models,” *Physical Review B*, vol. 97, no. 8, p. 085144, Feb. 2018. [Cited on pages 4 and 5.]
- [36] R. T. Scalettar, R. M. Noack, and R. R. P. Singh, “Ergodicity at large couplings with the determinant Monte Carlo algorithm,” *Physical Review B*, vol. 44, no. 19, pp. 10 502–10 507, Nov. 1991. [Cited on pages 4, 75, and 99.]
- [37] 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,” *Physical Review B*, vol. 40, no. 1, pp. 506–516, Jul. 1989. [Cited on pages 7 and 24.]
- [38] J. E. Hirsch, R. L. Sugar, D. J. Scalapino, and R. Blankenbecler, “Monte Carlo simulations of one-dimensional fermion systems,” *Physical Review B*, vol. 26, no. 9, pp. 5033–5055, Nov. 1982. [Cited on page 8.]
- [39] M. Suzuki, “Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems,” *Communications in Mathematical Physics*, vol. 51, no. 2, pp. 183–190, Jun. 1976. [Cited on page 9.]

- [40] —, “Quantum Monte Carlo methods and general decomposition theory of exponential operators and symplectic integrators,” *Physica A: Statistical Mechanics and its Applications*, vol. 205, no. 1-3, pp. 65–79, Apr. 1994. [Cited on page 9.]
- [41] F. F. Assaad, “SU(2)-spin Invariant Auxiliary Field Quantum Monte-Carlo Algorithm for Hubbard models,” in *High Performance Computing in Science and Engineering '98*, E. Krause and W. Jäger, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 105–111. [Cited on page 11.]
- [42] Z.-X. Li and H. Yao, “Sign-Problem-Free Fermionic Quantum Monte Carlo: Developments and Applications,” *Annual Review of Condensed Matter Physics*, vol. 10, no. 1, pp. 337–356, Mar. 2019. [Cited on pages 11 and 24.]
- [43] F. Assaad and H. Evertz, “World-line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons,” in *Computational Many-Particle Physics*, H. Fehske, R. Schneider, and A. Weiße, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, vol. 739, pp. 277–356. [Cited on pages 13, 18, 22, and 24.]
- [44] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of State Calculations by Fast Computing Machines,” *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, Jun. 1953. [Cited on page 17.]
- [45] W. K. Hastings, “Monte Carlo sampling methods using Markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, Apr. 1970. [Cited on page 17.]
- [46] D. V. Hinkley, “On the ratio of two correlated normal random variables,” *Biometrika*, vol. 56, no. 3, pp. 635–639, 1969. [Cited on page 23.]
- [47] A. Moreo, D. J. Scalapino, R. L. Sugar, S. R. White, and N. E. Bickers, “Numerical study of the two-dimensional Hubbard model for various band fillings,” *Physical Review B*, vol. 41, no. 4, pp. 2313–2320, Feb. 1990. [Cited on page 24.]
- [48] 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,” *International Journal of Modern Physics C*, vol. 16, no. 08, pp. 1319–1327, Aug. 2005. [Cited on page 25.]
- [49] C. Bauer, “Fast and stable determinant quantum Monte Carlo,” *SciPost Physics Core*, vol. 2, no. 2, p. 011, Jun. 2020. [Cited on pages 25 and 36.]

- [50] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed., ser. Johns Hopkins Studies in the Mathematical Sciences. Baltimore, MD: Johns Hopkins Univ. Press, 2013. [Cited on page 25.]
- [51] B. Cohen-Stead, S. M. Costa, J. Neuhaus, A. T. Ly, Y. Zhang, R. Scalettar, K. Barros, and S. Johnston, “Codebase release r0.3 for SmoQyDQMC.jl,” *SciPost Physics Codebases* 29-r0.3, May 2024. [Cited on pages 27 and 28.]
- [52] C.-R. Lee, S. Chiesa, C. N. Varney, E. Khatami, Z. Bai, E. D’Azevedo, M. Jarrell, Th. Maier, S. Savrasov, R. T. Scalettar, and K. Tomko, “QUEST: QUantum Electron Simulation Toolbox,” in *SciDAC 2010*, Chattanooga, Tennessee, USA, Jul. 2010. [Cited on page 30.]
- [53] L. E. Reichl, *A Modern Course in Statistical Physics*, 2nd ed. New York: Wiley, 1998. [Cited on page 42.]
- [54] B. A. Berg and T. Neuhaus, “Multicanonical algorithms for first order phase transitions,” *Physics Letters B*, vol. 267, no. 2, pp. 249–253, Sep. 1991. [Cited on page 45.]
- [55] F. Wang and D. P. Landau, “Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram,” *Physical Review E*, vol. 64, no. 5, p. 056101, Oct. 2001. [Cited on page 48.]
- [56] ———, “Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States,” *Physical Review Letters*, vol. 86, no. 10, pp. 2050–2053, Mar. 2001. [Cited on page 48.]
- [57] C. Zhou and R. N. Bhatt, “Understanding and improving the Wang-Landau algorithm,” *Physical Review E*, vol. 72, no. 2, p. 025701, Aug. 2005. [Cited on page 49.]
- [58] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 3rd ed. Cambridge University Press, Sep. 2020. [Cited on page 109.]
- [59] J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*, rev. ed., reprinted ed., ser. Wiley Series in Probability and Statistics. Chichester: John Wiley, 2002. [Cited on page 112.]
- [60] D. A. Harville, *Matrix Algebra from a Statistician’s Perspective*. New York: Springer, 1997. [Cited on page 115.]

- [61] J. Sherman and W. J. Morrison, “Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix,” *The Annals of Mathematical Statistics*, vol. 21, no. 1, pp. 124–127, Mar. 1950. [Cited on page 115.]
- [62] J. E. Hirsch, “Stable Monte Carlo algorithm for fermion lattice systems at low temperatures,” *Physical Review B*, vol. 38, no. 16, pp. 12 023–12 026, Dec. 1988. [Cited on page 122.]

A. DQMC derivations

A.1. Trace-determinant identity

The goal of this appendix is to derive the trace-determinant identity in Eq. (2.33), which is a central step in the determinantal formulation of the partition function, by converting a trace over the fermionic Fock space into a determinant over the single-particle Hilbert space. Let us then consider a collection of quadratic fermionic operators

$$\hat{A}_\ell = \sum_{i,j=1}^N \hat{c}_i^\dagger [A_\ell]_{ij} \hat{c}_j, \quad (\text{A.1})$$

where A_ℓ is the corresponding $N \times N$ single-particle matrix. We wish to show that

$$\text{Tr} \left[\prod_{\ell=1}^L e^{\hat{A}_\ell} \right] = \det \left(I + \prod_{\ell=1}^L e^{A_\ell} \right), \quad (\text{A.2})$$

where the products follow the convention adopted throughout this thesis,

$$\prod_{\ell=1}^L e^{A_\ell} \equiv e^{A_L} e^{A_{L-1}} \dots e^{A_1}, \quad (\text{A.3})$$

and similarly for the many-body operators $\hat{A}_\ell$.

The first step is to determine how the exponential of a quadratic operator acts on the fermionic creation operators. Using the canonical anticommutation relations,

$$\{\hat{c}_i, \hat{c}_j^\dagger\} = \delta_{ij}, \quad (\text{A.4})$$

one finds

$$\begin{aligned} [\hat{A}_\ell, \hat{c}_k^\dagger] &= \sum_{i,j} [A_\ell]_{ij} [\hat{c}_i^\dagger \hat{c}_j, \hat{c}_k^\dagger] \\ &= \sum_i \hat{c}_i^\dagger [A_\ell]_{ik}. \end{aligned} \quad (\text{A.5})$$

Introducing the row vector of creation operators

$$\hat{\mathbf{c}}^\dagger = \begin{pmatrix} \hat{c}_1^\dagger & \hat{c}_2^\dagger & \dots & \hat{c}_N^\dagger \end{pmatrix}, \quad (\text{A.6})$$

Eq. (A.5) may be written compactly as

$$[\hat{A}_\ell, \hat{\mathbf{c}}^\dagger] = \hat{\mathbf{c}}^\dagger A_\ell. \quad (\text{A.7})$$

Repeated commutation therefore gives

$$[\hat{A}_\ell, [\hat{A}_\ell, \dots [\hat{A}_\ell, \hat{c}^\dagger] \dots]] = \hat{c}^\dagger A_\ell^n \quad (\text{A.8})$$

after n commutators. Applying the Baker-Hausdorff lemma [58] then yields

$$\begin{aligned} e^{\hat{A}_\ell} \hat{c}^\dagger e^{-\hat{A}_\ell} &= \hat{c}^\dagger \sum_{n=0}^{\infty} \frac{A_\ell^n}{n!} \\ &= \hat{c}^\dagger e^{A_\ell}. \end{aligned} \quad (\text{A.9})$$

Thus, although $e^{\hat{A}_\ell}$ acts on the complete fermionic Fock space, its action on the creation operators is completely determined by the single-particle matrix e^{A_ℓ} .

Consider now the many-body operator

$$\hat{U} = \prod_{\ell=1}^L e^{\hat{A}_\ell}. \quad (\text{A.10})$$

Successive application of Eq. (A.9) gives

$$\hat{U} \hat{c}^\dagger \hat{U}^{-1} = \hat{c}^\dagger U, \quad (\text{A.11})$$

where

$$U = \prod_{\ell=1}^L e^{A_\ell} \quad (\text{A.12})$$

is the single-particle representation of $\hat{U}$.

To evaluate its trace, first suppose that U is diagonalizable, with eigenvalues $\lambda_1, \dots, \lambda_N$ and eigenbasis $\{|n_1, \dots, n_N\rangle : n_a \in \{0, 1\}\}$. Since each occupied single-particle state contributes one factor of its corresponding eigenvalue, the action of $\hat{U}$ on its eigenstates is

$$\hat{U} |n_1, n_2, \dots, n_N\rangle = \left(\prod_{a=1}^N \lambda_a^{n_a} \right) |n_1, n_2, \dots, n_N\rangle. \quad (\text{A.13})$$

The trace over the complete Fock space is therefore

$$\begin{aligned} \text{Tr } \hat{U} &= \sum_{\{n_a\}} \prod_{a=1}^N \lambda_a^{n_a} \\ &= \prod_{a=1}^N \left(\sum_{n_a=0}^1 \lambda_a^{n_a} \right) \\ &= \prod_{a=1}^N (1 + \lambda_a). \end{aligned} \quad (\text{A.14})$$

On the other hand,

$$\det(I + \mathcal{U}) = \prod_{a=1}^N (1 + \lambda_a). \quad (\text{A.15})$$

Combining Eqs. (A.14) and (A.15) gives

$$\text{Tr } \hat{\mathcal{U}} = \det(I + \mathcal{U}). \quad (\text{A.16})$$

Using Eqs. (A.10) and (A.12), we finally obtain the desired identity in Eq. (A.2).

A.2. Equal-time Green-function representation

In this appendix, we derive the representation of the Green function in terms of the short-time propagators $B_{\ell,\sigma}$ given in Eq. (2.59). For a fixed auxiliary-field configuration, we denote the imaginary-time propagator from $\tau = 0$ to the slice ℓ for the spin sector σ by

$$\hat{\mathcal{U}}_{\sigma}(\ell, 0) = \hat{B}_{\ell,\sigma} \cdots \hat{B}_{1,\sigma}, \quad (\text{A.17})$$

with corresponding single-particle matrix

$$\mathcal{U}_{\sigma}(\ell, 0) = B_{\ell,\sigma} \cdots B_{1,\sigma}. \quad (\text{A.18})$$

Since each $\hat{B}_{\ell,\sigma}$ is the fixed-auxiliary-field one-body slice propagator, the imaginary-time Heisenberg operators present in Eq. (2.58) are

$$\hat{c}_{i\sigma}(\ell\Delta\tau) = \hat{\mathcal{U}}_{\sigma}(\ell, 0)^{-1} \hat{c}_{i\sigma} \hat{\mathcal{U}}_{\sigma}(\ell, 0), \quad (\text{A.19})$$

$$\hat{c}_{i\sigma}^{\dagger}(\ell\Delta\tau) = \hat{\mathcal{U}}_{\sigma}(\ell, 0)^{-1} \hat{c}_{i\sigma}^{\dagger} \hat{\mathcal{U}}_{\sigma}(\ell, 0). \quad (\text{A.20})$$

Therefore, their product simplifies to

$$\hat{c}_{i\sigma}(\ell\Delta\tau) \hat{c}_{j\sigma}^{\dagger}(\ell\Delta\tau) = \hat{\mathcal{U}}_{\sigma}(\ell, 0)^{-1} \hat{c}_{i\sigma} \hat{c}_{j\sigma}^{\dagger} \hat{\mathcal{U}}_{\sigma}(\ell, 0). \quad (\text{A.21})$$

Using the complete discrete-time propagator for the fixed auxiliary-field configuration, the equal-time Green function can therefore be written as

$$\begin{aligned} G_{\sigma,ij}(\ell) &= \left\langle \hat{c}_{i\sigma}(\ell\Delta\tau) \hat{c}_{j\sigma}^{\dagger}(\ell\Delta\tau) \right\rangle_h \\ &= \frac{\text{Tr} \left[\hat{\mathcal{U}}_{\sigma}(L, 0) \hat{\mathcal{U}}_{\sigma}(\ell, 0)^{-1} \hat{c}_{i\sigma} \hat{c}_{j\sigma}^{\dagger} \hat{\mathcal{U}}_{\sigma}(\ell, 0) \right]}{\text{Tr} \left[\hat{\mathcal{U}}_{\sigma}(L, 0) \right]}. \end{aligned} \quad (\text{A.22})$$

Since the complete propagator factorizes at the imaginary-time cut as

$$\hat{\mathcal{U}}_{\sigma}(L, 0) = \hat{B}_{L,\sigma} \cdots \hat{B}_{\ell+1,\sigma} \hat{\mathcal{U}}_{\sigma}(\ell, 0), \quad (\text{A.23})$$

we obtain

$$G_{\sigma,ij}(\ell) = \frac{\text{Tr}[\hat{B}_{L,\sigma} \cdots \hat{B}_{\ell+1,\sigma} \hat{c}_{i\sigma} \hat{c}_{j\sigma}^\dagger \hat{B}_{\ell,\sigma} \cdots \hat{B}_{1,\sigma}]}{\text{Tr}[\hat{B}_{L,\sigma} \cdots \hat{B}_{1,\sigma}]}. \quad (\text{A.24})$$

Finally, using the cyclic property of the trace, the factor $\hat{B}_{\ell,\sigma} \cdots \hat{B}_{1,\sigma}$ can be moved from the right of the operator insertion to the beginning of the trace,

$$G_{\sigma,ij}(\ell) = \frac{\text{Tr}[\hat{\mathcal{B}}_\sigma(\ell) \hat{c}_{i\sigma} \hat{c}_{j\sigma}^\dagger]}{\text{Tr}[\hat{\mathcal{B}}_\sigma(\ell)]}, \quad (\text{A.25})$$

where we define

$$\hat{\mathcal{B}}_\sigma(\ell) \equiv \hat{B}_{\ell,\sigma} \cdots \hat{B}_{1,\sigma} \hat{B}_{L,\sigma} \cdots \hat{B}_{\ell+1,\sigma}, \quad (\text{A.26})$$

and denote the corresponding single-particle matrix by $\mathcal{B}_\sigma(\ell)$. Thus, changing the slice at which the equal-time Green function is evaluated corresponds to a cyclic permutation of the complete imaginary-time propagator.

It remains to evaluate the trace in Eq. (A.25). For this purpose, we introduce an auxiliary parameter η whose derivative generates the operator insertion $\hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}$. Define the matrix $E^{(ji)}$ by

$$[E^{(ji)}]_{ab} = \delta_{aj} \delta_{bi}, \quad (\text{A.27})$$

so that the corresponding many-body operator is

$$\sum_{a,b} \hat{c}_{a\sigma}^\dagger [E^{(ji)}]_{ab} \hat{c}_{b\sigma} = \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}. \quad (\text{A.28})$$

We then consider the function

$$Z_{ij}(\eta) = \text{Tr}[\hat{\mathcal{B}}_\sigma(\ell) e^{\eta \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}}]. \quad (\text{A.29})$$

Since

$$\left. \frac{d}{d\eta} e^{\eta \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}} \right|_{\eta=0} = \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}, \quad (\text{A.30})$$

differentiating Eq. (A.29) gives

$$\left. \frac{dZ_{ij}}{d\eta} \right|_{\eta=0} = \text{Tr}[\hat{\mathcal{B}}_\sigma(\ell) \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}]. \quad (\text{A.31})$$

On the other hand, the single-particle representation of $e^{\eta \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}}$ is $e^{\eta E^{(ji)}}$. Applying the trace-determinant identity derived in Appendix A.1 therefore gives

$$Z_{ij}(\eta) = \det \left[I + \mathcal{B}_\sigma(\ell) e^{\eta E^{(ji)}} \right]. \quad (\text{A.32})$$

Using Jacobi's formula [59],

$$\frac{d}{d\eta} \det A(\eta) = \det A(\eta) \cdot \text{Tr} \left[A^{-1}(\eta) \cdot \frac{dA(\eta)}{d\eta} \right], \quad (\text{A.33})$$

and evaluating at $\eta = 0$, we obtain

$$\begin{aligned} \left. \frac{dZ_{ij}}{d\eta} \right|_{\eta=0} &= \det [I + \mathcal{B}_\sigma(\ell)] \\ &\quad \times \text{Tr} \left[(I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) E^{(ji)} \right]. \end{aligned} \quad (\text{A.34})$$

Now, for an arbitrary matrix A ,

$$\text{Tr} [A \cdot E^{(ji)}] = \sum_{a,b} A_{ab} [E^{(ji)}]_{ba} = A_{ij}, \quad (\text{A.35})$$

and hence

$$\left. \frac{dZ_{ij}}{d\eta} \right|_{\eta=0} = \det [I + \mathcal{B}_\sigma(\ell)] \left[(I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) \right]_{ij}. \quad (\text{A.36})$$

The trace-determinant identity also gives

$$\text{Tr} [\hat{\mathcal{B}}_\sigma(\ell)] = \det [I + \mathcal{B}_\sigma(\ell)]. \quad (\text{A.37})$$

Dividing Eq. (A.36) by this expression therefore yields

$$\frac{\text{Tr} [\hat{\mathcal{B}}_\sigma(\ell) \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}]}{\text{Tr} [\hat{\mathcal{B}}_\sigma(\ell)]} = \left[(I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) \right]_{ij}. \quad (\text{A.38})$$

However, the Green function in Eq. (A.25) contains the operators in the opposite order.

Using the anticommutation relation,

$$\hat{c}_{i\sigma} \hat{c}_{j\sigma}^\dagger = \delta_{ij} - \hat{c}_{j\sigma}^\dagger \hat{c}_{i\sigma}, \quad (\text{A.39})$$

we find

$$\begin{aligned} G_{\sigma,ij}(\ell) &= \delta_{ij} - \left[(I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) \right]_{ij} \\ &= \left[I - (I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) \right]_{ij}. \end{aligned} \quad (\text{A.40})$$

The matrix appearing in brackets can be simplified as

$$\begin{aligned} I - (I + \mathcal{B}_\sigma(\ell))^{-1} \mathcal{B}_\sigma(\ell) &= (I + \mathcal{B}_\sigma(\ell))^{-1} [I + \mathcal{B}_\sigma(\ell) - \mathcal{B}_\sigma(\ell)] \\ &= (I + \mathcal{B}_\sigma(\ell))^{-1}. \end{aligned} \quad (\text{A.41})$$

Consequently,

$$G_\sigma(\ell) = [I + \mathcal{B}_\sigma(\ell)]^{-1}. \quad (\text{A.42})$$

Substituting Eq. (A.26), we finally obtain

$$G_\sigma(\ell) = [I + B_{\ell,\sigma} \cdots B_{1,\sigma} B_{L,\sigma} \cdots B_{\ell+1,\sigma}]^{-1}. \quad (\text{A.43})$$

This is the equal-time Green-function representation given in Eq. (2.59).

A.3. Local update determinant ratio

We now derive the determinant ratio associated with the inversion of a single HS variable shown in Eq. (2.69). Following the notation introduced in Sec. 2.3.3, the proposal

$$h_{\ell,i} \mapsto h'_{\ell,i} = -h_{\ell,i}, \quad (\text{A.44})$$

with all remaining auxiliary-field variables unchanged, modifies the one-slice propagator $B_{\ell,\sigma}$ as

$$B'_{\ell,\sigma} = [I + \Delta_{\ell,\sigma}(i)] B_{\ell,\sigma}, \quad (\text{A.45})$$

where $\Delta_{\ell,\sigma}(i)$ is diagonal and has a single non-zero element,

$$[\Delta_{\ell,\sigma}(i)]_{jk} = \left(e^{\nu\sigma(h'_{\ell,i} - h_{\ell,i})} - 1 \right) \delta_{ji} \delta_{ki}. \quad (\text{A.46})$$

Now, the determinant entering the fermionic weight is

$$\det M_\sigma = \det [I + B_{L,\sigma} \cdots B_{1,\sigma}]. \quad (\text{A.47})$$

To isolate the modified propagator, we cyclically shift the imaginary-time product so that $B_{\ell,\sigma}$ appears on the left. Using

$$\det(I + AB) = \det(I + BA), \quad (\text{A.48})$$

repeatedly gives

$$\begin{aligned}\det M_\sigma &= \det [I + B_{\ell,\sigma} \cdots B_{1,\sigma} B_{L,\sigma} \cdots B_{\ell+1,\sigma}] \\ &= \det [I + \mathcal{B}_\sigma(\ell)],\end{aligned}\tag{A.49}$$

where we have used the same notation as Eq. (A.26) for the cyclic product. After the proposed flip, Eq. (A.45) gives

$$\mathcal{B}'_\sigma(\ell) = [I + \Delta_{\ell,\sigma}(i)] \mathcal{B}_\sigma(\ell).\tag{A.50}$$

The determinant ratio for this spin sector is therefore

$$\begin{aligned}R_\sigma &\equiv \frac{\det M'_\sigma}{\det M_\sigma} \\ &= \frac{\det [I + (I + \Delta_{\ell,\sigma}(i)) \mathcal{B}_\sigma(\ell)]}{\det [I + \mathcal{B}_\sigma(\ell)]}.\end{aligned}\tag{A.51}$$

Expanding the matrix in the numerator and factoring $I + \mathcal{B}_\sigma(\ell)$ on the right gives

$$\begin{aligned}I + (I + \Delta_{\ell,\sigma}(i)) \mathcal{B}_\sigma(\ell) &= I + \mathcal{B}_\sigma(\ell) + \Delta_{\ell,\sigma}(i) \mathcal{B}_\sigma(\ell) \\ &= [I + \Delta_{\ell,\sigma}(i) \mathcal{B}_\sigma(\ell) (I + \mathcal{B}_\sigma(\ell))^{-1}] \\ &\quad \times [I + \mathcal{B}_\sigma(\ell)].\end{aligned}\tag{A.52}$$

The second determinant consequently cancels between numerator and denominator in Eq. (A.51), leaving

$$R_\sigma = \det [I + \Delta_{\ell,\sigma}(i) \mathcal{B}_\sigma(\ell) (I + \mathcal{B}_\sigma(\ell))^{-1}].\tag{A.53}$$

Using Eq. (A.42),

$$\begin{aligned}\mathcal{B}_\sigma(\ell) [I + \mathcal{B}_\sigma(\ell)]^{-1} &= [I + \mathcal{B}_\sigma(\ell) - I] [I + \mathcal{B}_\sigma(\ell)]^{-1} \\ &= [G_\sigma(\ell)]^{-1} - I \\ &= I - G_\sigma(\ell),\end{aligned}\tag{A.54}$$

and hence

$$R_\sigma = \det [I + \Delta_{\ell,\sigma}(i) (I - G_\sigma(\ell))].\tag{A.55}$$

Since $\Delta_{\ell,\sigma}(i)$ has only one non-zero element, the matrix added to the identity in Eq. (A.55) has rank one. More explicitly,

$$\Delta_{\ell,\sigma}(i) = [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T, \quad (\text{A.56})$$

where $\mathbf{e}_i$ is the i th basis vector. Therefore,

$$\Delta_{\ell,\sigma}(i) (I - G_\sigma(\ell)) = [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T (I - G_\sigma(\ell)). \quad (\text{A.57})$$

We may then apply the matrix determinant lemma [60],

$$\det(I + \mathbf{u} \mathbf{v}^T) = 1 + \mathbf{v}^T \mathbf{u}, \quad (\text{A.58})$$

where $\mathbf{u}$ and $\mathbf{v}$ are column vectors, thus yielding

$$\begin{aligned} R_\sigma &= 1 + [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i^T (I - G_\sigma(\ell)) \mathbf{e}_i \\ &= 1 + [\Delta_{\ell,\sigma}(i)]_{ii} [1 - G_{\sigma,ii}(\ell)], \end{aligned} \quad (\text{A.59})$$

which is exactly the result shown in Eq. (2.69).

A.4. Sherman-Morrison Green function update

We now consider an accepted single-spin update and derive the corresponding update of the equal-time Green function given in Eq. (2.72). In the notation used in Appendices A.2 and A.3, the matrix entering the Green function changes according to

$$M'_\sigma(\ell) = M_\sigma(\ell) + [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T \mathcal{B}_\sigma(\ell), \quad (\text{A.60})$$

with

$$G_\sigma(\ell) = [M_\sigma(\ell)]^{-1}. \quad (\text{A.61})$$

The matrix whose inverse defines the Green function therefore becomes

$$\begin{aligned} M'_\sigma(\ell) &= I + [I + \Delta_{\ell,\sigma}(i)] \mathcal{B}_\sigma(\ell) \\ &= M_\sigma(\ell) + [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T \mathcal{B}_\sigma(\ell). \end{aligned} \quad (\text{A.62})$$

This is a rank-one modification of $M_\sigma(\ell)$. The Sherman-Morrison formula [61],

$$(A + \mathbf{u} \mathbf{v}^T)^{-1} = A^{-1} - \frac{A^{-1} \mathbf{u} \mathbf{v}^T A^{-1}}{1 + \mathbf{v}^T A^{-1} \mathbf{u}}, \quad (\text{A.63})$$

where $\mathbf{u}$ and $\mathbf{v}$ are column vectors, can therefore be applied. Using $G_\sigma(\ell) = M_\sigma(\ell)^{-1}$ gives

$$G'_\sigma(\ell) = G_\sigma(\ell) - \frac{[\Delta_{\ell,\sigma}(i)]_{ii} G_\sigma(\ell) \mathbf{e}_i \mathbf{e}_i^T \mathcal{B}_\sigma(\ell) G_\sigma(\ell)}{1 + [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i^T \mathcal{B}_\sigma(\ell) G_\sigma(\ell) \mathbf{e}_i}. \quad (\text{A.64})$$

Since

$$\mathcal{B}_\sigma(\ell) G_\sigma(\ell) = I - G_\sigma(\ell), \quad (\text{A.65})$$

the denominator becomes

$$1 + [\Delta_{\ell,\sigma}(i)]_{ii} [1 - G_{\sigma,ii}(\ell)] = R_\sigma, \quad (\text{A.66})$$

which is precisely the single-spin determinant ratio obtained in Appendix A.3. Equation (A.64) thus reduces to

$$G'_\sigma(\ell) = G_\sigma(\ell) - \frac{1}{R_\sigma} G_\sigma(\ell) \Delta_{\ell,\sigma}(i) [I - G_\sigma(\ell)], \quad (\text{A.67})$$

as we wished to show.

A.5. Absence of the sign problem at half filling

We now show that the repulsive Hubbard model with the spin-channel HS transformation used in this thesis is free of the fermion sign problem at half filling on a bipartite lattice. In the particle-hole-symmetric convention adopted in Chapter 2, half filling corresponds to $\mu = 0$.

Consider a bipartite lattice divided into two sublattices, A and B , and define

$$\eta_i = \begin{cases} +1, & i \in A, \\ -1, & i \in B. \end{cases} \quad (\text{A.68})$$

We perform a particle-hole transformation on the down-spin fermions,

$$\hat{c}_{i\downarrow} \longrightarrow \eta_i \hat{c}_{i\downarrow}^\dagger, \quad \hat{c}_{i\downarrow}^\dagger \longrightarrow \eta_i \hat{c}_{i\downarrow}. \quad (\text{A.69})$$

The corresponding number operator transforms as

$$\hat{n}_{i\downarrow} = \hat{c}_{i\downarrow}^\dagger \hat{c}_{i\downarrow} \longrightarrow \hat{c}_{i\downarrow} \hat{c}_{i\downarrow}^\dagger = 1 - \hat{n}_{i\downarrow}. \quad (\text{A.70})$$

At $\mu = 0$, the quadratic part of the Hamiltonian contains only the hopping term,

$$\hat{T}_\downarrow = - \sum_{i,j} t_{ij} \hat{c}_{i\downarrow}^\dagger \hat{c}_{j\downarrow}. \quad (\text{A.71})$$

For the bipartite hopping considered here, t_{ij} is non-zero only between opposite sublattices, for which

$$\eta_i \eta_j = -1. \quad (\text{A.72})$$

Under the transformation in Eq. (A.69), the hopping term is transformed as

$$\hat{T}_\downarrow \longrightarrow - \sum_{i,j} t_{ij} \eta_i \eta_j \hat{c}_{i\downarrow}^\dagger \hat{c}_{j\downarrow} = - \sum_{i,j} t_{ij} \hat{c}_{j\downarrow}^\dagger \hat{c}_{i\downarrow} = \hat{T}_\downarrow, \quad (\text{A.73})$$

where in the last step we used the fermionic anticommutation relation, $t_{ij} = t_{ji}$ and re-labeled the site indices. The kinetic term is therefore invariant under the particle-hole transformation.

For a fixed auxiliary-field configuration h , the short-time propagator for the down-spin sector is

$$\hat{B}_{\ell,\downarrow} = e^{-\nu \sum_i h_{\ell,i} \hat{n}_{i\downarrow}} e^{-\Delta\tau \hat{T}_\downarrow}. \quad (\text{A.74})$$

Using Eqs. (A.70) and (A.73), its auxiliary-field factor transforms as

$$e^{-\nu \sum_i h_{\ell,i} \hat{n}_{i\downarrow}} \longrightarrow e^{-\nu \sum_i h_{\ell,i} (1 - \hat{n}_{i\downarrow})} = e^{-\nu \sum_i h_{\ell,i}} e^{\nu \sum_i h_{\ell,i} \hat{n}_{i\downarrow}}. \quad (\text{A.75})$$

Consequently,

$$\hat{B}_{\ell,\downarrow} \longrightarrow e^{-\nu \sum_i h_{\ell,i}} e^{\nu \sum_i h_{\ell,i} \hat{n}_{i\downarrow}} e^{-\Delta\tau \hat{T}_\downarrow}. \quad (\text{A.76})$$

Apart from the positive scalar prefactor, the transformed operator has exactly the same form as the up-spin propagator. Applying the transformation to the complete imaginary-time product therefore gives

$$\prod_{\ell=1}^L \hat{B}_{\ell,\downarrow} \longrightarrow \exp \left[-\nu \sum_{\ell,i} h_{\ell,i} \right] \prod_{\ell=1}^L \hat{B}_{\ell,\uparrow}, \quad (\text{A.77})$$

where the spin label on the transformed fermionic operators is irrelevant, since the two spin sectors have identical Fock spaces. Taking the trace gives

$$\text{Tr} \left[\prod_{\ell=1}^L \hat{B}_{\ell,\downarrow} \right] = \exp \left[-\nu \sum_{\ell,i} h_{\ell,i} \right] \text{Tr} \left[\prod_{\ell=1}^L \hat{B}_{\ell,\uparrow} \right]. \quad (\text{A.78})$$

Using the trace-determinant identity derived in Appendix A.1, Eq. (A.78) becomes

$$\det M_{\downarrow}(h) = \exp \left[-\nu \sum_{\ell,i} h_{\ell,i} \right] \det M_{\uparrow}(h). \quad (\text{A.79})$$

The DQMC weight of the auxiliary-field configuration is therefore

$$\begin{aligned} W(h) &= \det M_{\uparrow}(h) \det M_{\downarrow}(h) \\ &= \exp \left[-\nu \sum_{\ell,i} h_{\ell,i} \right] [\det M_{\uparrow}(h)]^2. \end{aligned} \quad (\text{A.80})$$

All matrices entering the present formulation are real, so $\det M_{\uparrow}(h)$ is real. Since the exponential factor is strictly positive, it follows that $W(h) \geq 0$ for every auxiliary-field configuration h . There are therefore no negative configuration weights, and the average sign satisfies

$$\langle s \rangle_{|W|} = 1. \quad (\text{A.81})$$

This result relies on both particle-hole symmetry and the bipartite structure of the hopping. Away from half filling, the chemical-potential contribution is not invariant under Eq. (A.69), since

$$-\mu \hat{n}_{i\downarrow} \longrightarrow -\mu + \mu \hat{n}_{i\downarrow}. \quad (\text{A.82})$$

The transformation therefore reverses the sign of the chemical potential, and the relation in Eq. (A.79) no longer holds for $\mu \neq 0$. Similarly, hopping between sites on the same sublattice would violate the condition $\eta_i \eta_j = -1$ and generally remove the positivity argument.

B. Determinant ratios for collective excitations

B.1. General collective determinant ratio

We first establish a general expression for the determinant ratio produced by a collective modification of the auxiliary field. Suppose that a collective excitation modifies the matrix M_σ according to

$$M'_\sigma = M_\sigma + U_\sigma V_\sigma^T, \quad (\text{B.1})$$

where U_σ and V_σ are $N \times r$ matrices. The integer r is the rank of the modification and depends on the structure of the excitation.

The determinant ratio for the spin sector σ is

$$R_\sigma \equiv \frac{\det M'_\sigma}{\det M_\sigma} = \frac{\det [M_\sigma + U_\sigma V_\sigma^T]}{\det M_\sigma}. \quad (\text{B.2})$$

Factoring M_σ gives

$$R_\sigma = \det [I + M_\sigma^{-1} U_\sigma V_\sigma^T]. \quad (\text{B.3})$$

Using the identity $\det(I + AB) = \det(I + BA)$ and $M_\sigma^{-1} = G_\sigma$, this finally becomes

$$R_\sigma = \det [I_r + V_\sigma^T G_\sigma U_\sigma], \quad (\text{B.4})$$

where I_r denotes the $r \times r$ identity matrix.

Equation (B.4) reduces the ratio of two $N \times N$ determinants to the determinant of an $r \times r$ matrix. For a rank-one modification, $r = 1$, it reduces to the matrix determinant lemma used in Appendix A.3. The collective excitations considered below differ only in the form of the matrices U_σ and V_σ that describe the corresponding modification of the imaginary-time propagators.

B.2. Complete time-slice excitations

We first consider the collective excitation in which all auxiliary-field variables at a fixed imaginary-time slice ℓ are inverted,

$$h_{\ell,i} \longmapsto -h_{\ell,i}, \quad i = 1, \dots, N. \quad (\text{B.5})$$

Since only the auxiliary fields at slice ℓ are modified, the excitation affects a single propagator $B_{\ell,\sigma}$. The updated propagator can be written as

$$B'_{\ell,\sigma} = \left[I + \Delta_{\sigma}^{(\ell)} \right] B_{\ell,\sigma}, \quad (\text{B.6})$$

where $\Delta_{\sigma}^{(\ell)}$ is diagonal, with

$$\left[\Delta_{\sigma}^{(\ell)} \right]_{ij} = \left(e^{-2\nu\sigma h_{\ell,i}} - 1 \right) \delta_{ij}. \quad (\text{B.7})$$

As in Appendix A.3, we cyclically shift the imaginary-time product so that the modified propagator lies at the cut, and define

$$\mathcal{B}_{\sigma}(\ell) = B_{\ell,\sigma} \cdots B_{1,\sigma} B_{L,\sigma} \cdots B_{\ell+1,\sigma}. \quad (\text{B.8})$$

The corresponding fermionic matrix and equal-time Green function are

$$M_{\sigma}(\ell) = I + \mathcal{B}_{\sigma}(\ell), \quad G_{\sigma}(\ell) = M_{\sigma}(\ell)^{-1}. \quad (\text{B.9})$$

After the collective inversion,

$$\mathcal{B}'_{\sigma}(\ell) = \left[I + \Delta_{\sigma}^{(\ell)} \right] \mathcal{B}_{\sigma}(\ell), \quad (\text{B.10})$$

so that

$$\begin{aligned} M'_{\sigma}(\ell) &= I + \left[I + \Delta_{\sigma}^{(\ell)} \right] \mathcal{B}_{\sigma}(\ell) \\ &= M_{\sigma}(\ell) + \Delta_{\sigma}^{(\ell)} \mathcal{B}_{\sigma}(\ell). \end{aligned} \quad (\text{B.11})$$

This is of the general form introduced in Eq. (B.1) with

$$U_{\sigma} = \Delta_{\sigma}^{(\ell)}, \quad V_{\sigma}^T = \mathcal{B}_{\sigma}(\ell). \quad (\text{B.12})$$

In the present case U_{σ} and V_{σ} are $N \times N$ matrices, so the update has rank at most N . Applying Eq. (B.4), the determinant ratio is

$$R_{\sigma}^{(\ell)} = \det \left[I + \mathcal{B}_{\sigma}(\ell) G_{\sigma}(\ell) \Delta_{\sigma}^{(\ell)} \right]. \quad (\text{B.13})$$

Using

$$\begin{aligned} \mathcal{B}_{\sigma}(\ell) G_{\sigma}(\ell) &= \left[[G_{\sigma}(\ell)]^{-1} - I \right] G_{\sigma}(\ell) \\ &= I - G_{\sigma}(\ell), \end{aligned} \quad (\text{B.14})$$

we obtain

$$R_{\sigma}^{(\ell)} = \det \left[I + (I - G_{\sigma}(\ell)) \Delta_{\sigma}^{(\ell)} \right]. \quad (\text{B.15})$$

The total determinant ratio for the collective excitation is therefore

$$R^{(\ell)} = R_{\uparrow}^{(\ell)} R_{\downarrow}^{(\ell)}. \quad (\text{B.16})$$

Unlike the single-spin update, for which the corresponding modification has rank one, the full time-slice inversion generally produces a rank- N modification. The determinant in Eq. (B.15) is therefore an $N \times N$ determinant, with computational cost $\mathcal{O}(N^3)$. Since there are L possible time slices, probing all fixed-time spatial-slice excitations of a given configuration has a leading cost $\mathcal{O}(LN^3)$.

This does not improve the asymptotic scaling relative to constructing the same collective excitation as a sequence of single-spin inversions. Although each single-spin determinant ratio can be evaluated in $\mathcal{O}(1)$ operations, the Green function must be updated after each inversion before the ratio associated with the next spin can be computed. Each such update costs $\mathcal{O}(N^2)$, so applying the N constituent inversions of one time slice requires $\mathcal{O}(N^3)$ operations. Repeating this procedure for all L time slices therefore also gives a total cost of $\mathcal{O}(LN^3)$. The direct collective formula in Eq. (B.15) nevertheless avoids the sequence of temporary Green function updates and the associated copies of the Green function, thereby reducing the constant computational overhead.

B.3. Complete fixed-site worldline excitations

We next consider the collective excitation in which the auxiliary field at a fixed spatial site i is inverted along the complete imaginary-time direction,

$$h_{\ell,i} \longmapsto -h_{\ell,i}, \quad \ell = 1, \dots, L. \quad (\text{B.17})$$

For a given configuration there are therefore N possible complete fixed-site worldline excitations.

Unlike the complete time-slice excitation, this transformation does not modify a single propagator. Instead, one spatial entry is modified in every time slice,

$$B'_{\ell,\sigma} = [I + \Delta_{\ell,\sigma}(i)] B_{\ell,\sigma}, \quad (\text{B.18})$$

where

$$[\Delta_{\ell,\sigma}(i)]_{jk} = (e^{-2\nu\sigma h_{\ell,i}} - 1) \delta_{ji} \delta_{ki}. \quad (\text{B.19})$$

The determinant ratio is therefore

$$R_{\sigma}^{(i)} = \frac{\det[I + B'_{L,\sigma} \cdots B'_{1,\sigma}]}{\det[I + B_{L,\sigma} \cdots B_{1,\sigma}]}. \quad (\text{B.20})$$

Because the modified propagators are distributed over all imaginary-time slices, Eq. (B.20) cannot be reduced to a low-rank modification of a single equal-time matrix in the same way as the fixed-time spatial-slice excitation. A useful alternative is to resort to a space-time formulation of DQMC [13, 62] by introducing the $LN \times LN$ space-time matrix

$$\mathcal{M}_{\sigma} = \begin{pmatrix} I & 0 & 0 & \cdots & B_{L,\sigma} \\ -B_{1,\sigma} & I & 0 & \cdots & 0 \\ 0 & -B_{2,\sigma} & I & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & -B_{L-1,\sigma} & I \end{pmatrix}. \quad (\text{B.21})$$

Successive block-row operations,

$$R_{\ell+1} \mapsto R_{\ell+1} + B_{\ell,\sigma} R_{\ell}, \quad \ell = 1, \dots, L-1, \quad (\text{B.22})$$

transform Eq. (B.21) into an upper block-triangular matrix,

$$\mathcal{M}_{\sigma} \longrightarrow \begin{pmatrix} I & 0 & 0 & \cdots & B_{L,\sigma} \\ 0 & I & 0 & \cdots & B_{1,\sigma} B_{L,\sigma} \\ 0 & 0 & I & \cdots & B_{2,\sigma} B_{1,\sigma} B_{L,\sigma} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & I + B_{L-1,\sigma} \cdots B_{1,\sigma} B_{L,\sigma} \end{pmatrix}, \quad (\text{B.23})$$

from which we can immediately read off that

$$\det \mathcal{M}_{\sigma} = \det[I + B_{L-1,\sigma} \cdots B_{1,\sigma} B_{L,\sigma}] = \det M_{\sigma}, \quad (\text{B.24})$$

thus confirming the equivalence with our current DQMC formulation.

The space-time representation is particularly useful because the complete worldline inversion becomes a low-rank modification of $\mathcal{M}_{\sigma}$. To see this, let $\mathbf{e}_{\ell}$ denote the ℓ th basis vector in the L -dimensional imaginary-time space. For $\ell = 1, \dots, L-1$, the modification of

the block $-B_{\ell,\sigma}$ is

$$-B'_{\ell,\sigma} = -B_{\ell,\sigma} - [\Delta_{\ell,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T B_{\ell,\sigma}, \quad (\text{B.25})$$

while the upper-right block changes according to

$$B'_{L,\sigma} = B_{L,\sigma} + [\Delta_{L,\sigma}(i)]_{ii} \mathbf{e}_i \mathbf{e}_i^T B_{L,\sigma}. \quad (\text{B.26})$$

The full modification may therefore be written as

$$\mathcal{M}'_{\sigma} = \mathcal{M}_{\sigma} + U_{\sigma} V_{\sigma}^T, \quad (\text{B.27})$$

where U_{σ} and V_{σ} are $LN \times L$ matrices whose columns are defined by

$$u_{\ell,\sigma} = -[\Delta_{\ell,\sigma}(i)]_{ii} (\mathbf{e}_{\ell+1} \otimes \mathbf{e}_i), \quad v_{\ell,\sigma}^T = \mathbf{e}_{\ell}^T \otimes (\mathbf{e}_i^T B_{\ell,\sigma}), \quad \ell = 1, \dots, L-1, \quad (\text{B.28})$$

and

$$u_{L,\sigma} = [\Delta_{L,\sigma}(i)]_{ii} (\mathbf{e}_1 \otimes \mathbf{e}_i), \quad v_{L,\sigma}^T = \mathbf{e}_L^T \otimes (\mathbf{e}_i^T B_{L,\sigma}). \quad (\text{B.29})$$

Thus, although the space-time matrix has dimension LN , the complete worldline inversion modifies it by a matrix of rank at most L .

Applying the general determinant identity derived in Appendix B.1 gives

$$R_{\sigma}^{(i)} = \det [I_L + V_{\sigma}^T \mathcal{M}_{\sigma}^{-1} U_{\sigma}]. \quad (\text{B.30})$$

The $LN \times LN$ inverse $\mathcal{M}_{\sigma}^{-1}$ is the corresponding space-time Green-function matrix. Equation (B.30) therefore expresses the complete worldline excitation as an $L \times L$ determinant, in direct analogy with the $N \times N$ determinant obtained for a fixed-time spatial-slice excitation. In this sense, the two constructions interchange the roles of space and imaginary time.

In our DQMC implementation, however, Eq. (B.30) is not used directly. Evaluating it would require constructing the relevant time-displaced Green functions contained in $\mathcal{M}_{\sigma}^{-1}$, together with their numerical stabilization. Moreover, in the parameter regime considered in this thesis one typically has $L \gg N$, so the evaluation of an $L \times L$ determinant is not necessarily advantageous.

Instead, the complete worldline excitation is evaluated exactly as a sequence of L local modifications on a temporary copy of the configuration. Starting from the original Green functions, the copy is propagated through imaginary time. At each slice ℓ , the field $h_{\ell,i}$

is inverted, the corresponding single-spin determinant ratio $r_{\ell,\sigma}^{(i)}$ is evaluated, and the temporary Green function is updated using the Sherman-Morrison formula. No Metropolis decision is made after the individual inversions. After all L slices have been visited, the product

$$R_{\sigma}^{(i)} = \prod_{\ell=1}^L r_{\ell,\sigma}^{(i)} \quad (\text{B.31})$$

is exactly the determinant ratio between the final collectively modified configuration and the original one. The full weight ratio is then

$$R^{(i)} = R_{\uparrow}^{(i)} R_{\downarrow}^{(i)}. \quad (\text{B.32})$$

For a single fixed site, the L local determinant ratios themselves require only $\mathcal{O}(L)$ operations, while the corresponding Green-function updates cost $\mathcal{O}(LN^2)$. The dominant contribution arises from propagating the Green functions through the L imaginary-time slices, which requires a total of $\mathcal{O}(LN^3)$ operations. Since the excitation must be probed independently for each of the N spatial sites, evaluating all complete fixed-site worldline excitations of one configuration has a leading computational cost $\mathcal{O}(LN^4)$.

B.4. Contiguous fixed-site segment excitations

We now consider the inversion of a contiguous segment of M auxiliary-field spins at a fixed spatial site i . For a starting slice ℓ_0 , the excitation acts on

$$h_{\ell,i} \mapsto -h_{\ell,i}, \quad \ell = \ell_0, \dots, \ell_f, \quad (\text{B.33})$$

where

$$\ell_f = \ell_0 + M - 1, \quad (\text{B.34})$$

with all imaginary-time indices understood modulo L . For $1 \leq M < L$, there are NL possible choices of (i, ℓ_0) .

For a given segment, we define the original product of the propagators contained inside the interval by

$$\mathcal{S}_{\sigma}(\ell_0) = B_{\ell_f,\sigma} \cdots B_{\ell_0,\sigma}, \quad (\text{B.35})$$

and the corresponding modified product by

$$\tilde{\mathcal{S}}_{\sigma}(\ell_0, i) = B'_{\ell_f,\sigma} \cdots B'_{\ell_0,\sigma}, \quad (\text{B.36})$$

where

$$B'_{\ell,\sigma} = [I + \Delta_{\ell,\sigma}(i)] B_{\ell,\sigma} \quad (\text{B.37})$$

is the propagator obtained after inverting $h_{\ell,i}$.

The remaining propagators form the complementary product

$$\mathcal{C}_\sigma(\ell_0) = B_{\ell_0-1,\sigma} \cdots B_{\ell_f+1,\sigma}, \quad (\text{B.38})$$

ordered such that the cyclic product at the end of the segment is

$$M_\sigma(\ell_f) = I + \mathcal{S}_\sigma(\ell_0) \mathcal{C}_\sigma(\ell_0). \quad (\text{B.39})$$

The corresponding equal-time Green function is therefore

$$G_\sigma(\ell_f) = [M_\sigma(\ell_f)]^{-1}. \quad (\text{B.40})$$

It is useful to introduce the relative replacement matrix

$$X_\sigma(\ell_0, i) = \tilde{\mathcal{S}}_\sigma(\ell_0, i) [\mathcal{S}_\sigma(\ell_0)]^{-1}, \quad (\text{B.41})$$

so that

$$\tilde{\mathcal{S}}_\sigma(\ell_0, i) = X_\sigma(\ell_0, i) \mathcal{S}_\sigma(\ell_0). \quad (\text{B.42})$$

The fermionic matrix after the collective excitation can then be written as

$$\begin{aligned} M'_\sigma(\ell_f) &= I + \tilde{\mathcal{S}}_\sigma(\ell_0, i) \mathcal{C}_\sigma(\ell_0) \\ &= I + X_\sigma(\ell_0, i) \mathcal{S}_\sigma(\ell_0) \mathcal{C}_\sigma(\ell_0) \\ &= M_\sigma(\ell_f) + [X_\sigma(\ell_0, i) - I] \mathcal{S}_\sigma(\ell_0) \mathcal{C}_\sigma(\ell_0). \end{aligned} \quad (\text{B.43})$$

This has the general form of Eq. (B.1) with

$$U_\sigma = X_\sigma(\ell_0, i) - I, \quad V_\sigma^T = \mathcal{S}_\sigma(\ell_0) \mathcal{C}_\sigma(\ell_0). \quad (\text{B.44})$$

Applying Eq. (B.4) and using

$$\mathcal{S}_\sigma(\ell_0) \mathcal{C}_\sigma(\ell_0) G_\sigma(\ell_f) = I - G_\sigma(\ell_f), \quad (\text{B.45})$$

finally yields the determinant ratio

$$R_\sigma(\ell_0, i) = \det [I + (I - G_\sigma(\ell_f)) (X_\sigma(\ell_0, i) - I)]. \quad (\text{B.46})$$

The complete weight ratio is

$$R(\ell_0, i) = R_{\uparrow}(\ell_0, i) R_{\downarrow}(\ell_0, i). \quad (\text{B.47})$$

Although Eq. (B.46) provides a compact expression for the determinant ratio, it is not used directly in the numerical implementation. The relative replacement matrix $X_{\sigma}(\ell_0, i)$ contains both the inverse of the original segment product and the modified segment product. Explicitly constructing this inverse and multiplying it with the product of propagators would be numerically unstable, particularly for long imaginary-time intervals. Instead, the modified segment product and its complement are constructed directly using the same stabilized matrix-product machinery employed for the DQMC propagation in Sec. 2.5.

We begin by directly constructing the modified cyclic product

$$M'_{\sigma}(\ell_f) = I + \tilde{S}_{\sigma}(\ell_0, i) C_{\sigma}(\ell_0). \quad (\text{B.48})$$

Both the modified product and the complement are evaluated using stabilized matrix products, after which the determinant magnitude and sign of the modified configuration can be obtained directly. The determinant of the original configuration needs to be evaluated only once.

The remaining computational problem is therefore the efficient construction of $\tilde{S}_{\sigma}(\ell_0, i)$ for every possible starting slice ℓ_0 . Recomputing this product independently for each starting position would be highly redundant. Indeed, two consecutive segments,

$$\tilde{S}_{\sigma}(\ell_0, i) = B'_{\ell_f, \sigma} \cdots B'_{\ell_0, \sigma}, \quad (\text{B.49})$$

$$\tilde{S}_{\sigma}(\ell_0 + 1, i) = B'_{\ell_f+1, \sigma} \cdots B'_{\ell_0+1, \sigma}, \quad (\text{B.50})$$

share $M - 1$ of their M propagators. Shifting the segment by one time slice therefore amounts to removing the oldest propagator and adding one new propagator.

Formally, one could perform this shift through

$$\tilde{S}_{\sigma}(\ell_0 + 1, i) = B'_{\ell_0+M, \sigma} \tilde{S}_{\sigma}(\ell_0, i) [B'_{\ell_0, \sigma}]^{-1}. \quad (\text{B.51})$$

This direct sliding update would, however, require repeated multiplication by both propagators and their inverses, which is numerically unstable. We therefore avoid removing the outgoing propagator explicitly. Instead, the segment is represented by two collections

of stored partial products, constructed in the same stabilized manner as the matrix products discussed in Sec. 2.5. This allows the oldest propagator to be removed simply by selecting a previously stored partial product, without evaluating any matrix inverse.

For a given segment $[\ell_0, \ell_f]$, we choose an intermediate slice ℓ_m and divide the ordered sequence of modified propagators into two parts. For the older part of the segment, extending from ℓ_0 to ℓ_m , we store the cumulative products

$$\mathcal{F}_\sigma(r) = B'_{\ell_m, \sigma} \cdots B'_{r, \sigma}, \quad r = \ell_0, \dots, \ell_m, \quad (\text{B.52})$$

so that $\mathcal{F}_\sigma(r)$ represents the product of all propagators between the slice r and the splitting point ℓ_m . We refer to this collection of stored partial products as the front stack.

For the newer part of the segment, extending from $\ell_m + 1$ to ℓ_f , we analogously store

$$\mathcal{Q}_\sigma(r) = B'_{r, \sigma} \cdots B'_{\ell_m+1, \sigma}, \quad r = \ell_m + 1, \dots, \ell_f, \quad (\text{B.53})$$

where $\mathcal{Q}_\sigma(r)$ contains the product from the beginning of the newer part up to the slice r . We refer to this second collection as the back stack.

The complete segment product is then obtained as

$$\tilde{\mathcal{S}}_\sigma(\ell_0, i) = \mathcal{Q}_\sigma(\ell_f) \mathcal{F}_\sigma(\ell_0), \quad (\text{B.54})$$

with the back factor omitted when the back stack is empty.

The advantage of this representation becomes apparent when the segment is shifted. Removing the oldest propagator does not require modifying a matrix product, as one simply replaces $\mathcal{F}_\sigma(\ell_0)$ by the already stored partial product $\mathcal{F}_\sigma(\ell_0 + 1)$. The new propagator entering at the opposite end of the segment is incorporated into the back stack through a single multiplication,

$$\mathcal{Q}_\sigma(\ell_f + 1) = B'_{\ell_f+1, \sigma} \mathcal{Q}_\sigma(\ell_f). \quad (\text{B.55})$$

Thus, advancing the window requires no inverse of the propagator that leaves the segment.

Eventually, all propagators stored in the front stack are removed. At this point, the matrices accumulated in the back stack are transferred to the front stack in reverse order, while the corresponding partial products are reconstructed. Consequently, over a complete traversal of the L possible starting slices, only $\mathcal{O}(L)$ matrix multiplications are required

rather than $\mathcal{O}(LM)$.

All partial products in Eqs. (B.52) and (B.53) are maintained in stabilized UDV form. As such, this two-stack construction avoids both the issue of repeated recomputation of strongly overlapping products and the use of the naive and numerically unstable formula in Eq. (B.51).

The complementary products $\mathcal{C}_\sigma(\ell_0, M)$ can be constructed in the same manner using a sliding window of length $L - M$. Unlike the modified segment products, they do not depend on the spatial site i . The L complementary products are therefore computed once for a given configuration and reused when all N spatial sites are probed.

For a fixed spatial site, generating the modified segment products for all L starting positions requires $\mathcal{O}(LN^3)$ operations, since each stabilized multiplication of dense $N \times N$ matrices costs $\mathcal{O}(N^3)$. For each of the L segments, the modified segment product must then be combined with its corresponding complement and the resulting determinant evaluated, which also gives a total cost $\mathcal{O}(LN^3)$ per spatial site. Repeating the procedure for all N sites therefore gives $\mathcal{O}(LN^4)$. The initialization of the first length- M segment before the sliding-window reuse can begin requires $\mathcal{O}(MN^3)$ operations for each site, so a more explicit estimate is $\mathcal{O}((L + M)N^4)$. Since $M \leq L$, this has the same asymptotic scaling $\mathcal{O}(LN^4)$. The important point is that the sliding-window construction removes the additional factor of M that would arise if every length- M segment product were constructed independently.