

Discrete Time Quantum Walks with Engineered Decoherence as Quantum Models for Diffusive Transport

PhD Thesis

Lara Janiurek

QLiMIT

Department of Physics

University of Strathclyde, Glasgow

November 9, 2026

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Abstract

Quantum walks provide a natural framework for quantum algorithms and quantum simulation, but their coherent unitary dynamics usually leads to ballistic spreading rather than the diffusive behaviour characteristic of many classical transport processes. This thesis investigates whether discrete time quantum walks, augmented by controlled decoherence, can be used as quantum models for macroscopic transport equations such as diffusion and advection-diffusion. The work is motivated by the lattice Boltzmann method, where macroscopic fluid and transport equations emerge from local streaming and collision steps on a discrete velocity lattice.

The central idea explored in this thesis is that coin decoherence can be used constructively as a relaxation mechanism, rather than treated only as a source of noise. By damping off-diagonal coin coherences, projective dephasing can define a local equilibrium sector analogous to the role of collision in classical lattice Boltzmann methods. Continuous spacetime limits are first derived for decohered quantum walks, including both coin and spatial decoherence, and lazy quantum walks with an additional rest state are introduced to provide a D1Q3 velocity structure. A Chapman Enskog style expansion is then developed and applied to a range of decohered quantum walk models in order to identify how relaxation, coherent transport, and diffusive terms appear in the continuum hierarchy.

The analysis shows that simple two state and three state Grover walks contain useful ingredients for diffusion, and that the addition of a rest state improves the diffusive character of the simulations. However, their Chapman Enskog hierarchies remain strongly coupled, with coherent transport and dephasing appearing at the same leading order. Twistless walks produce cleaner continuum expansions by cancelling

singular lower order terms, but still do not fully separate equilibrium selection from unitary transport. This motivates the construction of a parity echo model, inspired by the cancellation mechanisms of the twistless walk. The parity echo model gives the closest match to the desired lattice Boltzmann inspired ordering, in which dephasing acts as a leading relaxation mechanism before higher order coherent corrections generate diffusion.

Across the models studied, the numerical simulations improve as the Chapman Enskog hierarchy becomes cleaner and closer to the desired relaxation/transport/diffusion ordering. This suggests that the multiscale hierarchy is not only a post-hoc explanation, but a useful design tool for engineering quantum walks with controlled macroscopic transport behaviour. The thesis does not claim to provide a complete quantum lattice Boltzmann algorithm for nonlinear fluid dynamics. Rather, it identifies the structural ingredients needed for decohered quantum walks to reproduce linear transport behaviour: controlled coin dephasing, suitable velocity structure, symmetry, and low order unitary cancellation. These results provide a foundation for future quantum walk based approaches to diffusion, advection-diffusion, and quantum kinetic simulation.

Contents

Abstract **ii**

List of Figures **vii**

List of Tables **xv**

Preface/Acknowledgements **xvii**

1 Introduction **2**

1.1 Motivation and Overview 2

1.2 Thesis Outline 4

2 Background **6**

2.1 Classical Computational fluid dynamics and the lattice Boltzmann methods 7

2.1.1 High Performance Computing 7

2.1.2 Classical Random Walks 9

2.1.3 The lattice Boltzmann method 15

2.1.4 D1Q3 lattice Boltzmann model and Chapman-Enskog expansion for advection-diffusion 20

2.2 Quantum Physics and Quantum Computing 26

2.2.1 Basic Principles of Quantum Mechanics 26

2.2.2 Quantum Computing 28

2.3 Quantum Walks 32

2.3.1 Discrete Time Quantum Walks 33

2.3.2 Continuum Limits and the Dirac Equation 39
2.3.3 Decoherence in Quantum Walks 49
2.4 Quantum Computational Fluid Dynamics and Quantum Lattice Boltzmann Methods 54

3 Continuous Limits of Lazy Open Quantum Walks 63
3.1 Lazy Quantum Walks 64
3.1.1 Lazy Discrete Time Quantum Walks 67
3.1.2 Fourier Basis 69
3.1.3 Gell Mann Representation 71
3.2 Decoherence Model 73
3.3 Continuum Limit of the Lazy Quantum Walk with Decoherence 76
3.3.1 Scaling and Parameterisation 76
3.3.2 Unitary Expansion and Effective Hamiltonian 77
3.3.3 Including Decoherence 80
3.4 Discussion 89

4 Decohered Quantum Walks as Kinetic Models for Diffusion 92
4.1 Motivation and Connection to classical LBM 93
4.2 General Model 94
4.3 A General Quantum CE Style Framework 95
4.3.1 Closure Step 101
4.3.2 Equilibrium and Desired Hierarchy 103
4.4 Grover Quantum Walk Model 107
4.4.1 The two state Grover quantum walk model 108
4.4.2 The three state lazy Grover quantum walk model 110
4.4.3 Comparison of the two state and three state Grover walks 112
4.5 Plastic Scaling and Regularity 117
4.6 Twistless Quantum Walk Model 120
4.6.1 The 2 state twistless model 123
4.6.2 The 3 state twistless model 126

4.6.3 Comparison of the two state and three state twistless models . . 131
4.7 Discussion 134

5 The Parity Echo Model **139**
5.1 The Parity Echo Model 140
5.1.1 The unitary parity echo model 141
5.2 The Chapman-Enskog Expansion 148
5.2.1 Decoherence Map and Scalings 148
5.2.2 Overdamped coherence closure 155
5.3 Interpretation of the hierarchy 159
5.4 Diffusion Simulations 161
5.4.1 Decoherence Strength 162
5.4.2 Coin Angle 167
5.5 Advection-Diffusion Simulations 171
5.5.1 Advection-diffusion from a biased initial coin state 171
5.5.2 Advection-diffusion parameter sweeps 174
5.6 Dependence of the Effective Diffusion Coefficient on Initial Conditions . 180
5.7 Encoding and Computational Cost 183
5.8 Discussion 184

6 Discussion and Conclusion **189**
6.1 Summary of Results 189
6.1.1 Continuum limits of lazy open quantum walks 189
6.1.2 Grover and twistless quantum walks 190
6.1.3 The parity echo model 191
6.2 Overarching Discussion and Conclusions 192
6.3 Future Work 195

A Appendix **198**
A.1 Gell Mann Matrices 198
A.2 Numerical Simulations of the Continuum Equation 198
A.3 Full lattice Boltzmann CE expansion to recover the Navier Stokes equation 200

A.4 Classical Lattice Boltzmann Comparison Models 203
A.4.1 Diffusion comparison 205
A.4.2 Advection-diffusion comparison 206
A.5 Full Decoherence Strength Sweep for the Parity Echo Diffusion Simulations 207
A.6 Full Coin Angle Sweep for the Parity Echo Diffusion Simulations 211
A.7 Full Decoherence Strength Sweep for the Parity Echo Advection-Diffusion Simulations 215
A.8 Full Coin Angle Sweep for the Parity Echo Advection-Diffusion Simulations 219

Bibliography **220**

List of Figures

2.1 Probability density of classical random walks after $t = 100$ steps with initial position at $x = 0$. The left image shows the one dimensional random walk, shown as a smooth probability density obtained from an ensemble of walkers. The right image shows the two dimensional random walk, shown as a heatmap. 10

2.2 One dimensional lattice velocity sets for the D1Q2 and D1Q3 lattice Boltzmann models. 18

2.3 The probability distribution of a classical random walk versus a quantum walk. Both walks begin localised at position $x = 0$ and evolve for 50 time steps. The DTQW is evolved using a simple Hadamard coin given by Eqn. (2.66) with $\theta = \pi/4$. 35

2.4 Variance growth for an unbiased classical random walk and a discrete time Hadamard quantum walk on a one dimensional lattice. The variance of the classical walk grows linearly, while the quantum walk variance grows quadratically. 35

2.5 Probability distribution of a one dimensional discrete time quantum walk after $t = 100$ steps, for different choices of coin operator, with a fixed symmetric initial coin state. 37

2.6 Probability distributions of a one dimensional discrete time quantum walk after $t = 100$ steps for combinations of different coin operators (columns) and initial coin states (rows). 40

2.7 Schematic comparison of different continuum limit scalings. 47

2.8 Effect of projective decoherence on a one dimensional Hadamard quantum walk after $t = 100$ timestep for both coin and spatial decoherence. Decoherence strength γ is varied from 0 to 1. 52

3.1 Probability distribution of the three–state lazy DTQW and the standard two–state DTQW after $t = 100$ steps. The standard walk uses an initial state $|0\rangle \otimes \frac{1}{\sqrt{2}}(|L\rangle + i\,|R\rangle)$. The initial state of the lazy walk is an equal magnitude Fourier–symmetric superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, with $\omega = e^{2\pi i/3}$. The lazy walk exhibits ballistic side peaks together with a pronounced central rest state peak, illustrating the three channel structure of the dynamics. 68

3.2 Time evolution of the probability distribution of the standard $\mathcal{SU}(2)$ Hadamard walk versus the $\mathcal{SU}(3)$ lazy Grover quantum walk at times $t = 10, 20, 40, 60, 80, 100$. The standard walk uses an initial state $|0\rangle \otimes \frac{1}{\sqrt{2}}(|L\rangle + i\,|R\rangle)$. The initial state of the lazy walk is an equal magnitude Fourier symmetric superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, with $\omega = e^{2\pi i/3}$. 69

3.3 Effect of dephasing on the lazy DTQW for different coin angles θ and decoherence strengths γ. Panels (a,c) show coin dephasing and panels (b,d) show spatial dephasing, for $\theta = \pi/2$ (top row) and $\theta = \pi$ (bottom row). Coin dephasing has projector choice $\mathbb{P}_L = |L\rangle\langle L|, \mathbb{P}_S = |S\rangle\langle S|, \mathbb{P}_R = |R\rangle\langle R|$, while spatial dephasing uses $\mathbb{P}_x = |x\rangle\langle x| \otimes \mathbb{I}_c$. The initial state is a localised position state with a Fourier symmetric coin superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, where $\omega = e^{2\pi i/3}$. 75

3.4 Discrete time quantum walk versus derived continuum master equation for coin dephasing. Final position probability $P(x,t)$ after $N = 100$ steps of the lazy three state quantum walk with $\theta = \pi/2$, shown for several decoherence rates γ. Solid lines show the solution of the corresponding derived continuum Lindblad equation. The initial state is a spatially delocalised Gaussian wavepacket of width $\sigma_0 = 5$ with a Fourier-symmetric coin superposition $(1, \omega, \omega^2)/\sqrt{3}$, where $\omega = e^{2\pi i/3}$. 82

3.5 Absolute error $|P_{\text{cont}} - P_{\text{disc}}|$ for coin dephasing. 84

3.6 Coin dephasing fraction error for discrete vs continuum models. 85

3.7 DTQWs derived continuum master equation for spatial dephasing. Final position probability $P(x,t)$ after $N = 100$ steps of the lazy three state quantum walk with $\theta = \pi/2$, shown for several decoherence rates γ. Solid lines show the solution of the corresponding derived continuum Lindblad equation with spatial dephasing. The initial state is a spatially delocalised Gaussian wavepacket of width $\sigma_0 = 5$ with a Fourier-symmetric coin superposition $(1, \omega, \omega^2)/\sqrt{3}$, where $\omega = e^{2\pi i/3}$. 88

4.1 Comparison between the two state and lazy three state Grover walks with coin dephasing, together with their matched D1Q3 lattice Boltzmann diffusion profiles. Both models were simulated on a line of length $N = 101$ for $T = 150$ timesteps, with dephasing strength $\gamma = 0.2$ and initial Gaussian width $\sigma_0 = 10$. 116

4.2 Comparison of decohered twistless two state and lazy three state quantum walks against fitted classical lattice Boltzmann models. The two state model is compared with a D1Q2 LBM, while the lazy three state model is compared with a D1Q3 LBM. In both cases the walks are evolved on a lattice of size $N = 101$ for $T = 150$ timesteps with dephasing strength $\gamma = 0.2$ and initial Gaussian width $\sigma_0 = 10$. 135

5.1 Comparison of the unitary parity echo lazy walk and the standard lazy Grover walk for the complex symmetric coin state coin state $\chi_{\text{cs}} = (1, 1, i)^T/\sqrt{3}$. The three panels show the final probability density, the variance growth, and the standard deviation growth. 145

5.2 Comparison of the unitary parity echo lazy walk and the standard lazy Grover walk for the unbiased initial coin state $\chi_{\text{sym}} = (1, 1, 1)^T/\sqrt{3}$. The three panels show the final probability density, the variance growth, and the standard deviation growth. 147

5.3 Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths are $\gamma = 0$, $\gamma = 0.1$, $\gamma = 0.3$, and $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$, while the initial coin state is defined in Eq. (5.74). In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve. 163

5.4 Inverse dephasing fit for the effective diffusion coefficient. The fitted diffusion coefficients are compared with a curve of the form $D_{\text{eff}}(\gamma) = A/\gamma + B$. The $\gamma = 0$ point is excluded from the fit because the overdamped coherence closure requires a nonzero dephasing rate. 166

5.5 Probability density snapshots from the θ sweep for the unbiased parity echo quantum walk at fixed $\gamma = 0.3$. The quantum walk density is compared with the matched D1Q3 LBM diffusion solution at $t = 0$, $t = 90$, and $t = 180$. The rows show $\theta = 0.05$, $\theta = 0.2$, $\theta = 0.35$, and $\theta = 0.7$. The spatial initial state is a Gaussian of width $\sigma_0 = 10$, while the unbiased initial coin state is defined in Eq. (5.74). 168

5.6 Quadratic fit for the dependence of the effective diffusion coefficient on the parity echo coin angle. The fitted diffusion coefficients are compared with curves of the form $D_{\text{eff}}(\theta) = A\theta^2 + B$. The small to moderate angle fit supports the CE prediction $D_{\text{eff}} \propto \theta^2$, while the larger-angle data show deviations from the leading perturbative scaling. 171

5.7 Comparison between the biased parity echo quantum walk and the matched D1Q3 lattice Boltzmann advection-diffusion model at selected time steps. The quantum walk is evolved using the direct density matrix dephasing channel. 172

5.8 Variance growth of the biased parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model. 173

5.9 Mean position of the biased parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model. 174

5.10 Inverse dephasing fit for the effective diffusion coefficient for the advection diffusion simulation. The fitted diffusion coefficients are compared with a curve of the form $D_{\text{eff}}(\gamma) = A/\gamma + B$. The $\gamma = 0$ point is excluded from the fit because the overdamped coherence closure requires a nonzero dephasing rate. 175

5.11 Fitted drift velocity for the biased parity echo advection-diffusion simulations as a function of the dephasing strength γ. The velocity is extracted from a late time linear fit to the mean position, $\langle x(t)\rangle \simeq \langle x(0)\rangle + vt$. The $\gamma = 0$ point corresponds to the fully coherent walk and is more affected by coherent oscillations, while for nonzero dephasing the fitted drift velocity increases and begins to level off at larger γ. 176

5.12 Quadratic fit for the dependence of the effective diffusion coefficient on the parity echo coin angle for the advection diffusion model. The fitted diffusion coefficients are compared with curves of the form $D_{\text{eff}}(\theta) = A\theta^2 + B$. The small to moderate angle fit supports the CE prediction $D_{\text{eff}} \propto \theta^2$, while the larger angle data show deviations from the leading perturbative scaling. 177

5.13 Fitted drift velocity for the biased parity echo advection-diffusion simulations as a function of the coin angle θ. The decrease at the largest value of θ should be interpreted with care, since the packet spreads strongly and finite-size or boundary effects can contaminate the late-time mean-position fit. 178

5.14 Variance linearity diagnostic for the biased parity echo advection-diffusion θ sweep. 179

5.15 Effective diffusion coefficient D_{eff} as a function of the dephasing strength γ for different unbiased initial coin states in the parity echo model. The diffusion coefficient is extracted from the linear growth of the spatial variance, with $\theta = 0.25$ and $\sigma_0 = 10$. 181

5.16 Effective diffusion coefficient D_{eff} as a function of the dephasing strength γ for different strongly biased initial coin states in the parity-echo model. The diffusion coefficient is extracted from the linear growth of the spatial variance, with $\theta = 0.25$ and $\sigma_0 = 10$. 181

5.17 Parameter sweep of the effective diffusion coefficient D_{eff} in the parity echo model as a function of the coin angle θ and dephasing strength γ, for the initial coin state $\chi = (1, i, 1)^T/\sqrt{3}$. The effective diffusion coefficient is extracted from a linear fit to the variance growth over the chosen fitting interval. 182

A.1 Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths range from $\gamma = 0$ to $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve. 209

A.2 Variance growth of the parity echo quantum walk for different dephasing strengths γ. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. The high linear fit quality across the sweep indicates approximately diffusive second moment scaling, although the full density snapshots are still needed to assess whether the probability profile is close to classical diffusion. 210

A.3 Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing coin angle θ. From top to bottom, the coin angles range from $\theta = 0.05$ to $\theta = 0.9$. The decoherence rate is fixed at $\gamma = 0.3$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve. 213

A.4 Variance growth of the parity echo quantum walk for different coin angles θ, with $\gamma = 0.3$ and $\sigma_0 = 10$. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. Lower values of θ appear to give a higher linear fit, and extremely high angles seem to lead to sub-diffusive scaling. 214

A.5 Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths range from $\gamma = 0$ to $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve. 216

A.6 Variance growth of the parity echo quantum walk for different dephasing strengths γ for the advection-diffusion model. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. The high linear fit quality across the sweep indicates approximately diffusive second moment scaling, although the full density snapshots are still needed to assess whether the probability profile is close to classical diffusion. 218

A.7 Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model for increasing coin angle θ. From top to bottom, the coin angles range from $\theta = 0.05$ to $\theta = 0.9$. The decoherence rate is fixed at $\gamma = 0.3$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve. 221

A.8 Variance growth of the parity echo quantum walk for different coin angles θ, with $\gamma = 0.3$ and $\sigma_0 = 10$. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. Lower values of θ appear to give a higher linear fit, and extremely high angles seem to lead to sub-diffusive scaling. 222

List of Tables

A.1 Numerical parameters used in the discrete quantum walk and continuum RK4 simulation. 201

Preface/Acknowledgements

I would first like to give my biggest thanks to my supervisor, Professor Viv Kendon. I feel incredibly lucky to have had such a supportive, caring, insightful and hardworking supervisor throughout my PhD. She has supported me in so many different ways, both academically and personally, and I cannot thank her enough for everything she has done for me. I am so grateful for all of the knowledge, guidance and opportunities she has shared with me throughout this entire process. If I am fortunate enough to continue in academia, I hope to become the kind of researcher and supervisor that she has inspired me to be.

I would also like to thank the Zombies in Room 733 for all of the support, lunchtime conversations, distractions and many, many laughs. You made the PhD experience much more enjoyable! I would especially like to thank Kathryn McInroy and Dr Steph Foulds for all of their help, encouragement and support with my work. I really appreciate everything you have done for me.

I would like to thank Dr Helen Vaughan for being so kind and supportive, particularly while I was writing this thesis.

I would like to thank Lynsay, Lewis, everyone from Flat 11, Nel and all of my skateboarding friends for the support, chaos, distractions, and good times throughout my PhD. You kept me laughing, gave me plenty of reasons to step away from my laptop, and made the whole process far more fun.

I would like to thank Beth Molloy. Having grown up together with a shared love of science, I am incredibly grateful that we have always encouraged one another to pursue our interests in STEM. I would also like to thank Kirsten Cooper, the teacher who first sparked my love of physics and inspired me to pursue it further.

I would like to thank my parents and my brother for their constant love, support and encouragement throughout my PhD. Thank you for always believing in me, taking an interest in my work, attempting to understand my publications and encouraging me to continue with academic life. I am extremely lucky to have such a loving and supportive family, and I could not have done this without you!

I would also like to give a very special thank you to my wonderful Grandma, who was always there to answer the phone for a chat, listen to me and offer support and encouragement whenever I needed it. This meant more to me than I can properly express throughout the PhD, and I have so much love and gratitude for her. I would also like to remember my Grandad, who had such a special love for physics and was always so encouraging of my career. I wish he could have seen me finish this thesis, and I hope he would be proud.

Finally, I would like to thank Fin for always supporting me in every way he can, for making my life easier, for being my biggest cheerleader and for always knowing how to make me laugh. Thank you for being there for me, your support has meant the world to me.

Chapter 0. Preface/Acknowledgements

Chapter 1

Introduction

1.1 Motivation and Overview

Richard Feynman famously remarked, "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy" [Feynman, 1982]. Feynman was primarily referring to the simulation of quantum systems, whose Hilbert spaces grow exponentially and whose dynamics are often extremely difficult to reproduce efficiently on a classical computer. This observation helped motivate the modern field of quantum computing and quantum simulation, where the central idea is that quantum systems may be the most natural tools for simulating other quantum systems. However, Feynman's remark also raises a broader question. If nature is ultimately quantum mechanical, can quantum computers also provide useful ways of simulating the classical equations that emerge from quantum and statistical physics? Many physical systems of practical importance are described not directly by microscopic dynamics, but by macroscopic partial differential equations, such as those governing fluid flow, diffusion, transport, and wave propagation [Ronch et al., 2012]. The accurate numerical solution of these equations forms the foundation of computational physics and engineering. This motivates the question explored in this thesis: whether quantum walk dynamics, augmented by controlled decoherence, can be used to reproduce classical transport behaviour in a way that remains quantum mechanical.

Computational fluid dynamics (CFD) plays a crucial role in modelling systems across a wide range of scales, from aerodynamic design to climate modelling [Mani and Dorgan, 2023, Malinverno et al., 2025, Elsharkawy et al., 2025]. Generally, CFD simulations are often computationally intensive, requiring fine spatial and temporal resolution, large scale discretisations, and significant computational resources. Once such CFD method of interest is the lattice Boltzmann method (LBM), which will be central to this thesis. The lattice Boltzmann method is particularly relevant here because it is already a mesoscopic formulation of fluid dynamics. Rather than discretising the macroscopic PDEs directly, LBM evolves local populations associated with discrete lattice velocities through repeated streaming and collision steps. Many PDEs of interest, such as the Navier Stokes equations, are nonlinear and do not admit general analytical solutions for arbitrary flows, making numerical simulation the only viable approach.

This motivates the search for alternative computational frameworks that may offer improved efficiency, such as quantum computation. Aside from providing the opening quote of this thesis (amongst other things), Feynman also proposed the idea of harnessing quantum systems to simulate physical processes, an idea that later inspired the development of quantum walks [Feynman and Hibbs, 1965]. Originally introduced as the quantum analogue of classical random walks, quantum walks provide a natural and flexible model for constructing quantum algorithms and simulators [Ambainis, 2004]. Quantum walks exhibit rich dynamical behaviour arising from superposition and interference. The LBM structure is closely related to the structure of a discrete time quantum walk, where a shift operation transports amplitudes along lattice directions and a coin operation mixes the internal velocity states. Previous work has therefore drawn connections between quantum walks, quantum lattice gas methods, and quantum lattice Boltzmann schemes [Succi et al., 2015, Meyer, 1996, Wawrzyniak et al., 2025, Zamora et al., 2026]. Since LBM is both practically useful for CFD and structurally similar to lattice based quantum dynamics, it provides a natural classical reference point and source of inspiration for the models developed in this thesis.

The central aim of this thesis is to explore how discrete time quantum walks with controlled decoherence can be used as kinetic models for macroscopic transport equa-

tions such as the advection diffusion equation. In particular, I investigate the extent to which quantum walk dynamics can reproduce key features of lattice Boltzmann methods, such as transport, relaxation, and diffusion, while remaining quantum mechanical. This perspective provides a bridge between quantum dynamics and classical continuum models, and offers a novel approach to the simulation of transport phenomena.

1.2 Thesis Outline

This thesis is structured as follows. First, I introduce the necessary background in computational fluid dynamics, the lattice Boltzmann method, quantum computing, and discrete time quantum walks. This establishes both the classical simulation problem and the motivation and tools for exploring alternative computational architectures. I then review existing work on quantum algorithms for fluids, quantum lattice gas and quantum lattice Boltzmann approaches. Chapter 3 then develops the continuum limit framework used throughout the thesis. In this chapter, I derive continuous spacetime limits for lazy discrete time quantum walks with both coin and spatial decoherence. This extends the standard unitary continuum limit analysis by showing how controlled decoherence enters the effective large scale equations. The lazy structure provides a foundation for the rest state models introduced in later chapters. These limits provide the starting point for Chapter 4, where a Chapman-Enskog style expansion is applied to decohered quantum walk models in order to assess whether they can reproduce diffusive dynamics. In this chapter, both simple Grover type walks and modified stroboscopic models are explored. The aim is to identify the structural ingredients required for a quantum walk to simulate diffusion in a way that mirrors the relaxation mechanism of classical lattice Boltzmann methods. The models are assessed both analytically, through their continuum hierarchies, and numerically, by comparing their density profiles and variance growth with matched classical LBM simulations. Chapter 5 introduces the parity echo model. The Chapman-Enskog expansion is applied to this model, and it is shown to reproduce diffusion more effectively by producing a more useful ordering of relaxation, coherent transport, and diffusive terms. Numerical simulations are then used to demonstrate both diffusion and advection-diffusion behaviour,

with comparisons made against classical lattice Boltzmann simulations. Finally, in Chapter 6, I summarise the main findings of the thesis and discuss the general outlook for using quantum walks and decohered quantum dynamics as models for future quantum fluid simulation algorithms.

Chapter 2

Background

This chapter introduces the classical and quantum background required for the rest of the thesis. I begin with the classical setting: high performance computing, computational fluid dynamics, random walks, diffusion, advection-diffusion, and the lattice Boltzmann method. This establishes the transport problems that motivate the work and demonstrates why LBM provides a useful kinetic reference model. I then introduce the Chapman-Enskog expansion, which connects microscopic lattice updates to macroscopic PDEs and will later be adapted to the quantum walk setting.

The second part of the chapter introduces the quantum framework. I review the basic principles of quantum computing before introducing discrete time quantum walks as the quantum analogue of classical random walks. I describe the role of the coin, shift and initial state. I then introduce the continuous spacetime limit of quantum walks, which again connects the discrete model to the macroscopic behaviour of the system. Finally, I introduce decoherence, density matrices, and open quantum walk dynamics, before reviewing existing work on quantum algorithms for fluids, quantum lattice gas methods, and quantum lattice Boltzmann approaches. Together, these sections provide the background needed to develop decohered quantum walks as potential models for diffusion and advection-diffusion in the following chapters.

2.1 Classical Computational fluid dynamics and the lattice Boltzmann methods

2.1.1 High Performance Computing

High performance computing (HPC) refers to the use of supercomputers and large scale computing clusters to solve computationally intensive problems ([Barral et al., 2025]. These systems leverage parallel and distributed computing architectures, allowing many processors to operate simultaneously and enabling the simulation of large scale physical systems that would be intractable on conventional hardware [Elsharkawy et al., 2025]. Over the past two decades, HPC has evolved from specialised supercomputing platforms to more flexible and scalable cluster and grid based systems, significantly expanding its accessibility and application domain.

HPC underpins a wide range of scientific and engineering applications. These include aerodynamic simulations for the design of aircraft and automobiles, climate and weather modelling and molecular modelling for drug discovery and materials science [Mani and Dorgan, 2023, Malinverno et al., 2025, Liu et al., 2016].

A significant proportion of HPC workloads are devoted to the numerical solution of partial differential equations (PDEs), which form the mathematical foundation for modelling physical phenomena [Galizia et al., 2022]. In particular, computational fluid dynamics (CFD) relies heavily on solving PDEs to describe the behaviour of fluids in motion [Tennie, 2025]. Applications range from simulating airflow over wind turbines and aircraft wings, to modelling subsurface flows in geophysics [Ghasemian et al., 2017, Ronch et al., 2012, Huyakorn, 2012]. These problems are typically high dimensional, nonlinear, and time dependent, requiring fine spatial and temporal discretisation. As a result, they demand substantial memory and computational resources, making them ideally suited to parallel execution on HPC systems.

At the core of many CFD problems lies the Navier Stokes equation, which governs the dynamics of fluid flow. In its incompressible form, the equation may be written as

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\frac{1}{\rho_{\mathrm{LB}}}\nabla p + \nu \nabla^2 \mathbf{u}, \tag{2.1}$$

where $\mathbf{u}$ is the fluid velocity, p is the pressure, ρ_{LB} is the density, and ν is the kinematic viscosity [Krüger et al., 2016]. The nonlinear advective term $(\mathbf{u}\cdot\nabla)\mathbf{u}$, together with the coupling between velocity and pressure, makes the Navier Stokes equation particularly challenging to solve. In general, closed form analytical solutions of the Navier Stokes equations are not available for most physically relevant scenarios, and numerical simulation is therefore required. A variety of computational techniques have been developed for this purpose, including finite difference, finite volume, and finite element methods [Usikalu, 2025, Li et al., 2025].

Despite the continued success of high performance computing, further gains in computational power are increasingly constrained by the physical limits of conventional silicon based hardware. For several decades, improvements in performance were driven by Moore's law and Dennard scaling, which enabled higher transistor densities and clock speeds [Moore, 1965]. However, these trends have slowed significantly due to power dissipation, heat generation, and quantum scale effects in nanoscale devices. As a result, simply increasing clock frequency or packing more transistors onto a chip no longer yields proportional performance improvements [Tennie, 2025]. While advances in parallel architectures and specialised accelerators, such as GPUs, have extended the capabilities of classical computing, they do not fundamentally overcome these scaling limitations. This has motivated the exploration of alternative computational paradigms, including quantum computing, which leverages fundamentally different physical principles, such as superposition and interference, to process information and may offer new approaches to the simulation of complex dynamical systems [Barral et al., 2025]. In particular, quantum computing offers a promising framework for the simulation of complex physical systems, including those governed by partial differential equations, where classical methods face significant computational challenges [Tosti Balducci et al., 2022].

Within this context, the simulation of PDEs provides a natural testing ground for quantum algorithms. Fluid dynamics, and in particular the Navier Stokes equations, represents an important long term goal for quantum computational fluid dynamics because of its central role in engineering, climate modelling, and applied physics. However,

directly targeting fully nonlinear Navier Stokes dynamics is extremely challenging, both mathematically and algorithmically [Ray et al., 2019]. This thesis therefore focuses on simpler transport equations, primarily diffusion and advection-diffusion, as benchmark problems. These equations retain the essential features of transport, spreading, and drift, while providing a more controlled setting in which to test whether quantum walk dynamics can reproduce the desired classical continuum behaviour.

Before introducing the quantum models studied in this thesis, it is essential to first review the corresponding classical lattice based picture of transport. The simplest example is the classical random walk, where local stochastic motion on a lattice gives rise, after averaging over many walkers, to macroscopic diffusion [Schwarz, 2022]. By introducing a directional bias, the same picture can be extended to advection-diffusion, where the distribution both drifts and spreads. These processes provide the basic classical transport behaviour that this thesis will later attempt to reproduce using quantum walk dynamics. The lattice Boltzmann method is then formally introduced as a more structured kinetic framework for simulating transport PDEs. Although LBM is commonly associated with the recovery of the Navier Stokes equations, it can also be used to model simpler equations such as diffusion and advection-diffusion [Guo et al., 1999]. In this thesis, LBM is used as a classical source of inspiration and comparison.

2.1.2 Classical Random Walks

The classical random walk is a stochastic process that underpins a wide range of phenomena across the natural and applied sciences [Xia et al., 2019]. It describes the motion of a walker undergoing random steps on a line, lattice, or graph, and appears in contexts such as Brownian motion, diffusion processes, and search dynamics [Kac, 1947, Schwarz, 2022]. Random walks have applications across engineering and many scientific fields, including ecology, psychology, computer science, physics, chemistry, biology, economics, and sociology [Codling et al., 2008, Blurton et al., 2020, Weiss, 1983, Ficara et al., 2022]. These classical processes also form the foundation of many numerical methods for modelling transport phenomena, specifically diffusion and advection-diffusion, which will later be discussed in the context of lattice Boltzmann methods.

In its simplest form, a classical random walk takes place on a one dimensional lattice. The walker occupies a lattice site and is equipped with a binary coin which determines the direction of motion at each timestep. For example, heads may correspond to moving one lattice site to the left, while tails corresponds to moving one lattice site to the right. Repeating this process many times produces a probability distribution over possible walker positions. In higher dimensional walks, the coin may have more possible outcomes, allowing the walker to move along different directions of a graph or lattice [Lawler and Limic, 2010]. The coin may also be biased, meaning that some directions are chosen with higher probability than others. A biased coin produces a net drift in the direction favoured by the coin, while an unbiased coin produces no preferred direction on average [Codling et al., 2010].

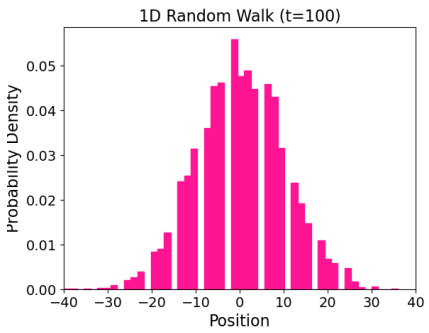

(a) Probability distribution of a one dimensional classical random walk.

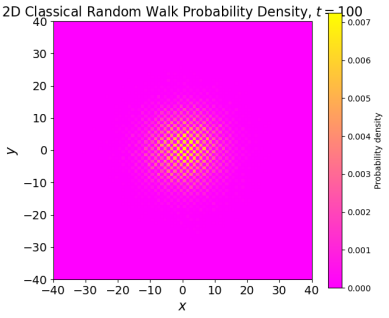

(b) Probability distribution of a two dimensional classical random walk. The gaps in the plot are a consequence of binning and not due to the walk itself.

Figure 2.1: Probability density of classical random walks after $t = 100$ steps with initial position at $x = 0$. The left image shows the one dimensional random walk, shown as a smooth probability density obtained from an ensemble of walkers. The right image shows the two dimensional random walk, shown as a heatmap.

More explicitly, consider an unbiased classical random walk on a one dimensional lattice with spatial lattice spacing $\Delta_x > 0$ and time step Δ_t. Let $P(x, t)$ denote the probability of finding the walker at position x at time t. If the walker moves left or right with equal probability $1/2$, then the probability of being at x at the next time step is obtained from walkers that were previously at the neighbouring sites $x - \Delta_x$

and $x + \Delta_x$ [Xia et al., 2019]. The update rule is therefore

$$P(x, t + \Delta_t) = \frac{1}{2}P(x - \Delta_x, t) + \frac{1}{2}P(x + \Delta_x, t). \tag{2.2}$$

Here, the first term represents walkers that were one lattice site to the right and moved left into position x, while the second term represents walkers that were one lattice site to the left and moved right into position x. Because the coin is unbiased, both contributions have equal weight. In the continuum limit, this update rule leads to the diffusion equation,

$$\partial_t P(x, t) = D\, \partial_x^2 P(x, t), \tag{2.3}$$

where $P(x, t)$ is now the continuous probability density and D is the diffusion coefficient [Wojnar, 2013]. The diffusion coefficient controls how quickly the distribution spreads. For the simple unbiased random walk, it is related to the lattice spacing and timestep by $D = \Delta_x^2/(2\Delta_t)$, in the appropriate continuum scaling.

Diffusion describes the spreading of a distribution due to random, uncorrelated microscopic motion [Schwarz, 2022]. A particle undergoing diffusion does not move persistently in one direction. Instead, each step partly erases the memory of the previous direction of motion. As a result, the mean position of an unbiased walk remains zero, $\langle x(t) \rangle = 0$, while the variance grows linearly in time,

$$\mathrm{Var}(x(t)) \propto t. \tag{2.4}$$

Equivalently, the width of the distribution grows like $\sqrt{t}$ [Xia et al., 2019]. The Gaussian shape of the long time distribution arises mathematically from the central limit theorem, since the final position of the walker is the sum of many independent random steps [Shlesinger, 1974]. Physically, the Gaussian profile reflects the loss of directional information over time. After many random steps, the walker is most likely to be found near its initial position, but with a spread that increases gradually as time passes. Figures 2.1 show the probability density of classical random walks after $t = 100$ steps in both the one and two dimensional cases. The one dimensional distribution approaches

a Gaussian profile as a consequence of the central limit theorem [Lawler and Limic, 2010]. The right image shows the two dimensional random walk, shown as a heatmap. The distribution is isotropic and approximately Gaussian in both spatial directions. In both cases, the spreading is diffusive, with the variance growing linearly in time. These characteristic markers of diffusive behaviour are what I later aim to recreate using quantum walk models.

Advection-Diffusion

A closely related process is advection-diffusion. In this case the motion is split into two effects, a deterministic advective drift and a stochastic diffusive displacement. The drift moves the centre of the distribution in a preferred direction, while diffusion causes the distribution to spread around this moving centre [Pierrehumbert, 2000]. At the level of many random walkers, the ensemble probability density converges to the macroscopic advection-diffusion equation,

$$\partial_t P(x,t) + u\,\partial_x P(x,t) = D\,\partial_x^2 P(x,t), \tag{2.5}$$

where $P(x,t)$ is the probability density, u is the advection velocity, and D is the diffusion coefficient [Mojtabi and Deville, 2015]. In the random walk picture, a biased coin can produce the advective drift, while the remaining random component of the motion produces diffusion. This behaviour is central to this work. The aim is not only to produce spreading that visually resembles diffusion, but to reproduce the characteristic macroscopic signatures of diffusion and advection-diffusion. In a purely diffusive process, the centre of the distribution remains fixed while the width grows in time. In an advection-diffusion process, the centre of the distribution moves approximately linearly with time, while the distribution also spreads diffusively around that moving centre [Vlahos et al., 2008]. These features provide simple quantitative tests for whether a microscopic lattice model is reproducing the desired continuum transport behaviour. Diffusion and advection-diffusion therefore provide simple but important examples of classical transport processes, and they form a natural benchmark to recreate with quan-

tum walk models before moving on to more complicated fluid dynamics. If this can be achieved in a computationally useful way, quantum walk models may offer advantages in terms of the resources required to represent and evolve the system, potentially reducing the number of qubits, operations, or measurements needed compared with more direct quantum implementations of fluid models. This could, in turn, make such algorithms less computationally and energetically demanding. Any potential quantum advantage would nevertheless depend strongly on the encoding, implementation, and hardware used. The precise meaning of quantum advantage in this context, and the conditions under which it may be retained, will be discussed in considerably more detail later in this thesis.

Mean Position and Variance as Transport Diagnostics

To quantify the spreading behaviour of a walk, it is useful to track the moments of the probability distribution. The mean position is defined as

$$\langle x(t) \rangle = \sum_x x P(x,t), \tag{2.6}$$

where $P(x,t)$ is the probability of finding the walker at position x at time t. The variance is defined as

$$\mathrm{Var}(x(t)) = \sum_x \left(x - \langle x(t) \rangle\right)^2 P(x,t). \tag{2.7}$$

The mean position measures the drift of the distribution, while the variance measures its spatial spread [Brun et al., 2003a]. For an unbiased classical random walk, there is no preferred direction, and the mean position remains fixed,

$$\langle x(t) \rangle = 0. \tag{2.8}$$

The variance grows linearly in time,

$$\mathrm{Var}(x(t)) \propto t, \tag{2.9}$$

which is again the characteristic signature of diffusion [Schwarz, 2022]). Equivalently, the width of the probability distribution grows like $\sqrt{t}$. For an advection-diffusion process with constant drift velocity u, the mean position instead grows linearly,

$$\langle x(t) \rangle \simeq ut, \tag{2.10}$$

while the variance still grows linearly,

$$\mathrm{Var}(x(t)) \simeq 2Dt. \tag{2.11}$$

Thus, the mean position can be used to estimate the advective velocity, while the variance can be used to estimate the diffusion coefficient.

These moment diagnostics will be used throughout this thesis. When comparing quantum walk models with classical diffusion or advection-diffusion, the variance provides a simple way to distinguish diffusive spreading from ballistic spreading. A classical diffusive process has $\mathrm{Var}(x(t)) \propto t$, whereas a coherent quantum walk often exhibits ballistic spreading, with $\mathrm{Var}(x(t)) \propto t^2$ [Brun et al., 2003a]. Therefore, recovering approximately linear variance growth is one of the main indicators that the quantum walk is simulating a diffusive regime. Similarly, in the advection-diffusion case, a linear mean position together with a linear variance provides evidence that the model is reproducing both drift and diffusion at the macroscopic level.

The classical random walk therefore provides the simplest microscopic picture of diffusion and advection-diffusion. The individual walkers undergo local stochastic motion, and the macroscopic transport equation emerges after averaging over many such trajectories. However, numerical methods for PDE simulation often do not track individual random walkers directly. Instead, they evolve distribution functions or populations associated with discrete lattice velocities, meaning more complex systems may be simulated and analysed. This is the methodology used in the lattice Boltzmann method. In this sense, LBM can be viewed as a more structured generalisation of random walk transport, in which local lattice updates are designed so that the desired macroscopic PDE is recovered through moment equations and continuum expansions. The following

subsection formally introduces the lattice Boltzmann method and demonstrates how it may be used to simulate advection diffusion.

2.1.3 The lattice Boltzmann method

Rather than following individual stochastic walkers, the LBM evolves populations moving along discrete lattice velocities, and uses collision and streaming steps to recover macroscopic continuum equations [Krüger et al., 2016]. The lattice Boltzmann method originates from lattice gas automata (LGA) models developed in the 1980s, most notably the work of Frisch et al. [1986], who introduced a discrete particle model capable of reproducing fluid behaviour, notably the Navier Stokes equation. While these early models successfully captured hydrodynamic phenomena, they suffered from statistical noise and limited numerical efficiency. To address these issues, the lattice Boltzmann method was introduced as a continuum analogue, replacing discrete particle occupation numbers with particle distribution functions. This reformulation, developed in the late 1980s and early 1990s by McNamara and Zanetti [1988] and further advanced by Succi et al. [1991], significantly improved numerical stability and efficiency. The modern LBM framework, particularly with the introduction of the Bhatnagar-Gross-Krook (BGK) collision operator, provided a simple yet powerful kinetic scheme in which macroscopic fluid equations emerge from local streaming and relaxation dynamics [Bhatnagar et al., 1954]. Since then, the LBM has become a widely used tool in computational fluid dynamics, valued for its simplicity, parallel scalability, and natural suitability for complex boundary conditions and multiscale flows [Aidun and Clausen, 2010].

Rather than discretising macroscopic equations such as the Navier Stokes equations directly, the LBM evolves a set of particle distribution functions on a lattice, whose low order moments reproduce the desired continuum dynamics [He and Luo, 1997]. This kinetic formulation provides the method with a natural separation between transport and relaxation processes. The locality makes it particularly well suited to parallel computation and multiscale modelling, as only a single lattice width around domains has to be swapped with neighbouring processors. In this work, the LBM is not introduced as a numerical scheme to be implemented directly, but rather as a conceptual and

structural template. The aim is to reproduce its key kinetic features, namely discrete velocity channels, streaming, relaxation toward equilibrium, and controlled dissipation within a quantum walk framework. Although the lattice Boltzmann method is most commonly associated with the numerical recovery of the Navier Stokes equations, it can also be used, or adapted, to reproduce a wider class of PDEs, including diffusion and advection-diffusion equations [Krüger et al., 2016]. The quantum walk models developed in this thesis are therefore LBM inspired rather than direct quantum versions of the full Navier Stokes solver. The aim is not yet to reproduce full nonlinear fluid dynamics, but to take a more controlled first step to investigate whether quantum walk dynamics, together with suitable coins, scalings, and decoherence, can reproduce the simpler transport equations such as diffusion and advection-diffusion, in a similar way to classical LBM.

The LBM evolves particle distribution functions on a discrete lattice through successive streaming and collision steps. Both the timestep and lattice spacing are discrete. However, a hierarchical continuous limit may be taken using the Chapman-Enskog (CE) expansion to determine the PDE being simulated by the mesoscopic model [Alexander et al., 1993, 1992]. Using the CE expansion, the discrete dynamics can be systematically analysed across multiple timescales, allowing the recovery of macroscopic partial differential equations from the underlying kinetic model. In particular, the leading order dynamics capture local equilibrium, while higher order corrections generate advective and diffusive behaviour, yielding equations such as the advection-diffusion or Navier Stokes equations.

The discrete time lattice Boltzmann equation governs the evolution of a set of particle populations $\{f_i\}$ associated with a finite number of discrete velocity channels. It may be written as

$$f_i(x + \mathbf{c}_i\Delta_t,\, t + \Delta_t) - f_i(x,t) = \Omega_i(f), \tag{2.12}$$

where x and t denote space and time, Δ_t is the discrete time step, and $f_i(x,t)$ represents the population associated with the ith discrete lattice velocity $\mathbf{c}_i$ [Krüger et al., 2016]. The left hand side of Eq. (2.12) describes the streaming process, in which populations

are advected along their corresponding velocity directions $\mathbf{c}_i$ over one time step. The right hand side encodes local interactions through the collision operator Ω_i, which acts pointwise in space and time.

In most practical implementations, the collision operator is taken to be of BGK form,

$$\Omega_i(f) = -\frac{\Delta_t}{\tau}\left[f_i(x,t) - f_i^{\mathrm{eq}}(x,t)\right], \tag{2.13}$$

where τ is the relaxation time. The BGK operator models local relaxation of the populations toward a prescribed equilibrium distribution f_i^{eq}, with τ controlling the rate at which this relaxation occurs [Qian et al., 1992]. The collision operator is selected depending on the specific PDE being simulated, and this is discussed in further paragraphs. The relaxation timescale τ is the operator that introduces viscosity to the system and is related to the kinematic viscosity ν via:

$$\nu = c_s^2(\tau - \frac{1}{2}). \tag{2.14}$$

Here, c_s is the lattice speed of sound and τ is a scalar number that determines the rate of relaxation. For $0 < \tau < 1$, the pre collision state relaxes to the equilibrium. For $\tau > 1$, it overshoots and for $\tau = 1$, the operator becomes a full projection onto the equilibrium state. Despite its simplicity, the BGK operator captures the essential dissipative physics of kinetic relaxation and gives rise, through a Chapman-Enskog expansion, to diffusive and viscous transport at the macroscopic level.

The discrete lattice speeds $\{\mathbf{c}_i\}$ define the velocity set of the model, and their number and symmetry determine which continuum equations can be recovered. In this work we restrict attention to the one dimensional, three velocity (D1Q3) lattice, for which

$$\mathbf{c}_i \in \{-1,\, 0,\, 1\}, \tag{2.15}$$

and therefore $i = 0, 1, 2$ corresponding to left moving, rest, and right moving populations. In the 2D case, a popular choice of lattice is the D2Q9 lattice, with $i = 0, ..., 8$ [Krüger et al., 2016]. Figure 2.2 shows the one dimensional velocity sets used in the simplest lattice Boltzmann models. In the D1Q2 model, particles move only to the



Figure 2.2: One dimensional lattice velocity sets for the D1Q2 and D1Q3 lattice Boltzmann models.

left or right at each timestep, with discrete velocities $\mathbf{c}_i \in \{-1, +1\}$. In the D1Q3 model, an additional rest population is included. This additional rest state is useful when constructing lattice models for diffusion and advection-diffusion, since it provides more freedom in the equilibrium distribution and in the resulting macroscopic transport coefficients.

The equilibrium distribution that is commonly used to simulate the Navier Stokes equation follows from the second order Hermite expansion of the Maxwell Boltzmann distribution:

$$f_i^{\mathrm{eq}}(x,t) = \omega_i \rho_{\mathrm{LB}} \left[1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{c}_i \cdot \mathbf{u})^2 - c_s^2(\mathbf{u} \cdot \mathbf{u})}{2c_s^4} \right] \tag{2.16}$$

Here, $\rho_{\mathrm{LB}} = \sum_i f_i$ is the macroscopic fluid density, $\mathbf{u}$ is the macroscopic fluid velocity, and $\mathbf{c}_i$ denotes the discrete lattice velocity associated with direction i [Krüger et al., 2016]. The lattice sound speed is denoted by c_s, and ω_i are the lattice weights. The precise form of the equilibrium distribution is model dependent as different choices of f_i^{eq} impose different moment constraints and therefore lead to different macroscopic PDEs in the continuum limit [Gebäck T, 2018]. For example, the second order equilibrium above is used for hydrodynamic models such as Navier Stokes, while simpler linear equilibria can be used to recover diffusion or advection diffusion equations. The lattice weights are denoted by ω_i, which for the D1Q3 example are given as:

$$\mathbf{c}_i \in \{-1, 0, +1\}, \qquad \omega_i = \left\{ \frac{1}{6}, \frac{4}{6}, \frac{1}{6} \right\}, \qquad c_s^2 = \frac{1}{3}. \tag{2.17}$$

Since the BGK collision operator depends on f_i^{eq}, the collision depends on the equilib-

rium distribution, which itself depends on the flow velocity **u** both linearly and quadratically as seen in Eqn. (2.16). This quadratic dependence is the source of nonlinearity in the LBM scheme and the behaviour that is so hard to reproduce in a quantum system. This feature highlights a fundamental difference between classical LBM and quantum walk dynamics. In classical LBM the collision step requires knowledge of local macroscopic moments at every timestep. Implementing such state dependent updates coherently within a quantum algorithm would require measurement and feedback, destroying quantum coherence and eliminating potential advantages.

However, if the equilibrium distribution is chosen to depend only linearly on the conserved moments, then the resulting collision term is linear and can recover linear PDEs such as diffusion or advection-diffusion. These linear transport equations are the focus of this work. I therefore do not attempt to reproduce the full nonlinear Navier Stokes collision structure directly; instead, I use the linear diffusion and advection-diffusion LBM as the classical reference hierarchy against which the decohered quantum walk models are compared.

The Chapman-Enskog Expansion

The connection between discrete LBM kinetic models and the macroscopic equations being simulated is established through the Chapman-Enskog expansion. This multiscale method separates the dynamics into different orders in a small parameter ε, which represents the separation between the microscopic lattice scale and the macroscopic continuum scale [Li, 2015]. At leading order, the system is assumed to be close to local equilibrium. At first order, the transport of this equilibrium distribution gives rise to advective terms, while second order corrections introduce diffusive and viscous effects. In this way, the Chapman-Enskog expansion shows how a simple microscopic update rule can reproduce a macroscopic PDE such as the diffusion, advection-diffusion, or Navier Stokes equations [Chapman, 1970].

This is important because in numerical modelling, one often does not have an exact analytical solution of the PDE available. For many physically relevant PDEs, especially nonlinear or high dimensional systems, the purpose of the numerical method is precisely

to approximate a solution that cannot be written down in closed form. The Chapman-Enskog expansion therefore provides a way of identifying what equation the microscopic model is actually simulating. Rather than only observing the numerical output, one expands the discrete kinetic dynamics and checks which macroscopic PDE emerges in the continuum limit. If the expansion recovers the desired equation, then the lattice model can be interpreted as a valid numerical scheme for that macroscopic process. This mechanism is central to the lattice Boltzmann method and will serve as a guiding framework for constructing quantum analogues. In the later chapters, a similar logic will be applied to quantum models. This is essential for showing that the quantum model is not merely producing visually similar spreading, but is genuinely reproducing diffusion or advection-diffusion behaviour at the level of the limiting PDE.

2.1.4 D1Q3 lattice Boltzmann model and Chapman-Enskog expansion for advection-diffusion

Since this thesis focusses on the simpler transport equations of diffusion and advection-diffusion, the CE expansion will now be applied to the lattice Boltzmann method in the advection-diffusion setting. For completeness, Appendix A.3 gives the corresponding Chapman-Enskog expansions for both the Navier Stokes equation and the diffusion only case.

To make the connection between lattice Boltzmann dynamics and continuum transport explicit, we consider the simplest one dimensional case, the D1Q3 model introduced in Section 2.1.3, specialised to a scalar density on a one dimensional lattice and neglects the higher order velocity dependent error term in the low Mach, constant velocity regime. The derivation in the subsection follows Krüger et al. [2016]. The discrete velocity set is given in Eqn. (2.17). Combining the lattice Boltzmann equation in Eqn. (2.12) with the D1Q3 BGK collision operator in Eqn. (2.13), the update rule becomes

$$f_i(x + \mathbf{c}_i\Delta_t, t + \Delta_t) - f_i(x,t) = -\frac{\Delta_t}{\tau}\left[f_i(x,t) - f_i^{\mathrm{eq}}(x,t)\right]. \tag{2.18}$$

The left hand side describes streaming along the lattice velocity $\mathbf{c}_i$, while the right hand side describes relaxation toward a local equilibrium distribution f_i^{eq}. The param-

eter τ controls the relaxation timescale and therefore determines the effective diffusion coefficient.

For an advection-diffusion model, the conserved macroscopic variable is the scalar density

$$\rho_{\mathrm{LB}}(x,t) = \sum_i f_i(x,t). \tag{2.19}$$

For the D1Q3 lattice, a common choice for modelling linear advection-diffusion is the linear equilibrium distribution

$$f_i^{\mathrm{eq}} = w_i \rho_{\mathrm{LB}} \left(1 + \frac{\mathbf{c}_i u}{c_s^2}\right). \tag{2.20}$$

Since this is a one dimensional D1Q3 model, the dot product $\mathbf{c}_i \cdot \mathbf{u}$ has been written simply as $\mathbf{c}_i u$. This equilibrium is chosen so that the lowest moments reproduce the desired macroscopic fields. In particular,

$$\sum_i f_i^{\mathrm{eq}} = \rho_{\mathrm{LB}}, \tag{2.21}$$

so the equilibrium preserves the density, while

$$\sum_i \mathbf{c}_i f_i^{\mathrm{eq}} = \rho_{\mathrm{LB}} u, \tag{2.22}$$

so the first moment produces an advective flux. The second equilibrium moment is

$$\sum_i \mathbf{c}_i^2 f_i^{\mathrm{eq}} = c_s^2 \rho_{\mathrm{LB}}. \tag{2.23}$$

Thus, the equilibrium distribution fixes the density, the advective flux, and the diffusive scale of the continuum model. In lattice Boltzmann methods, the form of f_i^{eq} is not just a numerical detail but instead determines which macroscopic PDE is recovered in the continuum limit. Changing the equilibrium changes the moment constraints of the kinetic model, and therefore changes the terms that appear in the resulting macroscopic equation. For example, including higher order velocity terms in the equilibrium is necessary when deriving fluid models such as the Navier Stokes equations, whereas the

linear equilibrium above is sufficient for recovering a linear advection-diffusion equation. If the equilibrium distribution is chosen to depend linearly on the scalar density, with the advection velocity treated as an externally imposed field, then the resulting BGK update is linear in the transported scalar population. This is sufficient to recover the linear diffusion or advection-diffusion equation in the continuum limit, with higher order corrections neglected in the CE expansion.

To derive the continuum equation, a Taylor expansion is performed of the left hand side of Eq. (2.18) about (x, t). Defining

$$\mathbb{D}_i \equiv \partial_t + \mathbf{c}_i \partial_x, \tag{2.24}$$

we obtain

$$\Delta_t \mathbb{D}_i f_i + \frac{\Delta_t^2}{2} \mathbb{D}_i^2 f_i = -\frac{\Delta_t}{\tau} \left(f_i - f_i^{\mathrm{eq}} \right) + \mathcal{O}(\Delta_t^3). \tag{2.25}$$

Dividing by Δ_t gives

$$\mathbb{D}_i f_i + \frac{\Delta_t}{2} \mathbb{D}_i^2 f_i = -\frac{1}{\tau} \left(f_i - f_i^{\mathrm{eq}} \right) + \mathcal{O}(\Delta_t^2). \tag{2.26}$$

We now introduce the Chapman-Enskog multiscale expansion,

$$\partial_t = \varepsilon \partial_{t_1} + \varepsilon^2 \partial_{t_2} + \cdots, \qquad \partial_x = \varepsilon \partial_{x_1}, \tag{2.27}$$

and expand the distribution functions as

$$f_i = f_i^{(0)} + \varepsilon f_i^{(1)} + \varepsilon^2 f_i^{(2)} + \cdots. \tag{2.28}$$

The leading order distribution is identified with the local equilibrium:

$$f_i^{(0)} = f_i^{\mathrm{eq}}. \tag{2.29}$$

At order $\mathcal{O}(\varepsilon)$, Eq. (2.26) gives

$$\mathbb{D}_{1i} f_i^{(0)} = -\frac{1}{\tau} f_i^{(1)}, \qquad \mathbb{D}_{1i} \equiv \partial_{t_1} + \mathbf{c}_i \partial_{x_1}. \tag{2.30}$$

Taking the zeroth moment by summing over i, and using conservation of density under collision,

$$\sum_i f_i^{(1)} = 0, \tag{2.31}$$

we find

$$\partial_{t_1} \rho_{\mathrm{LB}} + \partial_{x_1} \sum_i \mathbf{c}_i f_i^{(0)} = 0. \tag{2.32}$$

Using the equilibrium moment relation

$$\sum_i \mathbf{c}_i f_i^{(0)} = \mathbf{u} \rho_{\mathrm{LB}}, \tag{2.33}$$

this becomes

$$\partial_{t_1} \rho_{\mathrm{LB}} + \partial_{x_1} (u \rho_{\mathrm{LB}}) = 0. \tag{2.34}$$

Thus, the first order equation gives the advective part of the macroscopic dynamics. At order $\mathcal{O}(\varepsilon^2)$, the Chapman-Enskog expansion gives

$$\partial_{t_2} f_i^{(0)} + \mathbb{D}_{1i} f_i^{(1)} + \frac{\Delta_t}{2} \mathbb{D}_{1i}^2 f_i^{(0)} = -\frac{1}{\tau} f_i^{(2)}. \tag{2.35}$$

Taking the zeroth moment and using $\sum_i f_i^{(2)} = 0$, we obtain

$$\partial_{t_2} \rho_{\mathrm{LB}} + \partial_{x_1} \sum_i \mathbf{c}_i f_i^{(1)} + \frac{\Delta_t}{2} \sum_i \mathbb{D}_{1i}^2 f_i^{(0)} = 0. \tag{2.36}$$

This is the point where the lower order equations are substituted back into the hierarchy. The first order relation (2.30) gives the first nonequilibrium correction as

$$f_i^{(1)} = -\tau \mathbb{D}_{1i} f_i^{(0)}. \tag{2.37}$$

Therefore the nonequilibrium flux appearing in Eq. (2.36) is

$$\sum_i \mathbf{c}_i f_i^{(1)} = -\tau \sum_i \mathbf{c}_i \left(\partial_{t_1} + \mathbf{c}_i \partial_{x_1} \right) f_i^{(0)}. \tag{2.38}$$

Expanding the right hand side gives

$$\sum_i \mathbf{c}_i f_i^{(1)} = -\tau \left[\partial_{t_1} \sum_i \mathbf{c}_i f_i^{(0)} + \partial_{x_1} \sum_i \mathbf{c}_i^2 f_i^{(0)} \right]. \tag{2.39}$$

Using the equilibrium moments

$$\sum_i \mathbf{c}_i f_i^{(0)} = u\rho_{\mathrm{LB}}, \qquad \sum_i \mathbf{c}_i^2 f_i^{(0)} = c_s^2 \rho_{\mathrm{LB}}, \tag{2.40}$$

we obtain

$$\sum_i \mathbf{c}_i f_i^{(1)} = -\tau \left[\partial_{t_1} (u\rho_{\mathrm{LB}}) + c_s^2 \partial_{x_1} \rho_{\mathrm{LB}} \right]. \tag{2.41}$$

For a constant imposed velocity u, the first order advection Eqn. (2.34) may be used to replace the time derivative,

$$\partial_{t_1} \rho_{\mathrm{LB}} = -u \partial_{x_1} \rho_{\mathrm{LB}}. \tag{2.42}$$

Hence,

$$\partial_{t_1} (u\rho_{\mathrm{LB}}) = -u^2 \partial_{x_1} \rho_{\mathrm{LB}}. \tag{2.43}$$

In the linear advection-diffusion setting, the u^2 contribution is higher order in the low Mach expansion and is neglected. Therefore the leading nonequilibrium flux is

$$\sum_i \mathbf{c}_i f_i^{(1)} = -\tau c_s^2 \partial_{x_1} \rho_{\mathrm{LB}}. \tag{2.44}$$

The Taylor correction in Eq. (2.36) gives the usual lattice correction to this diffusive flux. Using

$$\sum_i \mathbb{D}_{1i}^2 f_i^{(0)} = \sum_i \left(\partial_{t_1} + \mathbf{c}_i \partial_{x_1} \right)^2 f_i^{(0)}, \tag{2.45}$$

and applying the same equilibrium moment relations and first order advection equation, the leading diffusive contribution is

$$\sum_i \mathbb{D}_{1i}^2 f_i^{(0)} = c_s^2 \partial_{x_1}^2 \rho_{\mathrm{LB}}. \tag{2.46}$$

Substituting this result and Eq. (2.44) into Eq. (2.36) gives

$$\partial_{t_2}\rho_{\mathrm{LB}} + \partial_{x_1}\left(-\tau c_s^2 \partial_{x_1}\rho_{\mathrm{LB}}\right) + \frac{\Delta_t}{2} c_s^2 \partial_{x_1}^2 \rho_{\mathrm{LB}} = 0. \tag{2.47}$$

Rearranging,

$$\partial_{t_2}\rho_{\mathrm{LB}} = c_s^2 \left(\tau - \frac{\Delta_t}{2}\right) \partial_{x_1}^2 \rho_{\mathrm{LB}}. \tag{2.48}$$

Thus, the diffusion term is not inserted by hand. It appears by substituting the first order nonequilibrium correction $f_i^{(1)}$ back into the second order equation. This is the key Chapman-Enskog mechanism, with lower order corrections generating higher order macroscopic transport terms. Combining the first and second order contributions and returning to physical variables yields the advection-diffusion equation

$$\partial_t \rho_{\mathrm{LB}} + \partial_x(u\rho_{\mathrm{LB}}) = D_{\mathrm{LB}}\, \partial_x^2 \rho_{\mathrm{LB}}, \tag{2.49}$$

with diffusion coefficient

$$D_{\mathrm{LB}} = c_s^2 \left(\tau - \frac{\Delta_t}{2}\right). \tag{2.50}$$

For constant u, this may be written as

$$\partial_t \rho_{\mathrm{LB}} + u\partial_x \rho_{\mathrm{LB}} = D_{\mathrm{LB}} \partial_x^2 \rho_{\mathrm{LB}}. \tag{2.51}$$

Strictly, the Chapman-Enskog expansion of the linear ADE equilibrium also produces an additional error term of the form

$$E = \partial_x \left[\left(\tau - \frac{\Delta_t}{2}\right) \partial_t(\rho_{\mathrm{LB}} u)\right]. \tag{2.52}$$

For the linear equilibrium this term can be interpreted as a velocity dependent correction to the diffusion coefficient. In the present work I neglect this contribution by working in the low Mach, constant velocity regime, and in the pure diffusion case $u = 0$ it vanishes. Thus the D1Q3 model is used here as a simple linear reference hierarchy for diffusion and advection-diffusion, rather than as a high accuracy ADE solver for arbitrary velocity fields.

The streaming step produces transport along the lattice directions, while relaxation toward equilibrium generates diffusion through the nonequilibrium correction $f_i^{(1)}$. This is the key structural feature that motivates the comparison with DTQWs. Both systems evolve through local lattice updates, and both admit continuum limits in which macroscopic transport equations emerge. The next step is therefore to ask whether an analogous construction can be developed using quantum dynamics to simulate classical diffusion and advection-diffusion. To make this comparison, the following sections first introduce the basics behind quantum computation and quantum walks. I review how quantum states, unitary operators, and measurement enter the quantum computing framework, before introducing discrete time quantum walks as the quantum analogue of classical random walks. This provides the foundation for the later chapters, where quantum walk models are engineered using coins, scalings, and decoherence in order to reproduce diffusion and advection-diffusion behaviour in suitable continuum limits. I then review the current literature on quantum lattice Boltzmann methods and related quantum approaches to computational fluid dynamics, placing the models developed in this thesis within the broader context of quantum algorithms for transport and fluid simulation.

2.2 Quantum Physics and Quantum Computing

2.2.1 Basic Principles of Quantum Mechanics

Quantum mechanics provides the fundamental framework for describing physical systems at atomic and subatomic scales. In classical mechanics, a system is described by a point in phase space, with observables given by real valued functions of position and momentum. In quantum mechanics, the corresponding structure is different. States are represented by vectors in a Hilbert space, observables by self adjoint operators, and closed system dynamics by unitary transformations [McKeever and Nazir, 2026]. Physical symmetries are also represented by unitary transformations on the Hilbert space.

A quantum state may be described by a complex valued wavefunction $\Psi(x,t)$. The

Born rule connects this wavefunction to measurable probabilities: $|\Psi(x,t)|^2$ gives the probability density for finding the system at position (x) [Born, 1926]. The wavefunction must therefore be normalised so that the total probability over all space is equal to one. In the absence of measurement, the state evolves deterministically according to the Schrödinger equation. Measurement, however, is intrinsically probabilistic and projects the system onto an eigenstate of the measured observable.

Two features of quantum mechanics are particularly important for quantum computation. The first is superposition, in which a quantum system can exist in a linear combination of basis states. The second is interference, where probability amplitudes can combine constructively or destructively. Quantum algorithms exploit these effects to amplify desired outcomes and suppress undesired ones. Quantum systems also obey uncertainty relations, meaning that certain observables, such as position and momentum, cannot in general be simultaneously specified with arbitrary precision.

State vectors and bra-ket notation

At a fixed time t, the wavefunction $\Psi(x,t)$ may be interpreted as the components of a vector in an infinite dimensional Hilbert space. In Dirac, or bra ket, notation, the state is written as

$$|\Psi(t)\rangle = \int_{-\infty}^{\infty} \Psi(x,t)\,|x\rangle\,dx \tag{2.53}$$

where $|x\rangle$ denotes the position eigenstate [Bhattacharyya and Guha, 2022]. More generally, a quantum state can be expanded in any chosen basis, with different representations, such as position and momentum space, corresponding to different basis choices. Transformations of quantum states are represented by linear operators. Of particular importance are unitary operators, which describe the time evolution of closed quantum systems. A unitary operator $\hat{\mathbf{U}}$ satisfies

$$\hat{\mathbf{U}}^\dagger\hat{\mathbf{U}} = \hat{\mathbf{U}}\hat{\mathbf{U}}^\dagger = \mathbf{I}, \tag{2.54}$$

so inner products, norms, and probabilities are preserved [Feng et al., 2023]. Throughout this work, $\mathbf{I}$ represents the identity operator. The time evolution of a state may

therefore be written as

$$|\Psi(t)\rangle = \hat{\mathbf{U}}(t)\,|\Psi(0)\rangle\,. \tag{2.55}$$

From a computational perspective, unitary operators act as reversible quantum gates. In this thesis, these ideas are mainly used in finite dimensional lattice models, where the relevant Hilbert space is built from position states and internal coin states.

2.2.2 Quantum Computing

Quantum computing is an emerging field that seeks to use fundamental principles of quantum mechanics, such as superposition and entanglement, to perform computations in ways that may offer advantages over classical approaches. Early developments in this direction include the introduction of the quantum Turing machine [Benioff, 1980], which provided a theoretical framework for computation based upon quantum mechanics. Around the same time, Manin [1980] and Feynman [1982] independently observed that simulating quantum systems on classical computers generally incurs an exponential overhead, suggesting that devices exploiting quantum phenomena could provide significant computational advantages. Subsequent work by Bennett and Brassard [1984] demonstrated the application of quantum mechanics to cryptography, most notably through quantum key distribution.

These foundational contributions helped establish quantum computing as a distinct and rapidly developing field at the interface of physics and computer science. In the following decade, several quantum algorithms were proposed for oracle based problems, including the Bernstein-Vazirani algorithm in 1993 and Simon's algorithm in 1994 [Bennett and Brassard, 1984, Simon, 1994]. Although these algorithms were mainly theoretical, they played an important role in showing that quantum computation could offer more than just a different hardware model. By exploiting superposition and interference, they demonstrated that a quantum computer could extract certain global properties of a function using fewer queries than a classical computer.

A major breakthrough came with Shor's algorithm in 1994, which showed that integer factorisation, an essential problem in modern encryption, could be solved exponentially faster on a quantum computer than by any known classical algorithm [Shor,

1997]. This was followed by Grover's algorithm in 1996, which provided a quadratic speedup for unstructured search problems [Grover, 1996]. In the same year, Lloyd [1996] proved that quantum computers can efficiently simulate quantum systems, thereby confirming Feynman's earlier conjecture and establishing quantum simulation as a natural application of quantum computation.

Despite these theoretical advances, it is important to note that quantum computing remains in its early stages of development. To date, no quantum computer has demonstrated a clear and practical advantage for broadly useful real world problems. Claims of quantum supremacy have often been followed by improved classical algorithms that reduce or eliminate the observed advantage, highlighting the continuing strength and adaptability of classical methods [Arute et al., 2019a, Pednault et al., 2019, Zhong et al., 2020, Oh et al., 2024, Tindall et al., 2024]. Nevertheless, this interplay has been highly productive. The pursuit of quantum algorithms has inspired the development of more efficient classical techniques, sometimes referred to as quantum inspired algorithms. As such, even in the absence of large scale fault tolerant quantum computers, the field has already contributed to an advance in computation and algorithmic efficiency.

In quantum computing, computation is performed through the application of unitary operators to quantum states. These operators act as quantum logic gates, transforming qubits in a reversible manner while preserving the norm of the state. For a single qubit, gates are represented by 2×2 unitary matrices acting on the computational basis $\{|0\rangle, |1\rangle\}$. For example, the Pauli X gate, which acts as a quantum analogue of the classical NOT gate, is given by

$$\hat{\mathbf{X}} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \tag{2.56}$$

and satisfies

$$\hat{\mathbf{X}}\,|0\rangle = |1\rangle\,, \qquad \hat{\mathbf{X}}\,|1\rangle = |0\rangle\,. \tag{2.57}$$

Another important example is the Hadamard gate, defined as

$$\hat{\mathbf{H}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \tag{2.58}$$

which creates superposition states:

$$\hat{\mathbf{H}} |0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \qquad \hat{\mathbf{H}} |1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle). \tag{2.59}$$

More generally, any single qubit unitary transformation can be written as a rotation on the Bloch sphere [Feng et al., 2023]. For example, rotations about the x axis are given by

$$R_x(\theta) = e^{-i\theta\sigma_x/2}, \tag{2.60}$$

where σ_x is the Pauli X matrix. For multiqubit systems, the state space grows exponentially. A system of n qubits is described by a 2^n dimensional Hilbert space, and a general pure state may be written as

$$|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle . \tag{2.61}$$

Here, $|i\rangle$ denotes a computational basis state, corresponding to the binary representation of the integer i. For example, for two qubits the basis states are $|00\rangle , |01\rangle , |10\rangle , |11\rangle$. The complex coefficients α_i are probability amplitudes, and the probability of measuring the basis state $|i\rangle$ is given by $|\alpha_i|^2$. Normalisation of the quantum state requires

$$\sum_{i=0}^{2^n-1} |\alpha_i|^2 = 1. \tag{2.62}$$

Thus, while n classical bits can represent one of 2^n possible bit strings at a time, an n qubit quantum state can exist in a superposition over all 2^n computational basis states. These multiqubit gates allow interactions between qubits. A fundamental example is

the controlled NOT, or CNOT, gate, which acts as

$$\text{CNOT}\,|00\rangle = |00\rangle\,, \quad \text{CNOT}\,|01\rangle = |01\rangle\,, \quad \text{CNOT}\,|10\rangle = |11\rangle\,, \quad \text{CNOT}\,|11\rangle = |10\rangle\,. \tag{2.63}$$

Such gates enable the creation of entanglement, a key property in quantum computation [Wong, 2024]. Together, single qubit and multiqubit gates form a universal set, meaning that any unitary operation can be approximated to arbitrary accuracy using a finite sequence of these gates.

A typical quantum computation consists of three main stages, state preparation, unitary evolution, and measurement. First, the input data must be encoded into a quantum state by a process known as state preparation. The system is then evolved using a sequence of unitary operations, which implement the desired algorithm. Finally, the output is obtained through measurement, and in many cases must be decoded to extract meaningful classical information. Mid circuit measurements may also be useful in some applications, as evolution does not necessarily have to be unitary, especially if ancilla qubits are used. While the unitary evolution can, in principle, be highly efficient, the encoding and decoding stages often present significant practical challenges [Nmaju et al., 2025]. In particular, preparing a quantum state that accurately represents classical data can be computationally expensive, and extracting full information from the quantum state typically requires repeated measurements. As a result, these stages can limit the overall advantage of quantum algorithms, especially for problems involving large amounts of classical data.

Current implementations of quantum hardware include superconducting circuits, ion traps, and cold atoms [Arute et al., 2019b, Ransford, 2026, Radnaev et al., 2025]. Many scientists believe that certain quantum computers may outperform classical computers for a specific subset of problems or tasks, although whether these tasks will lead to broadly useful real world applications remains an open question. Despite their potential advantages, quantum systems are highly sensitive to environmental interactions. These interactions lead to decoherence, which degrades quantum information and introduces noise into the system [Kendon, 2007]. Quantum computing therefore remains at

a relatively early stage of development compared to classical computing, with current hardware and algorithms still limited in scale, reliability, and practical applicability.

Within this broader landscape, quantum walks provide a particularly useful and well studied model for quantum algorithm design. For this reason, I now introduce discrete time quantum walks in more detail, before using them as the central modelling framework for the rest of this thesis.

2.3 Quantum Walks

Quantum walks constitute the quantum analogue of classical random walks, which were introduced in Section 2.1.2. Quantum walks form a versatile framework for modelling quantum transport and computation [Kendon, 2020, Molfetta, 2021]. Unlike their classical counterparts, quantum walks evolve through coherent unitary dynamics and exhibit interference effects. In particular, they do not converge to limiting distributions and due to quantum interference, they may spread significantly faster than their classical equivalents [Childs et al., 2002]. These effects can lead to behaviour that differs significantly from classical random walks, including ballistic spreading, localisation, and highly non trivial transport dynamics depending on the structure of the walk. They can also lead to quadratic speedups in hitting and mixing times, underpinning their use in quantum algorithms for search and graph analysis [Shenvi et al., 2003a, Ambainis, 2004, 2007]. Childs [2009a] showed that a quantum walk can be viewed as a universal computational primitive, in the sense that a desired quantum computation can be encoded into the structure of an underlying graph [Lovett et al., 2010]. The distinctive features of quantum walks arise from the interplay between superposition and interference. A quantum walker may occupy multiple positions simultaneously, and amplitudes associated with different propagation paths can combine constructively or destructively. This coherent interference is what distinguishes quantum walks from classical random walks, where probabilities, rather than amplitudes are propagated [Childs et al., 2002].

Early work by Aharonov et al. [1993] formulated the discrete time quantum walk, in which a unitary coin operation determines the direction of motion. Shortly thereafter,

continuous time quantum walks were introduced by Farhi and Gutmann [1998], providing an alternative formulation without an internal coin degree of freedom. Continuous time quantum walks are defined directly in terms of Hamiltonian evolution on a graph, whereas discrete time quantum walks introduce an additional internal coin degree of freedom, which governs the direction of motion. Discrete time quantum walks therefore provide the most direct quantum analogue of lattice based classical walks, and they will form the primary focus of this work. In the early 2000s, quantum walks gained significant attention in quantum computing following the work of Ambainis et al. [2001] and others, who demonstrated algorithmic advantages including faster hitting times and applications to search problems [Kempe, 2002, Shenvi et al., 2003b,a]. Since then, quantum walks have developed into a broad framework spanning quantum algorithms, transport phenomena, and quantum simulation [Qiang et al., 2024, Molfetta, 2021]. It has also been shown that DTQWs admit continuum limits that reproduce well known partial differential equations, such as the Dirac equation, establishing a connection between discrete quantum dynamics and relativistic wave propagation [Meyer, 1996].

Quantum walks are motivated by the central role of classical random walks in the design of randomised algorithms, and have become a key component of many quantum algorithms. In certain oracle based problems, quantum walks can achieve exponential speedups over any classical approach [Childs et al., 2003]. Furthermore, the well known Grover search algorithm can be interpreted within the framework of quantum walks, highlighting their unifying role in quantum algorithm design [Shenvi et al., 2003b]. The following subsections introduce the mathematical structure of discrete time quantum walks, before discussing their key dynamical properties and continuous spacetime limits.

2.3.1 Discrete Time Quantum Walks

DTQWs are the quantum analogue of the classical random walk and again involve a walker on a line, graph or lattice. The walker now possesses a quantum coin, which may be flipped and the coin tells the walker to go in a superposition of both directions at once. Again, in higher dimensions, the walker may have a multidimensional coin corresponding to the lattice degrees of freedom, or a bias coin that causes the walker to

drift in a certain direction [Kendon, 2006]. The walker evolves in a Hilbert space that includes both position and an internal coin degree of freedom, allowing for coherent superpositions of different propagation paths.

I consider a DTQW on the line with Hilbert space $\mathcal{H}_{xc} = \mathcal{H}_x \otimes \mathcal{H}_c$. The position space $\mathcal{H}_x$ is spanned by $\{|x\rangle : x \in \Delta_x \mathbb{Z}\}$ with lattice spacing $\Delta_x > 0$. The coin Hilbert space is two dimensional, $\mathcal{H}_c = \text{span}\{|L\rangle, |R\rangle\}$, corresponding to "left"and "right" moves. The global state evolves by

$$|\psi(t+1)\rangle = \hat{\mathbf{U}}\,|\psi(t)\rangle = \hat{\mathbf{S}}\,(\mathbb{I} \otimes \hat{\mathbf{C}})\,|\psi(t)\rangle \tag{2.64}$$

where $\hat{\mathbf{C}} \in \mathcal{U}(2)$ is the coin operator, $\hat{\mathbf{S}}$ is the conditional shift, and $\hat{\mathbf{U}}$ is the unitary evolution operator. The conditional shift operator moves the particle left or right:

$$\hat{\mathbf{S}} = \sum_x \Big(|x+\Delta_x\rangle\langle x| \otimes |R\rangle\langle R| + |x-\Delta_x\rangle\langle x| \otimes |L\rangle\langle L| \Big) \tag{2.65}$$

The coin operator is a unitary operator that controls the dynamics of the DTQW. A common choice is the Hadamard coin, with two dimensional form

$$\hat{\mathbf{C}}_H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \tag{2.66}$$

The coin operator is as unbiased as possible across both outcomes. The Hadamard coin provides a simple and highly symmetric choice for the internal dynamics of the walk. Different coin parametrisations would, in general, lead to distinct effective behaviours, which will be explored further in the following sections [Tregenna et al., 2003]. Figure 2.3 shows the probability distribution function for an unbiased classical random walk and a discrete time Hadamard quantum walk on a one dimensional lattice. Both walks are initialised at the origin $x = 0$ and evolved for $T = 120$ timesteps on a lattice of size $N = 245$. The classical walk uses an unbiased binary coin, so that the walker moves left or right with probability $1/2$ at each timestep. The quantum walk uses the Hadamard coin with initial coin state $\chi = (|L\rangle + i\,|R\rangle)/\sqrt{2}$. The classical random walk gives a Gaussian distribution function centred about the origin, whereas the quantum

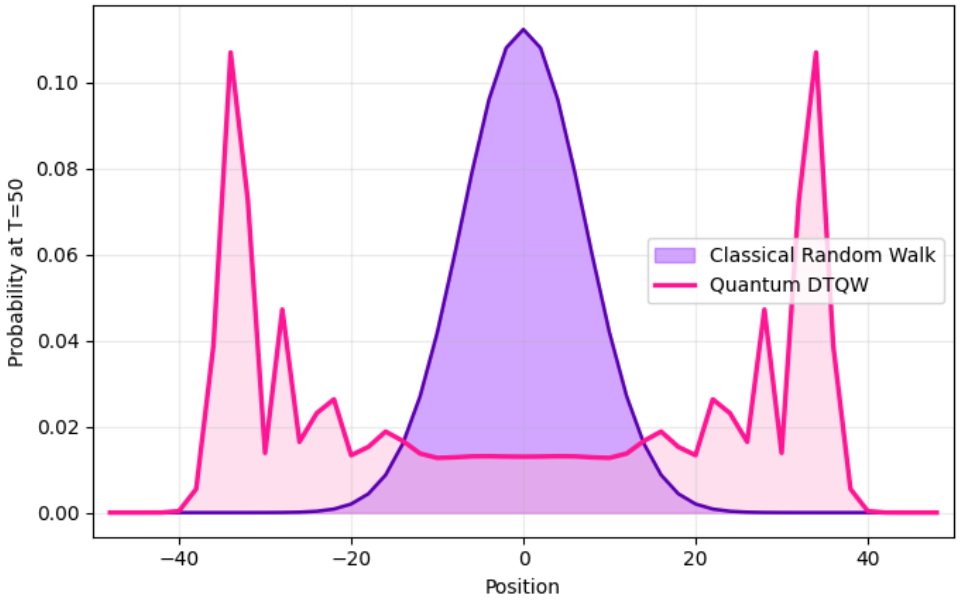

Figure 2.3: The probability distribution of a classical random walk versus a quantum walk. Both walks begin localised at position $x = 0$ and evolve for 50 time steps. The DTQW is evolved using a simple Hadamard coin given by Eqn. (2.66) with $\theta = \pi/4$.

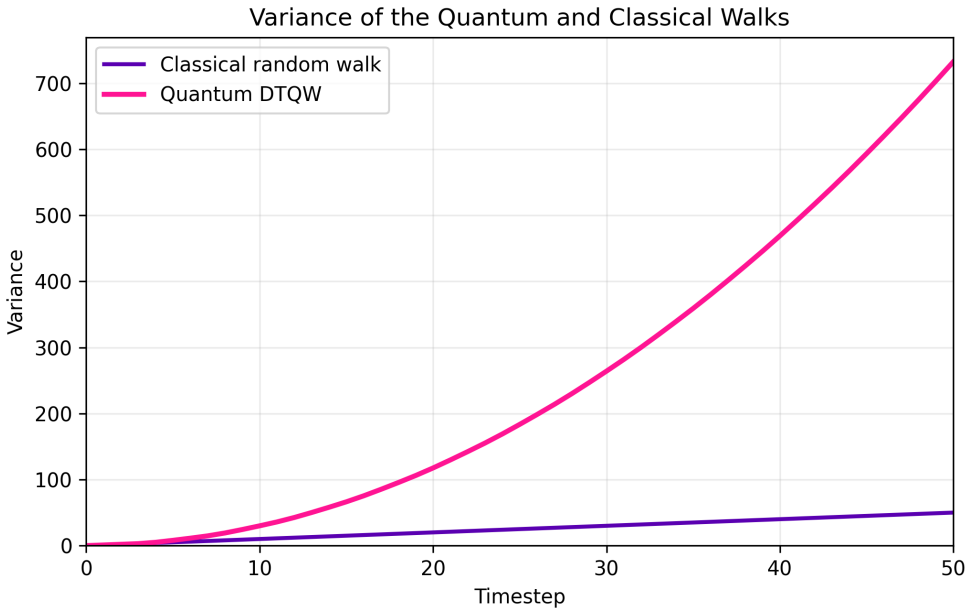

Figure 2.4: Variance growth for an unbiased classical random walk and a discrete time Hadamard quantum walk on a one dimensional lattice. The variance of the classical walk grows linearly, while the quantum walk variance grows quadratically.

walk spread significantly further and exhibits interference fringes. For visual clarity, the zero probability values are not shown in this plot. These zeros occur because the nearest neighbour walk can occupy only alternating lattice sites at each timestep, with the accessible parity determined by whether the timestep is even or odd. Figure 2.4 shows the variance growth for the same simulation. The classical random walk exhibits diffusive spreading, with variance growing linearly in time, $\mathrm{Var}(x) \propto t$. In contrast, the Hadamard quantum walk spreads ballistically due to coherent interference between propagation paths, leading to faster variance growth, approximately $\mathrm{Var}(x) \propto t^2$.

Coin operators and initial states

The dynamics of a DTQW are strongly influenced by the choice of coin operator, which governs the mixing between the internal propagation channels. In the two state case, the coin acts on the Hilbert space $\mathcal{H}_c = \mathrm{span}\{|L\rangle, |R\rangle\}$, and is represented by a unitary matrix $\hat{\mathbf{C}} \in U(2)$. A commonly used choice is the Hadamard coin, given by Eqn. (2.66), which generates an equal superposition of left and right moving components. This coin leads to symmetric ballistic spreading for suitably chosen initial states. More generally, one may consider a family of rotation coins of the form

$$\hat{\mathbf{C}}(\theta) = \begin{pmatrix} \cos\theta & \sin\theta \\ \sin\theta & -\cos\theta \end{pmatrix}, \tag{2.67}$$

where the parameter θ controls the degree of mixing between the two directions. For $\theta = \pi/4$, this reduces to the Hadamard coin, while other values introduce bias and modify the transport properties of the walk. In particular, asymmetric coins lead to drift and break the spatial symmetry of the probability distribution. The choice of coin operator determines the interference structure of the walk and therefore plays a central role in shaping the resulting probability distribution. Figure 2.5 shows the probability distributions of DTQWs on the line for multiple different coins. In this graph, the Hadamard coin is used, as well as the coin given by Eqn. (2.67) with $\theta = \pi/8$ and $\theta = \pi/3$. The choice of coin operator determines the interference structure of the walk and therefore plays a central role in shaping the resulting probability distribution.

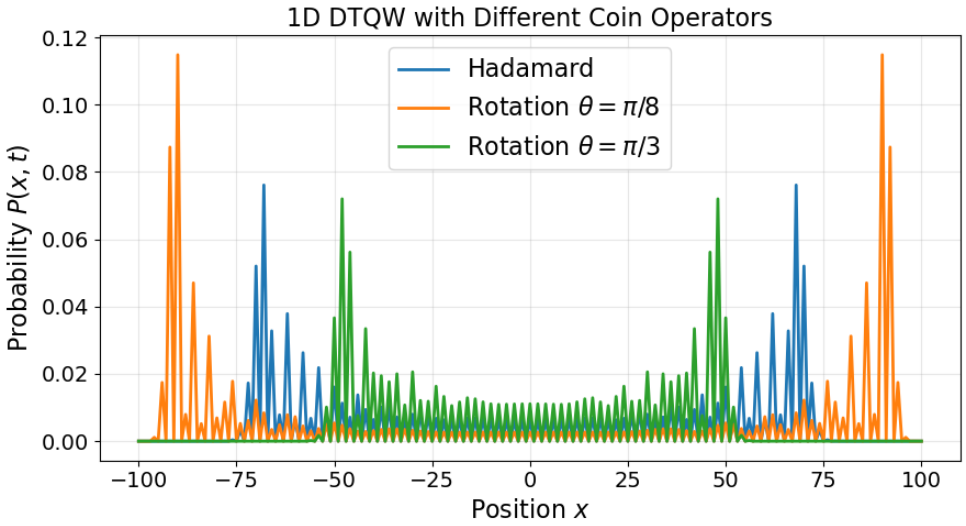

Figure 2.5: Probability distribution of a one dimensional discrete time quantum walk after $t = 100$ steps, for different choices of coin operator, with a fixed symmetric initial coin state.

While all cases exhibit ballistic spreading characterised by peaks propagating away from the origin, the relative amplitudes, symmetry, and internal interference patterns vary depending on the coin. In particular, biased or rotated coins break the symmetry of the walk, introducing directional preference and modifying the transport behaviour.

In addition to the coin operator, the initial coin state also plays a crucial role in determining the dynamics of the walk. The initial state of the walk determines the relative phase and amplitude between the left and right moving components, and therefore controls the symmetry of the resulting distribution [Tregenna et al., 2003]. Real superpositions produce symmetric spreading, while complex relative phases introduce asymmetry due to interference effects. This highlights the strong dependence of quantum walk dynamics on the initial state, in contrast to classical random walks. For an initial state of the form

$$|\psi(0)\rangle = |0\rangle \otimes (a|L\rangle + b|R\rangle), \tag{2.68}$$

where $a, b \in \mathbb{C}$ are complex probability amplitudes for the left and right moving coin states, respectively. The relative amplitudes and phases between $|L\rangle$ and $|R\rangle$ control the

interference between left and right moving components. Symmetric superpositions lead to symmetric probability distributions, while complex phase differences can introduce asymmetry and directional bias. This sensitivity to both the coin operator and the initial state highlights the flexibility of quantum walks as a framework for engineering transport behaviour.

To see explicitly how the initial coin state affects the dynamics, consider the first two steps of a Hadamard walk starting from

$$|\psi(0)\rangle = |0\rangle \otimes (a\,|L\rangle + b\,|R\rangle), \qquad |a|^2 + |b|^2 = 1. \tag{2.69}$$

After applying the Hadamard coin,

$$H\,|L\rangle = \frac{1}{\sqrt{2}}(|L\rangle + |R\rangle), \qquad H\,|R\rangle = \frac{1}{\sqrt{2}}(|L\rangle - |R\rangle), \tag{2.70}$$

the state becomes

$$|\psi(1)\rangle = \frac{a+b}{\sqrt{2}}\,|-1, L\rangle + \frac{a-b}{\sqrt{2}}\,|1, R\rangle\,. \tag{2.71}$$

Thus the probabilities at the first timestep are

$$P(-1,1) = \frac{|a+b|^2}{2}, \qquad P(1,1) = \frac{|a-b|^2}{2}. \tag{2.72}$$

Already at this stage, the relative phase between a and b determines whether probability is enhanced on the left or right. After a second timestep, the state is

$$|\psi(2)\rangle = \frac{a+b}{2}\,|-2, L\rangle + \frac{a+b}{2}\,|0, R\rangle + \frac{a-b}{2}\,|0, L\rangle - \frac{a-b}{2}\,|2, R\rangle\,. \tag{2.73}$$

The probability at the origin is therefore

$$P(0,2) = \frac{|a-b|^2}{4} + \frac{|a+b|^2}{4} = \frac{1}{2}, \tag{2.74}$$

while the outer probabilities are

$$P(-2,2) = \frac{|a+b|^2}{4}, \qquad P(2,2) = \frac{|a-b|^2}{4}. \tag{2.75}$$

This simple calculation shows that the initial amplitudes and phases directly control the interference between left and right moving components. For example, if $a = b = 1/\sqrt{2}$, then $P(-1, 1) = 1$ and the first step is entirely left moving. By contrast, if $a = 1/\sqrt{2}$ and $b = i/\sqrt{2}$, then

$$|a + b|^2 = |a - b|^2 = 1, \tag{2.76}$$

so the first step is balanced. The phase of the initial coin state therefore affects the symmetry and directionality of the walk through quantum interference.

Figure 2.6 shows the probability distributions for multiple different quantum walk coins and initial states. This figure illustrates the impact of the coin choice and the initial state in determining the dynamics of the walk. The coin operator governs the underlying mixing and transport properties, while the initial state controls the interference between propagation channels. Together, these determine the symmetry, localisation, and overall structure of the distribution. In all cases, the walk exhibits ballistic spreading, with variance growing quadratically in time, in contrast to the diffusive behaviour of classical random walks. In the case of the classical random walk, the coin usually determines a probability of moving in each direction, and a biased coin produces a corresponding drift. In the quantum case, however, different choices of coin can change not only the average direction of motion, but also the interference between different paths. Similarly, changing the initial coin state can make the same walk appear symmetric, biased, or localised depending on how the initial amplitudes and phases are chosen. This sensitivity to the coin and initial state will be used throughout the later chapters of this thesis. Rather than treating the coin as a fixed ingredient, it will be used as a tunable part of the model. By modifying the coin, the initial state, and later the decoherence structure, the large scale properties of quantum walks can be adjusted to produce desired transport behaviour.

2.3.2 Continuum Limits and the Dirac Equation

A key motivation for studying DTQWs is that their continuum limits can reproduce partial differential equations [Molfetta, 2021]. A DTQW is defined through a discrete update rule on a lattice, but in the limit of the small lattice spacing and timestep

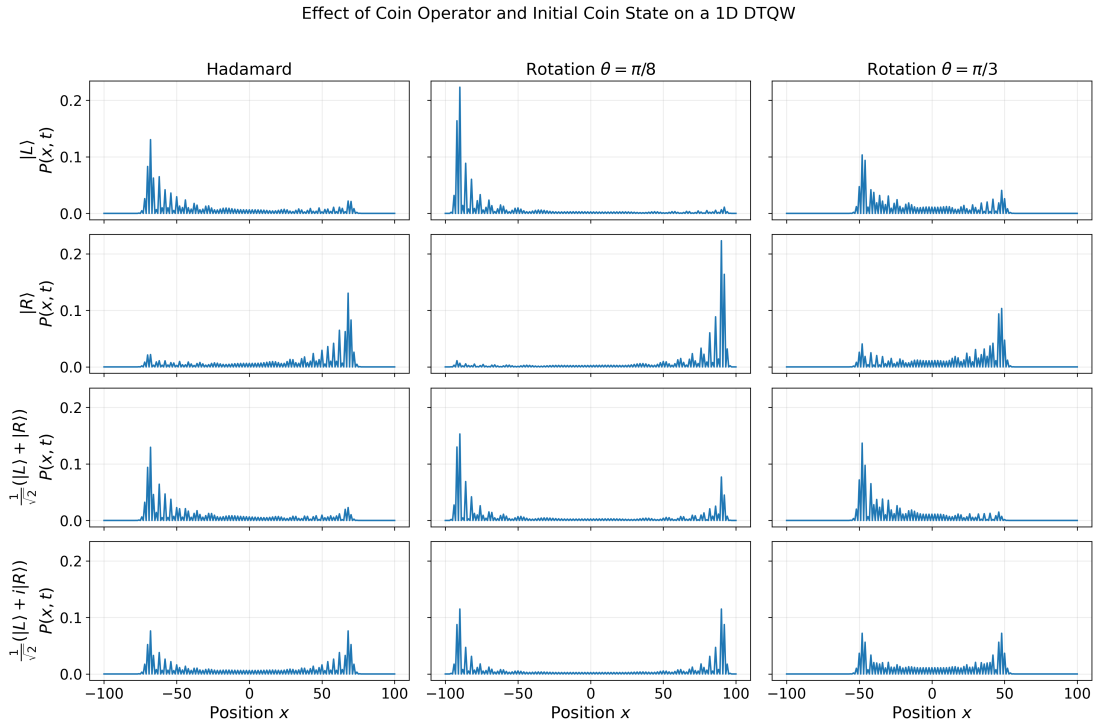

Figure 2.6: Probability distributions of a one dimensional discrete time quantum walk after $t = 100$ steps for combinations of different coin operators (columns) and initial coin states (rows).

approaching zero, the large scale evolution may approach a continuous PDE. The best known example is the Dirac equation. A suitable quantum walk is able to recover Dirac dynamics in the limit of small lattice spacings and discrete time steps [Yamagishi et al., 2023]. In this setting, the coin degree of freedom is interpreted as an internal spin, while the shift operator represents propagation on the lattice. Rotations in the coin space then generate the mass term in the limiting Dirac equation [Molfetta and Debbasch, 2012]. This provides a powerful means to simulate continuous relativistic equations using discrete quantum dynamics.

Continuous spacetime limits have been derived for many different kinds of quantum walks and lattices, including walks on triangular and honeycomb lattices [Arrighi et al., 2018]. DTQWs therefore provide a natural lattice model for coherent quantum transport. Early work by Meyer [1996] demonstrated that quantum lattice gas automata can reproduce the Dirac equation in the continuum limit, establishing one of the first connections between discrete quantum dynamics and relativistic wave equations. Subsequent work showed that DTQWs themselves admit continuum limits corresponding to relativistic quantum equations. In particular, Strauch [2006] demonstrated that the one dimensional DTQW reproduces the Dirac equation in the continuous limit, while later work by Di Molfetta [2019] developed a systematic framework for deriving continuous spacetime limits of quantum walks.

Within this framework, the structure of the coin operator, shift operators, and scaling assumptions determine which PDE emerges in the continuum limit. The connection has been extended beyond the simplest Dirac limits. Di Molfetta et al. [2014] showed that one dimensional DTQWs with space and time dependent coin parameters can recover, in the continuous limit, the dynamics of a Dirac fermion coupled to artificial electric and gravitational fields. Work by Jolly and Di Molfetta [2023] shows that modifications to the standard walk on the line can admit limits that may be useful for simulating fluid properties such as dispersion, which will be discussed further in later sections. In these models, two different coin operators can be used within a single timestep, changing the structure of the effective continuum equation. Related work by Hatifi et al. [2019] has also explored Madelung transformations of quantum walk or

Dirac dynamics, further connecting quantum walk continuum limits to hydrodynamic descriptions. By viewing the quantum walk through a nonlinear hydrodynamic lens, the authors were able to simulate nonlinear fluid behaviour, such as shock, using a quantum walk model. These developments are important for this thesis because they show that continuum limits are not fixed, but can be tuned by modifying the walk.

This perspective is especially useful because exact analytical solutions exist only for a limited class of PDEs with highly idealised boundary conditions. Most physically relevant equations do not admit closed form solutions in general, so one usually relies on numerical simulation [Tennie, 2025]. This requires discretising space and time and evolving the system over many degrees of freedom, leading to large computational costs. The resolution must be sufficiently fine to capture the relevant physical scales, while stability constraints often impose small timesteps. If a DTQW is constructed such that its continuum limit reproduces a PDE of interest, then the walk provides a natural discrete model for simulating that equation. Rather than starting from the PDE and discretising it directly, the desired dynamics are generated intrinsically by the underlying quantum evolution, with the target continuum behaviour emerging in the limit of small lattice spacing and timestep. This suggests that DTQWs may be engineered not only to reproduce relativistic equations such as the Dirac equation, but also to simulate other continuum transport equations by modifying the coin, shift, lattice, scaling, and, in this work, the decoherence structure.

This perspective is similar in spirit to the LBM, as one begins with a discrete kinetic model on a lattice, evolves microscopic populations through local update rules, and then studies the continuum limit to determine which macroscopic PDE is recovered. The Chapman-Enskog expansion provides the formal multiscale procedure for making this connection, as it separates the microscopic lattice dynamics into different orders in a small parameter and shows how equilibrium, advection, diffusion, and viscosity emerge at successive orders. A similar philosophy may be applied to DTQWs. Although the microscopic variables are now quantum amplitudes or density matrices rather than classical populations, the aim is again to understand how a local lattice update produces a continuum equation in an appropriate limit. In this sense, the continuum limit

calculations used for DTQWs can be viewed as a quantum analogue of the kinetic reasoning behind LBM, and a Chapman-Enskog style expansion that utilises these continuum limits may provide a framework for extracting the macroscopic transport equations generated by the quantum walk.

Since continuous spacetime limits of quantum walks are fundamental to this work, the following subsections derive the limit of the one dimensional walk on the line and demonstrate that this model produces the one dimensional Dirac equation as its continuum limit. I will then introduce plasticity in DTQW models and explains how the continuum limit is taken in that setting, providing the foundation for the quantum walk models developed in the later chapters.

1D Continuous Spacetime Limit

We now derive the continuum limit of the standard two state DTQW on the line. This derivation provides the simplest example of how a discrete quantum walk may reproduce a relativistic continuum equation and follows the same scaling logic used later for the different quantum walk models [Molfetta and Arrighi, 2019]. I first consider a one dimensional DTQW with Hilbert space $\mathcal{H} = \mathcal{H}_x \otimes \mathcal{H}_c$, where the position space is spanned by lattice states $\mathcal{H}_x = \text{span}\{|x\rangle : x \in \Delta_x \mathbb{Z}\}$, and the two dimensional coin space is $\mathcal{H}_c = \text{span}\{|L\rangle, |R\rangle\}$. The states $|L\rangle$ and $|R\rangle$ label the two internal propagation channels of the walker. For a well defined continuous spacetime limit to exist, the one step unitary operator of the walk must approach the identity in the limit $\varepsilon \to 0$ [di Molfetta and Debbasch, 2012]. More precisely, the unitary must be expandable around the identity, so that it can be written in the form

$$\hat{\mathbf{U}}(\varepsilon) = \mathbb{I}_2 + \varepsilon U_1 + \varepsilon^2 U_2 + \cdots , \tag{2.77}$$

where the matrices $U_1, U_2, \ldots$ encode the successive orders of the continuum expansion. This requirement imposes restrictions on the class of unitary walks that can admit a continuous spacetime limit. If the unitary does not approach the identity as $\varepsilon \to 0$, the difference between consecutive timesteps does not become infinitesimal, and the

limiting evolution cannot be interpreted as a standard continuous time PDE in the usual way.

A single step of the walk is defined by Eqn. (2.64) and the conditional shift operator is given by Eqn. (2.65). For the coin operator, I choose

$$\hat{\mathbf{C}}(\theta) = \begin{pmatrix} \cos\theta & -i\sin\theta \\ -i\sin\theta & \cos\theta \end{pmatrix}. \tag{2.78}$$

The angle θ controls the mixing between the left and right moving channels. This coin can also be written compactly in terms of the Pauli X matrix,

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{2.79}$$

Using the identity

$$e^{-i\theta\sigma_x} = \cos\theta\,\mathbb{I}_2 - i\sin\theta\,\sigma_x, \tag{2.80}$$

the coin operator becomes

$$\hat{\mathbf{C}}(\theta) = \cos\theta\,\mathbb{I}_2 - i\sin\theta\,\sigma_x = e^{-i\theta\sigma_x}. \tag{2.81}$$

Thus, this coin corresponds to a rotation generated by σ_x in the two dimensional coin space. To take the continuum limit, it is convenient to move to Fourier space. We define

$$|\psi(t)\rangle = \int_{-\pi/\Delta_x}^{\pi/\Delta_x} \psi(k,t)|k\rangle\, dk, \tag{2.82}$$

where $\psi(k,t)$ is a two component spinor. In Fourier space, the conditional shift becomes diagonal:

$$\tilde{\mathbf{S}}(k) = \begin{pmatrix} e^{-ik\Delta_x} & 0 \\ 0 & e^{ik\Delta_x} \end{pmatrix} = e^{-ik\Delta_x\sigma_z}. \tag{2.83}$$

Therefore the Fourier space walk operator is

$$\tilde{\mathbf{U}}(k) = \tilde{\mathbf{S}}(k)\hat{\mathbf{C}}(\theta) = e^{-ik\Delta_x\sigma_z}e^{-i\theta\sigma_x}. \tag{2.84}$$

We now introduce the continuum scaling

$$\Delta_x = \varepsilon, \qquad \Delta_t = \varepsilon, \qquad \theta = \varepsilon m, \tag{2.85}$$

where m remains finite as $\varepsilon \to 0$. Under this scaling, the shift and coin operators become

$$\tilde{\mathbf{S}}(k,\varepsilon) = e^{-ik\varepsilon\sigma_z}, \tag{2.86}$$

$$\hat{\mathbf{C}}(m,\varepsilon) = e^{-i\varepsilon m\sigma_x}. \tag{2.87}$$

These are then expand as

$$e^{-ik\varepsilon\sigma_z} = \mathbb{I}_2 - i\varepsilon k\sigma_z + O(\varepsilon^2), \tag{2.88}$$

$$e^{-i\varepsilon m\sigma_x} = \mathbb{I}_2 - i\varepsilon m\sigma_x + O(\varepsilon^2). \tag{2.89}$$

Hence the full walk operator becomes

$$\tilde{\mathbf{U}}(k) = \left(\mathbb{I}_2 - i\varepsilon k\sigma_z + O(\varepsilon^2)\right)\left(\mathbb{I}_2 - i\varepsilon m\sigma_x + O(\varepsilon^2)\right) \tag{2.90}$$

$$= \mathbb{I}_2 - i\varepsilon\left(k\sigma_z + m\sigma_x\right) + O(\varepsilon^2). \tag{2.91}$$

Thus, to leading order in ε, the discrete walk has the form

$$\tilde{\mathbf{U}}(k) = \mathbb{I}_2 - i\varepsilon\hat{\mathbf{H}}(k) + O(\varepsilon^2), \tag{2.92}$$

with effective Hamiltonian

$$\hat{\mathbf{H}}(k) = k\sigma_z + m\sigma_x. \tag{2.93}$$

After expanding the unitary operator, the evolution equation now may be written in Fourier form, as

$$\psi(k, t+\varepsilon) = \tilde{\mathbf{U}}(k)\psi(k,t). \tag{2.94}$$

Expanding the left hand side gives

$$\psi(k, t+\varepsilon) = \psi(k,t) + \varepsilon \partial_t \psi(k,t) + O(\varepsilon^2). \tag{2.95}$$

Substituting the expansion of $\tilde{\mathbf{U}}(k)$ in Eqn. (2.92), we obtain

$$\psi(k,t) + \varepsilon \partial_t \psi(k,t) = \left[\mathbb{I}_2 - i\varepsilon \left(k\sigma_z + m\sigma_x\right)\right] \psi(k,t) + O(\varepsilon^2). \tag{2.96}$$

Cancelling the zeroth order terms and dividing by ε gives to first order

$$\partial_t \psi(k,t) = -i \left(k\sigma_z + m\sigma_x\right) \psi(k,t). \tag{2.97}$$

Equivalently,

$$i\partial_t \psi(k,t) = \left(k\sigma_z + m\sigma_x\right) \psi(k,t). \tag{2.98}$$

Finally, transforming back to real space using $k \mapsto -i\partial_x$, we find

$$i\partial_t \psi(x,t) = \left(-i\sigma_z \partial_x + m\sigma_x\right) \psi(x,t). \tag{2.99}$$

This is the one dimensional Dirac equation in Hamiltonian form,

$$i\partial_t \psi(x,t) = \hat{\mathbf{H}}_D \psi(x,t), \qquad \hat{\mathbf{H}}_D = -i\sigma_z \partial_x + m\sigma_x. \tag{2.100}$$

Thus, the basic one dimensional discrete time quantum walk admits the Dirac equation as its continuum limit. The conditional shift generates the spatial derivative term, while the coin rotation generates the mass term.

Plasticity in quantum walks

The previous subsection considered the standard continuous spacetime limit of a DTQW, where the lattice spacing and timestep are sent to zero at the same rate, $\Delta_x = \varepsilon, \Delta_t = \varepsilon$. This is not the only possible continuum scaling. A further important development is the introduction of plasticity in DTQW models, where the spatial and temporal discretisation parameters are allowed to tend to zero at different rates [Di Molfetta, 2024].

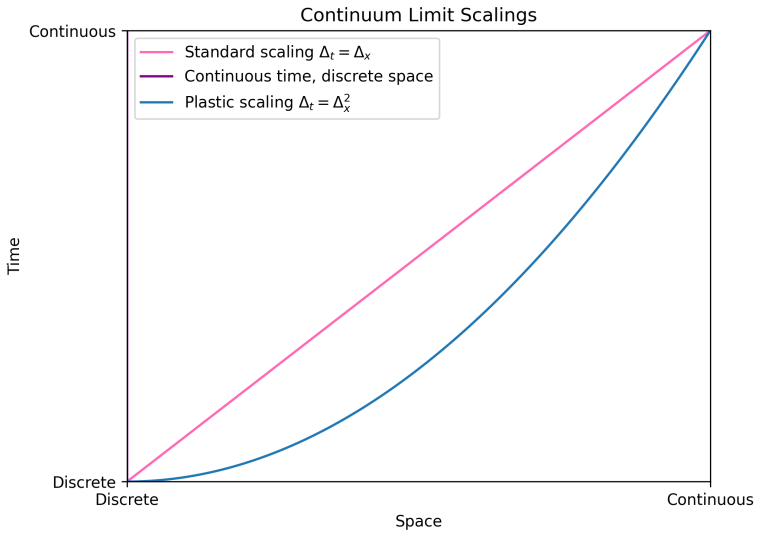

Figure 2.7: Schematic comparison of different continuum limit scalings.

Following the notation used throughout this thesis, I write

$$\Delta_t = \varepsilon, \qquad \Delta_x = \varepsilon^{1-\kappa}, \qquad 0 \leq \kappa \leq 1. \tag{2.101}$$

The parameter κ controls the anisotropy of the spacetime grid. When $\kappa = 0$, space and time scale together and one recovers the usual continuous spacetime limit. When $\kappa = 1$, the timestep tends to zero while the spatial lattice spacing remains finite, giving a continuous time, discrete space limit. Intermediate values $0 < \kappa < 1$ interpolate between these two regimes [Molfetta, 2021]. Figure 2.7 is a graphical demonstration of this scaling. The standard continuous spacetime limit, shown in pink, sends the spatial and temporal lattice spacings to zero at the same rate, $\Delta_t \propto \Delta_x$. The vertical purple line represents the continuous time, discrete space limit, where $\Delta_t \to 0$ while Δ_x remains finite. The blue curve illustrates a plastic scaling, here shown schematically as $\Delta_t \propto \Delta_x^2$, where time tends to the continuum faster than space. Plastic scalings therefore interpolate between the usual spacetime limit and the continuous time, discrete space regime by allowing space and time to shrink anisotropically.

Plasticity is therefore a way of making the continuum limit less rigid. Instead of forcing the walk to converge only under one fixed relationship between Δ_x and Δ_t, a

plastic walk can adapt to different scalings of the underlying spacetime grid. In the terminology of Di Molfetta [2019], a plastic quantum walk is a quantum walk which admits both a continuous spacetime limit and a continuous time/discrete space limit, depending on the scaling parameters. Strong plasticity refers to the case where the same walk, with the same leading coin structure, admits both limits. Weak plasticity requires the coin scaling itself to be adjusted depending on the scaling.

For a one dimensional DTQW, Molfetta [2021] shows that there is no nontrivial plasticity when the continuum limit is taken after a single timestep, $t = 1$. In that case, the anisotropic scaling does not produce both a continuous spacetime limit and a continuous time/discrete space limit and only the usual spacetime limit survives, leading to a continuum equation. Nontrivial plasticity appears when the walk is viewed stroboscopically, for example after two timesteps. One then considers

$$|\psi(t + 2\Delta_t)\rangle = \hat{\mathbf{U}}^2|\psi(t)\rangle. \tag{2.102}$$

This allows the walk to support both the continuous spacetime limit and the continuous time/discrete space limit, which is the defining feature of a plastic quantum walk.

As discussed above, a continuous limit requires the effective walk operator to approach the identity as $\varepsilon \to 0$. Some choices of coin and shift are unitary, but their one step operator does not expand directly around $\mathbb{I}$. In such cases, the evolution over a single timestep is too large to be interpreted as an infinitesimal time update. A stroboscopic construction also provides a way around this by considering the walk only after two or more timesteps. The product of several unitary steps may cancel the leading order terms that prevent the one step operator from approaching the identity, so that the effective multistep operator admits a controlled expansion.

Overall, plasticity provides an additional tuning mechanism for DTQWs. The choice of coin, shift, timestep, lattice spacing, and scaling exponent κ determines which terms remain in the limiting equation. As in the standard continuum limit, the resulting PDE is not put in by hand, but emerges from the asymptotic expansion of the microscopic walk. However, the plastic scaling changes the ordering of spatial and temporal terms, making it possible to access limiting regimes that are not visible in the ordinary $\Delta_x \propto$

Δ_t limit.

A related example that is central to this work is the twisted quantum walk introduced by Jolly and Di Molfetta [2023]. This model uses a plastic scaling together with a scaled coin parameter. This corresponds to the plastic case with $\kappa = 1/2$ in the notation used here. By introducing a twist rotation in the coin basis, the authors show that the continuum limit is no longer just the usual Dirac equation, but a generalised Dirac equation with an additional dispersive term. In particular, the coin parameter controls the linear Dirac term, while the twist angle controls the quadratic dispersive contribution. Thus, the twist modifies the plastic continuum limit by allowing a second order spatial contribution to survive at leading order. When the twist angle is set to zero, the quadratic term disappears and the usual Dirac limit is recovered, while for non-zero values the walk interpolates between Dirac and dispersive behaviour. The continuous spacetime limit of this model is derived explicitly in Section 4.6.1 before being modified and used in later sections.

Plasticity will be important in future chapters because diffusion and transport terms often appear at different asymptotic orders. By choosing an appropriate scaling between space, time, coin parameters, and decoherence strength, one can control whether advection, diffusion, or coherent quantum terms dominate the continuum equation. This provides a useful framework for engineering quantum walk models whose large scale behaviour resembles classical transport equations such as diffusion and advection-diffusion.

2.3.3 Decoherence in Quantum Walks

In realistic settings, quantum systems are not perfectly isolated and interact with their environment, leading to decoherence. This process suppresses quantum coherence and drives the system toward classical behaviour as the system loses information into the environment [Kendon, 2007]. While decoherence is often regarded as detrimental, as it suppresses quantum coherence and drives the system toward classical behaviour, it can also play a constructive role in quantum walks. Decoherence provides a mechanism for interpolating between ballistic and diffusive regimes, enabling the engineering of

effective transport properties [Brun et al., 2003a, Love and Boghosian, 2005]. In this sense, decoherence should not be viewed solely as an unwanted perturbation, but rather as a tunable resource that can be used to control the dynamics of quantum walks.

Open quantum walks (OQWs) generalise the framework of quantum walks to include environmental noise or engineered dissipation through completely positive trace preserving (CPTP) maps. Environment induced errors in discrete quantum walks were first studied by Dür et al. [2002] to model errors in proposed experimental implementations. Kendon and Tregenna [2003] found that decoherence could improve the algorithmic properties of quantum walks by providing a smoother uniform sampling that was still quadratically faster than classical. Dissipation has been leveraged more generally for quantum computing and quantum state preparation [Beige et al., 2000, Verstraete et al., 2009]. More general open system quantum walk models were introduced by Attal et al. [2012a,b], Sinayskiy and Petruccione [2012]. They provide a natural setting for describing transport in open quantum systems and have found applications in dissipative computing, state transfer, and biological energy transport [Mohseni et al., 2008, Sinayskiy and Petruccione, 2014].

It is important to note that there are multiple ways to incorporate decoherence into quantum walks. Projective decoherence involves decoherence acting on the coin or spatial subspace, however you may also use higher dimensional or multiple coins to dilute the effects of interference [Brun et al., 2003a]. This work focusses on projective decoherence, and to mathematically describe such decoherence we need a new notation. Decoherence is usually described as a non unitary evolution of the quantum walk, and it is necessary to introduce the density matrix formalism to describe these mixed states [Blum, 2012]. All operators act on same the Hilbert space $\mathcal{H}_{xc} = \mathcal{H}_x \otimes \mathcal{H}_c$ as described in the previous section. The system is described by a time dependent density operator, expressed in the computational basis as:

$$\hat{\boldsymbol{\rho}}(t) = \sum_{x,c} \sum_{x',c'} \rho_{x,c}^{x',c'}(t) |x, c\rangle\langle x', c'| \tag{2.103}$$

This operator is Hermitian, positive semi definite, of unit trace, and bounded. It

governs the evolution of both the internal coin and spatial degrees of freedom. The discrete time evolution of the system includes both unitary and non-unitary dynamics. A single timestep, Δ_t, is governed by the map:

$$\hat{\boldsymbol{\rho}}(t+\Delta_t) = (1-\gamma)\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}(t)\hat{\mathbf{U}}^\dagger + \gamma\sum_j \mathbb{P}_j\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}(t)\hat{\mathbf{U}}^\dagger\mathbb{P}_j^\dagger \tag{2.104}$$

where $\gamma \in [0,1]$ is the probability of decoherence per full step, and the operator $\hat{\mathbf{U}}$ is the full unitary evolution operator. This decoherence model was explored in [Kendon, 2007] and evolved numerically for various choices of $\mathbb{P}_j$.

The Lindbladian form of the discrete evolution equation is necessary to take the continuum limit of a DTQW with decoherence. Rewriting the stepwise update in Lindblad form ensures that the resulting generator remains both trace preserving and completely positive, regardless of the specific choice of spatial step scaling [Annabestani et al., 2010]. The Lindblad form given below is the infinite temperature Lindblad, whose stationary state is the maximally mixed state and treats all energy levels symmetrically. The exact form of the Lindblad CPTP map is given by:

$$\hat{\boldsymbol{\rho}}(t+\Delta_t) = \hat{\boldsymbol{\rho}}(t) - i[\hat{\mathbf{H}},\rho(t)] + \gamma\sum_j \big(L_j\hat{\boldsymbol{\rho}}(t)L_j^\dagger - \frac{1}{2}\{L_j^\dagger L_j, \hat{\boldsymbol{\rho}}(t)\}\big). \tag{2.105}$$

Here, the jump operators model decoherence as a projective noise. The projectors, $\{\mathbb{P}_j\}$, specify the particular decoherence channel. The projectors satisfy the usual relations

$$\mathbb{P}_j^\dagger = \mathbb{P}_j, \quad \mathbb{P}_j^2 = \mathbb{P}_j, \quad \sum_j \mathbb{P}_j = \mathbb{I}. \tag{2.106}$$

In the absence of decoherence, quantum walks produce ballistic spreading with two counter propagating peaks in position space. However, decoherence modifies this structure, introducing damping and potential crossover to classical behaviour depending on the model, strength and type of environmental coupling [Brun et al., 2003a]. This relationship between coherence and noise will play a central role in the constructions developed later in this thesis.

Figure 2.8 illustrates the effect of projective decoherence on a simple one dimensional

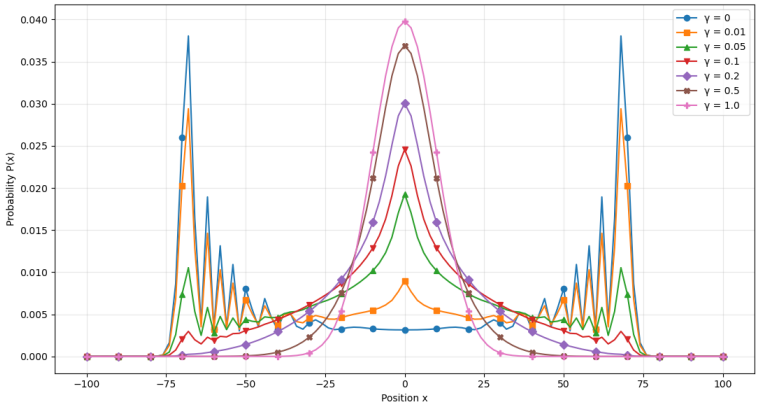

(a) Coin decoherence

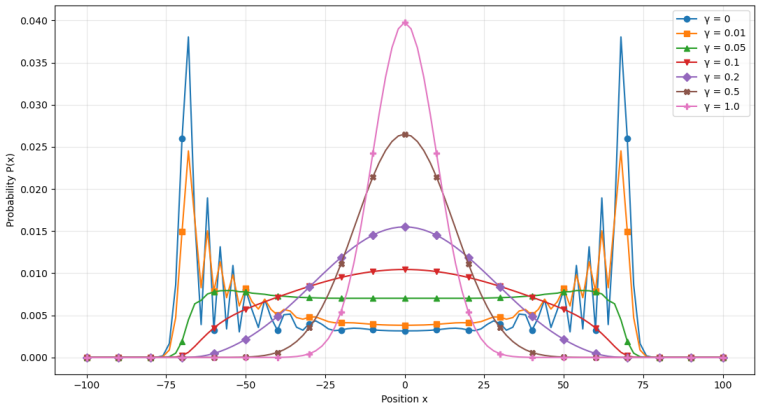

(b) Spatial decoherence

Figure 2.8: Effect of projective decoherence on a one dimensional Hadamard quantum walk after $t = 100$ timestep for both coin and spatial decoherence. Decoherence strength γ is varied from 0 to 1.

Hadamard walk. The walk is evolved for $t = 100$ timesteps from an initially localised spatial state with coin state $(|0\rangle + i\,|1\rangle)/\sqrt{2}$, while the decoherence strength γ is varied. The projectors used in this example are given by $\mathbb{P}_L = |L\rangle\langle L|$ and $\mathbb{P}_R = |R\rangle\langle R|$. When $\gamma = 0$, the walk is fully coherent and displays the usual ballistic quantum walk profile, with two dominant propagating peaks and oscillatory interference structure. As γ is increased, these interference fringes are progressively suppressed and the distribution becomes smoother and more classical. The two decoherence choices have similar qualitative effects, but they act on different parts of the walk. Coin decoherence removes coherences between the internal left and right moving coin states, while leaving spa-

tial coherences intact. As a result, for small and moderate γ, the coin dephased walk still retains visible remnants of the ballistic peaks and interference structure. Spatial decoherence, by contrast, directly suppresses coherences between different lattice positions. This damps the spatial interference pattern more directly, producing a smoother and more classical looking distribution at intermediate decoherence strengths. In both cases, increasing γ drives the walk away from the coherent ballistic regime and towards a centred Gaussian profile.

The trace of the density matrix plays a fundamental role in the description of quantum states. For any physical state, the density operator must satisfy the normalisation condition

$$\mathrm{Tr}(\hat{\boldsymbol{\rho}}) = 1, \tag{2.107}$$

which ensures that the total probability of all possible measurement outcomes is equal to one [Patel, 2023]. In this sense, the trace of the density matrix represents the total probability contained in the system. More generally, the trace operation provides a natural way to compute expectation values of observables. For an observable $\hat{A}$, the expectation value is given by

$$\langle A \rangle = \mathrm{Tr}(\boldsymbol{\rho}\hat{A}), \tag{2.108}$$

which generalises the standard expression from pure states to mixed states. The preservation of the trace under time evolution is therefore a necessary condition for any physically valid quantum dynamics, ensuring that probability is conserved. In open quantum systems, CPTP maps are required to ensure that the trace of the density matrix remains unity throughout the evolution, even in the presence of decoherence and dissipation.

The previous sections have now introduced the main ingredients needed for this thesis, quantum walks, continuum limits as a way of extracting PDEs, plastic scalings as a way of tuning which terms survive, and decoherence. These ideas lead to the question of whether quantum walk models can be used for computational fluid dynamics and transport simulation. The following section therefore reviews existing work on quantum computational fluid dynamics and quantum lattice Boltzmann methods. This places

the present approach in context, highlighting both the motivation for using quantum lattice models and the difficulties that arise when trying to reproduce classical kinetic methods such as LBM within a coherent quantum framework.

2.4 Quantum Computational Fluid Dynamics and Quantum Lattice Boltzmann Methods

Quantum computational fluid dynamics and quantum lattice Boltzmann methods (QLBM) aim to investigate whether quantum computation can provide advantages for the simulation of transport phenomena, kinetic equations, and fluid dynamics [Au-Yeung et al., 2024a]. The development of these approaches has been strongly motivated by the increasing computational demands of large scale classical fluid simulations and the growing interest in exascale and post exascale computing architectures [Barral et al., 2025]. Classical LBM is already highly efficient and naturally parallelisable, making it exceptionally well suited to high performance computing systems. Consequently, no clear general quantum advantage over classical LBM is currently known for standard fluid simulations. Instead, QLBM has largely developed as a conceptual and exploratory framework investigating how transport dynamics, kinetic theory, and hydrodynamic behaviour may emerge from quantum systems.

Potential advantages of QLBM may arise in specialised settings rather than from a general replacement of classical CFD methods. For example, amplitude encoding can represent a classical field over (N) lattice sites using only $\mathcal{O}(\log N)$ qubits, suggesting a possible memory advantage over storing the full distribution explicitly [Nmaju et al., 2025]. Quantum algorithms may also be useful when the goal is to sample from the state, extract selected global observables, or evolve particular classes of dynamics without reconstructing the full field. However, this compression does not automatically imply an overall computational speedup. Preparing the encoded quantum state may be expensive, and recovering a full classical flow field generally requires many measurements, which can remove the apparent advantage. Moreover, quantum dynamics is fundamentally linear and unitary, whereas classical fluid dynamics and LBM rely

on nonlinear, state dependent collision operators. Thus, the usefulness of a QLBM encoding depends strongly on the structure of the problem, the desired output, and the constraints of the target quantum platform [Rathore et al., 2026].

The origins of QLBM can be traced back to early work on quantum lattice gas automata and quantum walks. Initial proposals were inspired by quantum lattice gas models developed by researchers such as Meyer [1996] and Lloyd [1996] who demonstrated that discrete quantum lattice systems could reproduce wave propagation and hydrodynamic behaviour. These early models established a close relationship between quantum walks, lattice gas automata, and kinetic transport equations. Since then, the field has expanded significantly, with a large body of work attempting to formulate fluid and transport dynamics within quantum computational frameworks.

A major contribution to the conceptual development of the field was made by Sauro Succi and collaborators. Succi [1996] argued that lattice Boltzmann methods possess structural similarities to quantum transport models, particularly through their discrete streaming and local collision structure. In later work by Succi et al. [2015], it was explicitly shown that the lattice Boltzmann streaming step can be interpreted as a quantum walk. This work demonstrated that the Dirac equation, quantum walks, and lattice Boltzmann methods share a common mathematical structure based on discrete velocity transport and local updates. However, an important distinction remains. Classical LBM relies on state dependent equilibrium distributions and therefore introduces nonlinearity through the collision operator, whereas quantum evolution is intrinsically linear. This makes it difficult to directly reproduce nonlinear fluid behaviour within a purely coherent quantum framework.

One early direction in this area focused on quantum lattice gas algorithms and the direct derivation of hydrodynamic equations from microscopic quantum dynamics. Yepez [2002] introduced a quantum lattice gas model in which particle populations are encoded into qubit amplitudes and evolved through sequences of local collision and streaming operations. Using a Chapman-Enskog expansion of the resulting kinetic equation, Yepez showed that the macroscopic dynamics reproduce Burgers equations containing both nonlinear advection and viscous diffusion. This work represented some

of the earliest attempts to derive nonlinear fluid equations directly from quantum computational dynamics. Although Yepez's model is an important early example of a quantum algorithm reproducing a nonlinear hydrodynamic equation, it also illustrates one of the central difficulties in quantum lattice gas approaches to fluid simulation. The algorithm encodes the occupation probabilities into qubit amplitudes and uses a local unitary collision gate to generate the effective Burgers dynamics. However, the macroscopic flow field must still be known across the lattice at each timestep, so the method requires measurement or ensemble averaging to recover the occupation probabilities. This means that, although the local memory requirement is reduced compared with the classical lattice gas model, the algorithm is not a fully coherent quantum simulation of a macroscopic fluid field. The repeated extraction of classical information weakens the possible quantum advantage and makes the scheme closer to a hybrid kinetic algorithm.

Steijl and Barakos [2018] proposed an early hybrid quantum-classical approach for computational fluid dynamics, using the quantum Fourier transform as part of a Poisson solver within a vortex in cell method for incompressible flow. The work demonstrated how quantum subroutines could be embedded into an otherwise classical CFD algorithm and studied the effect of quantum noise on the resulting flow solution. However, the method still requires repeated exchange of information between the classical and quantum parts of the computation, so state preparation, measurement noise, and readout overhead remain major limitations. Todorova and Steijl [2020] later developed a quantum algorithm for the collisionless Boltzmann equation, representing the transport step as a quantum walk streaming process implemented through controlled NOT circuits. This avoids repeated measurement at every timestep and is aimed at high dimensional rarefied flow problems, where quantum encoding may reduce memory requirements. However, the method treats only the collisionless/free-molecular regime, so it does not include the nonlinear collision physics required for general hydrodynamic or Navier Stokes behaviour.

Budinski [2021] then proposed a quantum lattice Boltzmann algorithm in which the individual steps of the classical LBM method are mapped onto quantum circuit operations. His first construction treated the advection-diffusion equation using an

LBM inspired quantum algorithm, while later work extended the approach to a stream function vorticity formulation of the Navier Stokes equations [Ljubomir, 2022]. In these approaches, the algorithm follows the structure of classical LBM. One encodes a mesoscopic distribution, applies collision operations, implements streaming, and recovers macroscopic quantities. Although this closely mirrors the classical LBM workflow, it also inherits some of the central difficulties of QLBM, including the cost of state preparation, the depth of the streaming and collision circuits, and the need to extract classical information from the quantum state.

A particularly important issue in QLBM is the treatment of the nonlinear collision operator. Classical LBM uses a state dependent equilibrium distribution, and this state dependence is what gives rise to nonlinear hydrodynamic behaviour such as that found in Navier Stokes dynamics. Since unitary quantum evolution is linear, one common strategy is to linearise the nonlinear collision step before implementing it quantum mechanically. Carleman linearisation provides one such route. It embeds a nonlinear system into a higher dimensional linear system, allowing linear quantum evolution techniques to be applied. Itani et al. [2023] proposed a quantum algorithm for LBM simulation of incompressible fluids with a nonlinear collision term using this type of Carleman linearisation approach. This is an important development because it addresses the nonlinear collision problem more directly than purely linear advection-diffusion QLBM models. However, it also illustrates the cost of this strategy. The nonlinear dynamics are represented only after lifting the system to a larger linear space, so truncation, resource scaling, and state preparation become important limitations.

Another recent Carleman linearisation approach was proposed by Zamora et al. [2026], who developed a quantum lattice Boltzmann algorithm for several timesteps using a local encoding of the Carleman variables. Rather than directly implementing the nonlinear LBM collision operator, the method again uses Carleman linearisation to embed the nonlinear dynamics into a larger linear system, which can then be mapped onto a quantum circuit. The work focuses on a D2Q9 lattice Boltzmann model and introduces an encoding that preserves locality of the collision rule, improving the scaling of the circuit compared with earlier nonlocal Carleman QLBM constructions. The algo-

rithm combines state preparation, collision, propagation, and measurement steps, and the authors validate the method against classical LBM simulations, including a Taylor Green vortex test case. This approach is important because again it addresses the nonlinear collision problem more directly than purely linear advection-diffusion QLBM models. However, it also highlights the remaining difficulties of Carleman based QLBM. The method requires an enlarged Carleman state space, the collision operator is still non-unitary and must be embedded, and the probability of measuring the correct result remains low, of order 10^{-2} in the examples considered. In addition, state preparation and measurement remain significant bottlenecks for general fluid fields. Thus, while the method is a valuable step towards nonlinear QLBM, this goal has still yet to be reached.

Other works have explored alternative QLBM formulations aimed at reducing measurement overhead and improving resource efficiency. Xu et al. [2025] proposed an ancilla-free quantum lattice Boltzmann method for advection-diffusion equations in which nonunitary collision operators are replaced by carefully designed unitary operations. Wawrzyniak et al. [2025] developed a quantum algorithm for the lattice Boltzmann approximation of the linear advection-diffusion equation. Their work is useful because it provides a flexible QLBM framework with separate quantum blocks for initialization, collision, streaming, and recovery of macroscopic quantities. The algorithm is designed to adapt across different spatial dimensions and velocity sets, including D1Q3, D2Q5, D2Q9, and D3Q27, and is validated against both analytical solutions and digital LBM implementations. They also emphasise an important numerical point that the LBM scheme does not solve the exact advection-diffusion equation directly, but rather a modified advection-diffusion equation whose diffusion term is altered by the LBM discretisation. However, the approach still has several limitations. The algorithm assumes the simplified relaxation choice $\tau = 1$, so the collision step reduces to immediate relaxation to equilibrium. It also requires full state measurement and reinitialisation after every timestep, meaning that the method does not preserve a fully coherent quantum evolution over multiple timesteps. This introduces a substantial readout and state preparation overhead and makes the algorithm closer to a hybrid

quantum-classical implementation of LBM than to a purely quantum kinetic model. The authors also note that extending the method to variable relaxation times, non-linear equilibrium distributions, and more general Navier Stokes dynamics remains an important direction for future work.

Another recent approach was proposed by Lăcătuş and Möller [2026], who developed a surrogate quantum circuit design for the lattice Boltzmann collision operator. The core algorithm follows the same lattice Boltzmann structure, with a quantum state which encodes the square root of classical populations, a local collision step updates these populations, and a unitary streaming step shifts them across the lattice. More specifically, the method uses amplitude encoding of populations, a learned unitary circuit acting on the velocity register to approximate the BGK collision operator, measurement to recover classical populations, and reinitialisation for the next timestep. As in previous works, this repeated measurement and reinitialisation breaks global coherence and makes the algorithm effectively hybrid quantum-classical.

A closely related line of work by Di Molfetta and collaborators was discussed in the previous section, where quantum walks were shown to recover modified Dirac, dispersive, and hydrodynamic equations through suitable continuum limits and Madelung transformations [Jolly and Di Molfetta, 2023, Hatifi et al., 2019]. These works are relevant here because they also treat quantum walks as microscopic lattice models for emergent transport dynamics.

A recent QLBM work has begun to focus not only on formal algorithm design, but also on hardware realisability. Tiwari et al. [2025] identify several obstacles in previous QLBM proposals, including the cost of preparing the initial density function, the high circuit depth required for streaming and collision operations, and the inefficiency of full state readout during the time evolution. To address these limitations, they introduce a tensor network based state preparation method, an improved streaming operator using one hot encoding of the direction register, a revised collision operator based on linear combinations of unitaries, observable based readout to avoid full state tomography, and error mitigation. Their work demonstrates a 2D advection-diffusion example on IonQ trapped ion hardware and extends the method in simulation to 3D and non uniform ve-

locity fields. This represents an important step toward hardware QLBM design, as it is the first example of QLBM being implemented on a real quantum computer. However, this hardware oriented QLBM approach also highlights the remaining limitations of the field. To keep circuits shallow enough for NISQ hardware, the state must often be read out, approximated, and reloaded between batches of timesteps. The demonstrated hardware example focuses on smooth Gaussian advection-diffusion, where the density can be efficiently parameterised, rather than on fully general fluid fields. Nonlinearity, complex boundary conditions, long time evolution, and full field reconstruction remain significant challenges. Thus, although this work makes QLBM substantially more practical, it still operates within a restricted transport setting and does not yet provide a general quantum advantage for nonlinear CFD.

Schalkers and Möller [2025] address an important readout problem, as many QLBM algorithms prepare a quantum state representing the flow field, but do not provide an efficient way to extract useful physical quantities without reconstructing the full state. They introduce a quantum momentum exchange method for measuring the force on an immersed object, together with quantum bounce back boundary conditions, providing a route to relevant observables such as drag and lift.

Another recent direction is quantum algorithms for another particle based fluid method. Au-Yeung et al. [2024b] developed a quantum algorithm for smoothed particle hydrodynamics (SPH), where the SPH kernel sums are encoded into quantum registers and evaluated through inner products. The work is similar in spirit to the approach taken here because it uses a quantum algorithm as a route to simulating PDEs, in particular advection and diffusion, rather than attempting to solve the full nonlinear Navier Stokes equations directly. However, the mechanism is different, as their method is based on quantum encoded SPH interpolation operators. They assess their algorithm by comparing quantum SPH simulations with analytical solutions and classical numerical schemes, using RMS error for advection and direct comparison with finite difference diffusion profiles. In this sense, the work provides a useful point of comparison, because both approaches focus on whether quantum dynamics can reproduce the correct macroscopic transport behaviour before moving to more complicated

fluid models.

Across these approaches, two broad limitations appear repeatedly. Much of the work on QLBM algorithms either linearises the collision step of the LBM method or restricts attention to equations, such as advection-diffusion, where the collision structure is already effectively linear. This approach makes the quantum algorithm easier to design, but does not capture the full nonlinear fluid flow behaviour needed for useful large scale Navier Stokes simulation. Other methods use measurement, post selection, or explicit classical feedback as a way to introduce nonlinearity or recover classical populations. These methods can reproduce meaningful transport behaviour, but they require measurement and reinitialisation at every timestep or after short batches of timesteps. This breaks global coherence, reduces the possibility of quantum advantage, and introduces costly state preparation and readout overheads. Thus, although these methods provide important conceptual and algorithmic progress, they also highlight the difficulty of combining coherent quantum evolution with the nonlinear and dissipative structure of classical fluid dynamics.

The present work takes a slightly different perspective. Rather than directly quantising the classical lattice Boltzmann update rule, I instead construct a lattice based PDE simulator starting from quantum walk dynamics themselves. The approach is inspired by the kinetic structure of lattice Boltzmann methods, including discrete velocity channels, streaming, relaxation, and multiscale ordering, but these features emerge naturally from the dynamics of an open quantum walk. In particular, I employ lazy quantum walks generating D1Q3 kinetic structures, together with coin space decoherence acting as a relaxation mechanism and continuum scaling choices controlling the ordering of transport and diffusion terms. Applying a Chapman-Enskog inspired multiscale expansion to these open quantum walks yields effective continuum equations containing advection and diffusive transport contributions.Unlike many existing QLBM proposals, the aim is not to repeatedly measure and reinitialise a classical lattice Boltzmann distribution, nor to implement a state dependent BGK collision operator directly. Instead, the aim is to explore how effective continuum PDEs may emerge directly from open quantum walk dynamics. However, the scope of this thesis

remains deliberately restricted to linear transport equations, primarily diffusion and advection-diffusion. These models provide a controlled setting in which to understand how coherent transport, decoherence, and continuum scaling combine to produce desired macroscopic behaviour. Possible routes towards nonlinear dynamics may include hydrodynamic transformations such as the Madelung representation, or different decoherence models, but these are left as future directions rather than treated as the main target of this work.

The following chapter introduces lazy discrete time quantum walks and derives their continuous spacetime limits in the presence of both coin and spatial decoherence. This provides an important foundation for the rest of the thesis. The inclusion of a rest state gives the walk a D1Q3 structure, which is later used as inspiration for diffusion and advection-diffusion models. The decoherence terms introduced here also form the basis for the open quantum walk models studied in later chapters, where dephasing is used as a relaxation mechanism. By deriving the continuum equations explicitly, the chapter shows how the large scale behaviour of the discrete walk reflects the structure of the underlying lattice model. These results are then built upon in the later Chapman-Enskog analysis and numerical diffusion simulations.

Chapter 3

Continuous Limits of Lazy Open Quantum Walks

In this chapter, I present my analysis of a lazy DTQW subject to two distinct noise channels: one acting on the coin degree of freedom and another on the spatial register. By performing a systematic scaling analysis of the discrete evolution, I derive the corresponding continuum master equation governing the coarse grained state. This provides insight into the behaviour of lazy quantum walks with decoherence and complements existing probabilistic approaches, which establish Gaussian limits but not explicit macroscopic dynamics. For completeness, I derive the continuum limit for both the coin and spatial decoherence models. However, previous studies have shown that coin dephasing preserves spatial coherence while suppressing only internal superpositions, making it more useful for modelling controlled decoherence in quantum transport and algorithmic settings [López and Paz, 2003, Love and Boghosian, 2005, Kendon and Tregenna, 2003]. Here I present the continuous spacetime limit of the lazy three state walk in a full $\mathcal{SU}(3)$ representation, where the unitary and dephasing generators emerge naturally from the Gell Mann structure. This offers a clear route to continuum models for multichannel quantum walks and highlights how internal symmetry and coupling shape large scale dynamics.

The resulting continuum equations are not introduced with the aim of obtaining explicit analytical solutions. Rather, they serve as effective macroscopic descriptions

of the underlying discrete dynamics, providing a bridge between quantum walk models and target partial differential equations. This perspective is particularly relevant in regimes where the continuum equations do not admit compact closed form solutions. The key advantage of this approach is that it is constructive: by establishing a systematic mapping between discrete quantum walk dynamics and continuum PDEs, I obtain a physically motivated and implementable framework for numerical simulation. Specifically, the discrete model itself may be used to simulate the dynamics of the continuum PDE numerically on classical or, potentially, on quantum hardware. This is especially important for more complex systems, where analytical solutions are not accessible and numerical simulation is essential to fully understand the behaviour of the system.

The chapter should therefore be read as a foundation for what follows. The lazy walk is simple enough that the continuum limit can be derived explicitly, but rich enough to contain the essential ingredients needed later. Multiple velocity channels, a rest state, coherent mixing, and controlled decoherence. The limits obtained here are then built upon in the following chapters, where I investigate whether similar ideas can be used to engineer diffusive and fluid inspired behaviour using discrete time quantum walks.

3.1 Lazy Quantum Walks

The variant of quantum walk relevant to this chapter is the lazy quantum walk, originally introduced by Childs [2009b] to demonstrate that the continuous time quantum walk arises as a limit of the discrete model. The inclusion of a rest state introduces a non zero probability of remaining at the current site, modifying the spectrum of the walk operator and enabling a smooth connection between discrete and continuous time dynamics. Unlike the approach of Childs [2009b], where the lazy walk was introduced to embed continuous time dynamics at the spectral level, my analysis presented here includes dephasing, and derives an explicit continuous spacetime master equation by expanding the discrete Lindblad map in space and time. This yields a genuine coarse grained PDE description rather than a spectral correspondence between discrete and continuous models.

Beyond foundational interest, lazy walks have also been employed in algorithmic contexts, such as Grover's search, where self loops at vertices act as effective rest states, influencing success probabilities [Wong, 2015, 2018, Giri and Korepin, 2019]. Including a rest state extends the walk to a three channel model that more closely resembles discrete velocity schemes in kinetic theory [Succi, 2018], such as those used in classical LBM. The stationary component provides a minimal setting where propagating and stationary dynamics coexist [Molfetta and Debbasch, 2012]. This may therefore serve as a foundation for using open quantum walks to model coarse grained transport or hydrodynamic behaviour.

As described in the previous chapter, continuous spacetime limits provide a bridge between discrete quantum walks and continuum physical theories [Manighalam and Kon, 2020, Molfetta and Arrighi, 2019]. By expanding the discrete step operator in powers of a small parameter that controls the time and lattice spacing, one can identify the effective PDE governing large scale dynamics. This approach makes it possible to interpret microscopic unitary rules as discrete simulators of macroscopic behaviour.

A related approach is that of Li et al. [2015], who analyse transport in lazy discrete time quantum walks via long time asymptotics in the absence of decoherence. In that regime, moments of the position distribution can be computed analytically, allowing effective transport coefficients and scaling behaviour to be extracted. In my work, I consider quantum walks with controlled decoherence and focuses on the continuum limit, where the lattice spacing and time step vanish while physical time remains finite. This leads to coarse grained partial differential equations which do not, in general, admit compact closed form solutions. For the purposes of PDE simulation, particularly on quantum hardware, this continuum scaling is the relevant regime, as it captures finite time dynamics rather than asymptotic behaviour. My approach is therefore complementary to that of Li et al. [2015]: asymptotic analysis characterises long time transport of idealised unitary walks, whereas the continuum limit for finite time establishes a direct connection between open quantum walk dynamics and macroscopic PDEs.

Extensions of lazy OQW models have been developed. In particular, Kemp et al.

[2020], constructed a lazy OQW model derived from a Lindblad master equation and proved a central limit theorem demonstrating Gaussian asymptotics with an analytically computable covariance. Such analyses, however, remain microscopic, characterising asymptotic statistics without yielding explicit continuum PDEs for the coarse grained evolution. Tude and de Oliveira [2022] examine multiple forms of decoherence acting on the three state walk and report how these noise mechanisms change its spreading profile, localisation behaviour, and sensitivity to the rest state. Another closely related line of work is that of Sinayskiy and Petruccione [2015], who derived a discrete time open quantum walk from a microscopic system-bath model. Both of these analyses stay within the discrete time framework.

The inclusion of decoherence or dephasing in DTQWs provides a natural bridge between coherent quantum transport and classical diffusive behaviour. In this work, I focus on the intermediate regime, where the quantum walk exhibits partial coherence and remains amenable to continuum analysis. Such limits are particularly interesting because they form the conceptual foundation of more complex open and lazy quantum walks, which have been shown to interpolate between ballistic quantum spreading and diffusive or Gaussian transport [Kemp et al., 2020]. Previous studies have already demonstrated that standard quantum walks can admit a hydrodynamic interpretation in the continuum limit, with emergent hydrodynamic or Madelung type equations governing coarse grained dynamics [Hatifi et al., 2019]. Building on these ideas, dephased or lazy variants provide a controlled way to introduce dissipation, which is central to hydrodynamic and kinetic descriptions of fluids. Consequently, the simple dephased limit provides a controlled setting for probing the onset of hydrodynamic behaviour in quantum lattice models. While the broader application to quantum or quantum inspired fluid algorithms remain open, these results connect naturally to ongoing work on lattice and particle based quantum transport methods [Au-Yeung et al., 2025, Itani and Succi, 2021]

This is also where the connection to lattice Boltzmann methods becomes useful. Classical LBM evolves local populations associated with a discrete set of velocities, and the collision step relaxes these populations towards a local equilibrium distribution.

A lazy three state quantum walk has an analogous minimal velocity structure: left moving, stationary, and right moving components. Coin dephasing may then provides a quantum relaxation mechanism by damping internal coherences and selecting a local equilibrium sector. Although the analogy is not exact, it motivates the use of the lazy walk as a starting point for the more fluid inspired constructions developed in the following chapters.

3.1.1 Lazy Discrete Time Quantum Walks

I consider a lazy DTQW on the line with Hilbert space $\mathcal{H}_{xc} = \mathcal{H}_x \otimes \mathcal{H}_c$. The position space $\mathcal{H}_x$ is spanned by $\{|x\rangle : x \in \Delta_x \mathbb{Z}\}$ with lattice spacing $\Delta_x > 0$. The coin Hilbert space is three dimensional, $\mathcal{H}_c = \text{span}\{|R\rangle, |S\rangle, |L\rangle\}$, corresponding to "right", "stay", and "left" moves. The global state evolves by

$$|\psi(t+1)\rangle = \hat{\mathbf{U}}\,|\psi(t)\rangle = \hat{\mathbf{S}}\,(\mathbb{I} \otimes \hat{\mathbf{C}})\,|\psi(t)\rangle \tag{3.1}$$

where $\hat{\mathbf{C}} \in \mathcal{U}(3)$ is the coin operator, $\hat{\mathbf{S}}$ is the conditional shift, and $\hat{\mathbf{U}}$ is the unitary evolution operator. The conditional shift operator moves the particle left, right, or leaves it in place:

$$\hat{\mathbf{S}} = \sum_x \Big(|x+\Delta_x\rangle\langle x| \otimes |R\rangle\langle R| + |x\rangle\langle x| \otimes |S\rangle\langle S| + |x-\Delta_x\rangle\langle x| \otimes |L\rangle\langle L| \Big) \tag{3.2}$$

A common choice of coin operator is the Grover coin, with three dimensional form

$$\hat{\mathbf{C}}_G = 2|u\rangle\langle u| - \mathbb{I}_3 = \frac{1}{3} \begin{pmatrix} -1 & 2 & 2 \\ 2 & -1 & 2 \\ 2 & 2 & -1 \end{pmatrix}, \tag{3.3}$$

where

$$|u\rangle = \tfrac{1}{\sqrt{3}}(1,1,1)^{\mathsf{T}}. \tag{3.4}$$

The coin operator is as unbiased as possible across all three outcomes [Saha et al., 2021], as it is symmetric under permutations of the three coin states and therefore in-

troduces no directional preference between the left and right moving components. The Grover coin provides a simple and highly symmetric choice for the internal dynamics of the walk. Different coin parametrisations would, in general, lead to distinct effective behaviours in the continuum limit. In the present work, I focus on the symmetric Grover-type coin in order to preserve translational invariance and maintain a balanced coupling between all internal states. This choice ensures that any emergent large scale dynamics arise solely from the interplay between coherent evolution and decoherence, rather than from asymmetries in the underlying coin operation. Nonetheless, the symmetry can be intentionally broken if one wishes to introduce directional bias or model anisotropic effects. The lazy DTQW exhibits three distinct propagation branches, left,

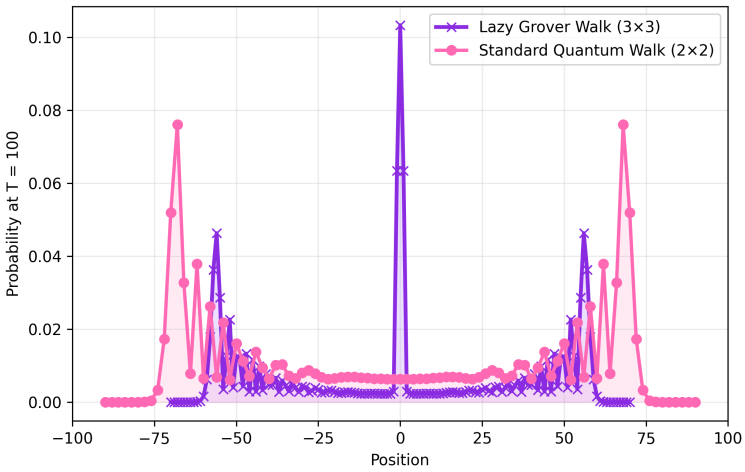

Figure 3.1: Probability distribution of the three–state lazy DTQW and the standard two–state DTQW after $t = 100$ steps. The standard walk uses an initial state $|0\rangle \otimes \frac{1}{\sqrt{2}}(|L\rangle + i\,|R\rangle)$. The initial state of the lazy walk is an equal magnitude Fourier–symmetric superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, with $\omega = e^{2\pi i/3}$. The lazy walk exhibits ballistic side peaks together with a pronounced central rest state peak, illustrating the three channel structure of the dynamics.

right, and a central peak associated with the rest subspace. Figure 3.1 shows the probability distribution of the three state lazy walk. For comparison, a standard two state quantum walk is also shown, generated by a 2×2 Hadamard coin acting on $\{|L\rangle, |R\rangle\}$, following equation (3.1) but with a two dimensional coin space, producing the familiar two ballistic peaks moving symmetrically away from the origin. In contrast, the three

state lazy quantum walk generates a characteristic triplet structure with two ballistic side peaks together with a pronounced central peak associated with the rest state. The presence of this rest component alters both the symmetry and the overall dispersion of the walk, and forms the basic microscopic structure that the continuum limit must reproduce.

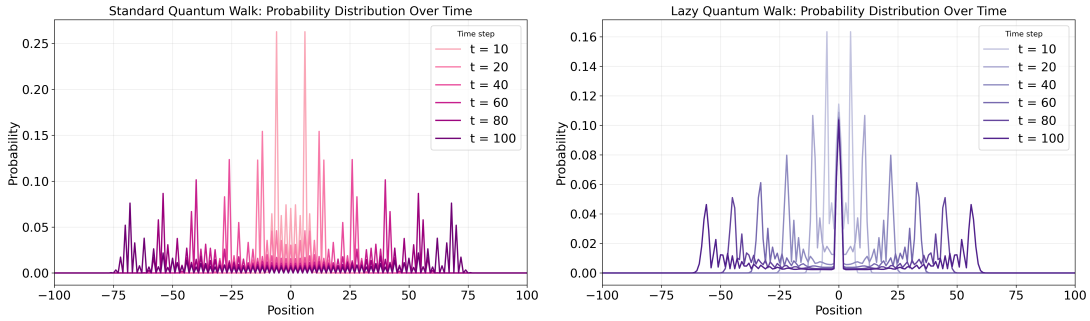

(a) Variance growth.

(b) Mean position

Figure 3.2: Time evolution of the probability distribution of the standard $\mathcal{SU}(2)$ Hadamard walk versus the $\mathcal{SU}(3)$ lazy Grover quantum walk at times $t = 10, 20, 40, 60, 80, 100$. The standard walk uses an initial state $|0\rangle \otimes \frac{1}{\sqrt{2}}(|L\rangle + i\,|R\rangle)$. The initial state of the lazy walk is an equal magnitude Fourier symmetric superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, with $\omega = e^{2\pi i/3}$.

Figure 3.2 shows the time evolution of the position probability distributions for the standard and lazy quantum walks, both initialised at the origin with their respective coin states and evolved without decoherence for up to 100 time steps, providing intuition the differences in their spatial spreading and probability distributions.

3.1.2 Fourier Basis

In order to analyse the continuum limit, it is convenient to work in the momentum representation, where the shift operator becomes diagonal and each momentum mode can be treated independently [Ambainis et al., 2001]. I therefore first transform both operators to Fourier space, before expressing them in the $\mathcal{SU}(3)$ Gell–Mann basis used for the continuum expansion. To obtain the momentum–space form of the walk, I apply

the spatial Fourier transform:

$$\tilde{\boldsymbol{\psi}}(k,t) = \sum_x e^{ikx}\,\boldsymbol{\psi}(x,t), \tag{3.5}$$

to the real space update rule implied by (3.1). Here k is the dimensionless lattice wavenumber, related to the the physical momentum p via the de Broglie relation $p = \hbar k$. Using the action of the shift operator in (3.2), the evolution yields the local recurrence:

$$\boldsymbol{\psi}(x,t{+}1) = \hat{\mathbf{M}}_+\,\boldsymbol{\psi}(x{-}\Delta_x,t) + \hat{\mathbf{M}}_0\,\boldsymbol{\psi}(x,t) + \hat{\mathbf{M}}_-\,\boldsymbol{\psi}(x{+}\Delta_x,t), \tag{3.6}$$

where $\hat{\mathbf{M}}_- = |L\rangle\langle L|\hat{\mathbf{C}}_G$, $\quad \hat{\mathbf{M}}_+ = |R\rangle\langle R|\hat{\mathbf{C}}_G$ and $M_0 = |S\rangle\langle S|\,\hat{\mathbf{C}}_G$. Applying the Fourier transform gives:

$$\tilde{\boldsymbol{\psi}}(k,t{+}1) = \Big(e^{+ik\Delta_x}\hat{\mathbf{M}}_+ \;+\; \hat{\mathbf{M}}_0 \;+\; e^{-ik\Delta_x}\hat{\mathbf{M}}_-\Big)\tilde{\boldsymbol{\psi}}(k,t) \tag{3.7}$$

Defining the diagonal momentum space shift operator

$$\tilde{\mathbf{S}}(k) = e^{-ik\Delta_x}|L\rangle\langle L| \;+\; |S\rangle\langle S| \;+\; e^{+ik\Delta_x}|R\rangle\langle R|, \tag{3.8}$$

I therefore obtain the Fourier–space walk operator in the compact form

$$\tilde{\boldsymbol{\psi}}(k,t{+}1) = \tilde{\mathbf{S}}(k)\,\hat{\mathbf{C}}_G\,\tilde{\boldsymbol{\psi}}(k,t). \tag{3.9}$$

The conditional shift operator given by equation (3.2) becomes diagonal in the momentum basis:

$$\tilde{\mathbf{S}}(k) = \begin{pmatrix} e^{-ik\Delta_x} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{ik\Delta_x}, \end{pmatrix} \tag{3.10}$$

which shows explicitly that the $|R\rangle$ component shifts forward, the $|L\rangle$ component shifts backward, and the $|S\rangle$ component remains localised.

The Grover coin, given by equation (3.3) admits a simple spectral decomposition, which enables us to parameterise it as the exponential of a generator. I define $G =$

$\mathbb{I}_3 - |u\rangle\langle u|$ with $|u\rangle$ being defined in equation (3.4), then

$$G = \frac{1}{3}\begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix} \tag{3.11}$$

A convenient parametrisation is given by

$$\hat{\mathbf{C}}(\theta) = e^{-i\theta G} = e^{-i\theta(\mathbb{I}_3 - |u\rangle\langle u|)} = |u\rangle\langle u| + e^{-i\theta}\,(\mathbb{I}_3 - |u\rangle\langle u|). \tag{3.12}$$

In this representation, the eigenstate corresponding to the uniform vector $|u\rangle$ has eigenvalue $+1$, while the two orthogonal Fourier modes acquire a common phase $e^{-i\theta}$. Thus, the coin operator $\hat{\mathbf{C}}(\theta)$ forms a smooth one parameter family that interpolates around the Grover operator. In particular, the standard Grover coin is recovered at the special value $\theta = \pi$, so that $\hat{\mathbf{C}}_G = \hat{\mathbf{C}}(\pi)$. This parametrisation introduces a tunable coin angle that controls the relative phase between the symmetric mode and the two transverse modes, and is therefore the natural small parameter for the continuum expansion. The parameter θ controls the strength of mixing between the three coin states. In the continuum limit this coupling plays the role of an effective mass term. Small θ yields nearly ballistic propagation, while larger values increase local mixing and open a mass-like gap in the dispersion. The full walk operator in momentum space can be written as:

$$\tilde{\mathbf{U}}(k,\theta) = \tilde{\mathbf{S}}(k)\,\hat{\mathbf{C}}(\theta) \tag{3.13}$$

3.1.3 Gell Mann Representation

I now express $\tilde{S}(k)$ and $\hat{\mathbf{C}}(\theta)$ in the $\mathcal{SU}(3)$ Gell Mann basis. This representation makes the underlying generators explicit and is the natural starting point for performing the continuum limit expansion of the walk operator in section 3.3.

For two state quantum walks, the two dimensional coin operators can be decomposed in the Pauli basis, any $\mathcal{SU}(2)$ operator can be written as a linear combination of Pauli matrices. In the lazy quantum walk, the coin space is three dimensional, so the

natural generalisation is to use the eight Gell Mann matrices $\{\lambda_i\}_{i=1}^{8}$, traceless Hermitian operators, which, together with the identity $\mathbb{I}_3$, span the space of 3×3 Hermitian operators [Wilson, 2024]. The explicit form of each Gell Mann matrix can be found in Appendix A.1. They satisfy the orthogonality relation $\mathrm{Tr}(\lambda_i \lambda_j) = 2\delta_{ij}$, so that any $\mathcal{SU}(3)$ operator $A \in \mathbb{C}^{3\times 3}$ can be expanded as

$$A = \frac{1}{3}\,\mathrm{Tr}(A)\,\mathbb{I}_3 + \frac{1}{2}\sum_{i=1}^{8} \mathrm{Tr}(A\lambda_i)\,\lambda_i. \tag{3.14}$$

Having expressed both the shift and coin in a separable form in Fourier space, we now rewrite these operators in the $\mathcal{SU}(3)$ Gell Mann basis. I first define the $\mathcal{SU}(3)$ generator:

$$J_z^{(3)} \;=\; \tfrac{1}{2}\lambda_3 + \tfrac{\sqrt{3}}{2}\lambda_8 \;=\; \mathrm{diag}(1, 0, -1) \tag{3.15}$$

Then the shift operator, given by equation (3.10), may be written as:

$$\tilde{\mathbf{S}}(k) = \exp\big[-\,i\,\Delta_x k\,(\mathbb{I}_x \otimes J_z^{(3)})\big]. \tag{3.16}$$

For notational simplicity, I omit tensor product identity operators $\mathbb{I}_x$ and $\mathbb{I}_c$ from this point onwards, writing $J_z^{(3)}$ and G to denote the corresponding operators acting on the coin subspace.

The symmetric action of the Grover type coin is captured by the generator G given by equation (3.11). The generator G may be written in the Gell Mann basis as:

$$G \;=\; \frac{2}{3}\mathbb{I}_3 - \frac{1}{3}\big(\lambda_1 + \lambda_4 + \lambda_6\big), \tag{3.17}$$

which mixes the three internal states.

The above representations are suitable to be used in the expansion of the unitary operator in the continuum limit (section 3.3). Having established the basic formalism of the lazy DTQW, the next subsection introduces decoherence and outlines a model employed to incorporate it into the quantum walk framework.

3.2 Decoherence Model

Decoherence is usually described as a non-unitary evolution using the density matrix formalism to describe the mixed states that arise. Simple decoherence models for quantum walks were introduced in Kendon and Tregenna [2003]. Here, I use a similar simple Lindblad master equation to examine the basic features of decoherence in the lazy quantum walk, providing a baseline against which more complex decoherence mechanisms can be compared. This approach assumes a Markovian open system description, in which the environment has a short correlation time and no memory effects are retained. The Lindblad master equations provides dynamics that are completely positive and trace preserving, ensuring that the density operator remains physical at all times, regardless of the step size, as I take the limits [Annabestani et al., 2010]. The coupling to the environment is treated phenomenologically through a single decoherence strength γ, without specifying a microscopic bath model or spectral density. Initial system-environment correlations are neglected, and the system follows a weak decoherence regime [Breuer and Petruccione, 2002].

Section 2.3.3 has already introduced the density matrix formalism required to describe the DTQW with decoherence. The system is again described by the time dependent density operator, and a single timestep discrete time evolution is described using Eqn. (2.105). The decoherence is modelled as projective noise using jump operators constructed as

$$L_j = \mathbb{P}_j \, \hat{\mathbf{U}}, \tag{3.18}$$

where the projectors $\{\mathbb{P}_j\}$ specify the particular decoherence channel. For coin decoherence, the projectors act on the coin space, so that off-diagonal coherences between coin states are suppressed while diagonal populations are preserved. For spatial decoherence, the projectors act on the position space, which destroys coherence between different spatial locations while leaving the coin degree of freedom unaffected. In both cases the decoherence parameter γ controls the strength of the noise, setting the probability of a projective decoherence event at each timestep. The projectors satisfy the

usual relations

$$\mathbb{P}_j^\dagger = \mathbb{P}_j, \quad \mathbb{P}_j^2 = \mathbb{P}_j, \quad \sum_j \mathbb{P}_j = \mathbb{I}. \tag{3.19}$$

At this stage, the projectors are kept general. Specific choices corresponding to coin and spatial decoherence explored in more detail in Section 3.3.3, although these represent only two examples of possible decoherence channels. Decoherence modifies the characteristic spreading peaks of the unitary quantum walk, and potentially induces a crossover to classical behaviour, depending on the model, strength and type of environmental coupling [Kendon, 2007]. The continuum limit of equation (2.105) may be taken after specifying the structure of the unitary operator and decoherence model. I first recap the effects of both coin and spatial decoherence on the discrete time lazy quantum walk.

Figure 3.3 compares the effects of coin and spatial dephasing on the lazy DTQW for several values of the decoherence strength γ and for different coin angles θ. Coin dephasing, shown in panels (a,c), selectively damps the off-diagonal coherences between the internal coin states $\{|L\rangle, |S\rangle, |R\rangle\}$ while leaving spatial coherences intact. In contrast, spatial dephasing, shown in panels (b,d), suppresses off-diagonal density matrix elements $\rho(x, x')$ for $x \neq x'$. As γ increases, the ballistic side peaks collapse more rapidly than in the coin dephased case and the probability distribution evolves more rapidly toward a broad, approximately Gaussian profile centred at the origin. This transition occurs already at relatively small values of γ, reflecting the sensitivity of ballistic transport to spatial coherence. For the spatially dephased case, the probability distribution seems to flatten out at low γ values before tending towards the Gaussian profile. This is consistent with the behaviour of of the two state quantum walk under projective decoherence. In the limit $\gamma \to 1$, both decoherence mechanisms yield a similar Gaussian distribution, corresponding to a fully classical diffusive regime. Comparing different coin angles further reveals the role of the coin in controlling transport and localisation. For $\theta = \pi$, the unitary walk exhibits a pronounced central peak, indicating that the coin strongly favours population of the rest state and suppresses propagation. For $\theta = \pi/2$, the coin induces stronger mixing between the propagating

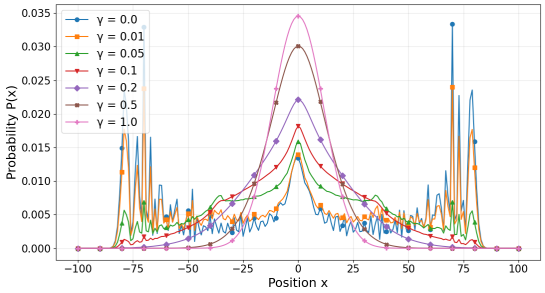

(a) Probability distribution for the case of coin decoherence over several values of γ for coin angle $\theta = \frac{\pi}{2}$.

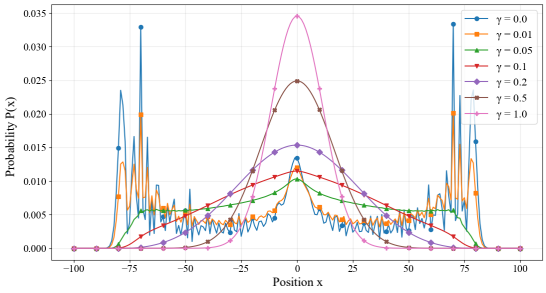

(b) Probability distribution for the case of spatial decoherence over several values of γ for coin angle $\theta = \frac{\pi}{2}$.

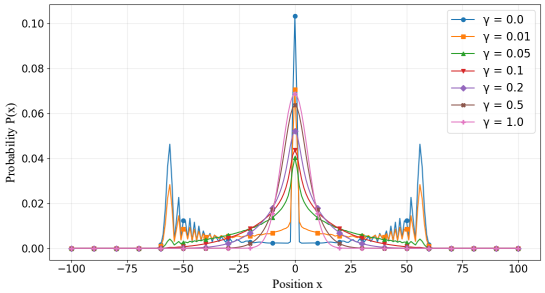

(c) Probability distribution for the case of coin decoherence over several values of γ for coin angle $\theta = \pi$.

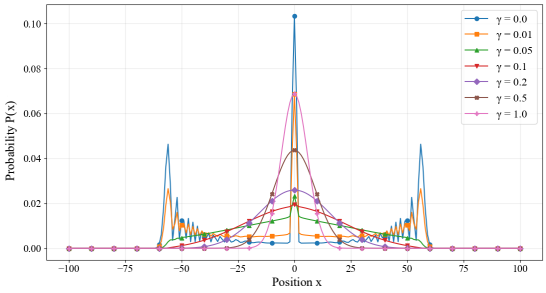

(d) Probability distribution for the case of spatial decoherence over several values of γ for coin angle $\theta = \pi$.

Figure 3.3: Effect of dephasing on the lazy DTQW for different coin angles θ and decoherence strengths γ. Panels (a,c) show coin dephasing and panels (b,d) show spatial dephasing, for $\theta = \pi/2$ (top row) and $\theta = \pi$ (bottom row). Coin dephasing has projector choice $\mathbb{P}_L = |L\rangle\langle L|, \mathbb{P}_S = |S\rangle\langle S|, \mathbb{P}_R = |R\rangle\langle R|$, while spatial dephasing uses $\mathbb{P}_x = |x\rangle\langle x| \otimes \mathbb{I}_c$. The initial state is a localised position state with a Fourier symmetric coin superposition, $|\psi_0\rangle = |0\rangle \otimes \frac{1}{\sqrt{3}}\big(|L\rangle + \omega\,|S\rangle + \omega^2\,|R\rangle\big)$, where $\omega = e^{2\pi i/3}$.

and non-propagating components, resulting in a less pronounced central peak and a higher probability of reaching the ballistic fronts. This behaviour persists under both decoherence mechanisms, although spatial dephasing again accelerates the transition toward a Gaussian profile. Overall, increasing θ modifies the internal coin mixing and interference structure of the walk, leading to systematically reduced spatial spreading and enhanced localisation near the origin. Note that the difference in smoothness between the probability distributions in Figures 3.3 and 3.2 arises from their initial conditions. Figure 3.3 uses a smooth initial state, whereas Figure 3.2 uses a localised initial state, producing sharper peaks due to quantum interference.

3.3 Continuum Limit of the Lazy Quantum Walk with Decoherence

The previous section introduced the discrete Lindblad formulation of the lazy quantum walk with decoherence. In this section, that framework is applied to specific choices of decoherence models, and their continuum limits are derived to identify the effective dynamics simulated by each case. The following subsections derive the continuous spacetime limit of lazy quantum walks with both coin and spatial decoherence. Although both models share the same underlying unitary dynamics, they differ in the structure of the Lindblad term, leading to a different effective PDE in the limit.

3.3.1 Scaling and Parameterisation

The continuum limit of equation (2.104) is obtained by Taylor expanding the discrete Lindblad evolution operator in powers of a small parameter ε. This parameter simultaneously scales the lattice spacing, timestep, and coin angle. The continuum limit is then triggered by taking $\varepsilon \to 0$. Following the standard approach [Molfetta, 2021], I set:

$$\begin{aligned} \Delta_t &= \varepsilon \\ \Delta_x &= \varepsilon \\ \theta &= \varepsilon\,\bar{\theta} \end{aligned} \tag{3.20}$$

where $\bar{\theta}$ is a constant and remains finite in the limit. The decoherence rate is also weakly scaled:

$$\gamma = \varepsilon\,\bar{\gamma}, \tag{3.21}$$

This scaling corresponds to the weak decoherence regime, in which the noise acts only as a small perturbation on the otherwise coherent dynamics. The constant $\bar{\gamma}$ remains finite in the limit and represents the effective decoherence rate per unit time step. If instead, the unscaled rate γ were kept fixed as $\varepsilon \to 0$, the decoherence would remain strong in the continuum limit, causing the Lindbladian term to dominate over the

unitary advection and mixing contributions. Such a regime leads to rapid relaxation and trivial stationary behaviour, and is therefore not considered further here. With the scaling defined above, I can now expand the unitary evolution operator to obtain the effective Hamiltonian governing the coherent part of the dynamics.

3.3.2 Unitary Expansion and Effective Hamiltonian

Before introducing decoherence, the unitary part of the continuum limit is derived. This component is common to both models considered below, since each decoherence mechanism employs the same underlying shift and coin operators. The differences between models arise solely through the choice of projectors acting within the Lindblad term, while the unitary evolution remains identical. The discrete unitary step operator is defined as:

$$\tilde{\mathbf{U}}(k) = \tilde{\mathbf{S}}(\varepsilon, k)\, \hat{\mathbf{C}}(\varepsilon, \bar{\theta}), \tag{3.22}$$

where both the shift and coin operators are expressed as exponentials of $\mathcal{SU}(3)$ generators. The coin operator given by equation (3.12) using the generator given by equation (3.11) may now be expanded around ε, as:

$$\hat{\mathbf{C}}(\varepsilon, \bar{\theta}) = e^{-i\varepsilon\bar{\theta}G} \approx \mathbb{I}_3 - i\varepsilon\bar{\theta}G - \frac{\varepsilon^2\bar{\theta}^2}{2}G^2 + \mathcal{O}(\varepsilon^3) \tag{3.23}$$

The shift operator, given by equation (5.11), can then be expanded around ε as:

$$\tilde{\mathbf{S}}(\varepsilon, k) = e^{-i\varepsilon k J_z^{(3)}} \approx \mathbb{I}_3 - i\varepsilon k J_z^{(3)} - \frac{\varepsilon^2 k^2}{2}\big(J_z^{(3)}\big)^2 + \mathcal{O}(\varepsilon^3) \tag{3.24}$$

Evaluating the first order expansion of $\tilde{\mathbf{U}}(k) = \tilde{\mathbf{S}}(\varepsilon, k)\hat{\mathbf{C}}(\varepsilon, \bar{\theta})$ small values of k and $\bar{\theta}$, the unitary operator can be expanded around the identity to find the Hamiltonian $\hat{\mathbf{H}}_{\text{phys}}$:

$$\tilde{\mathbf{U}}(k) \simeq \mathbb{I}_3 - i\varepsilon\, \hat{\mathbf{H}}_{\text{phys}} + \mathcal{O}(\varepsilon^2), \tag{3.25}$$

where $\hat{\mathbf{H}}_{\text{phys}}$ is the effective Hamiltonian governing the continuum time evolution. The unitary Hamiltonian is then found to be:

$$\hat{\mathbf{H}}_{\text{phys}}(k) = k\, J_z^{(3)} + \bar{\theta}\, G = k\Big(\tfrac{1}{2}\lambda_3 + \tfrac{\sqrt{3}}{2}\lambda_8\Big) + \frac{\bar{\theta}}{3}\big(\lambda_1 + \lambda_4 + \lambda_6\big), \tag{3.26}$$

where $J_z^{(3)}$ generates directional propagation along the lattice and G produces symmetric mixing between the three internal states. The three state twistless Hamiltonian can equivalently be written in terms of spin-1 operators as

$$H_{\text{tw},3}^{(2)} = 4\bar{\alpha}k\{J_x, J_y\} + \sqrt{2}\bar{\eta}\left(\{J_x, J_z\} + \{J_y, J_z\}\right), \tag{3.27}$$

where $\{A, B\} = AB + BA$ denotes the anticommutator. Here, J_x, J_y, and J_z are the spin-1 angular momentum operators, given in the $(|L\rangle, |S\rangle, |R\rangle)$ basis by

$$J_x = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad J_y = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & -i & 0 \\ i & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \quad J_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{3.28}$$

The Hamiltonian contains quadratic spin couplings, representing coherent mixing between the moving and stationary coin states.

To obtain a real space description, I invert the Fourier representation of the walk and take the continuum limit, in which the quasi-momentum k enters only through smooth phase factors generated by lattice translations. Expanding these phase factors to leading order in k and transforming back to position space maps the momentum dependence of the generator onto spatial derivatives acting on the wavefunction. As a consequence, the momentum variable is represented in real space by the differential operator $k \to -i\,\partial_x$, yielding the effective real space Hamiltonian

$$\hat{\mathbf{H}}_{\text{phys}} = -i\, J_z^{(3)}\, \partial_x + \bar{\theta}\, G, \tag{3.29}$$

so that the corresponding continuum dynamics satisfy

$$\partial_t \psi(x,t) = -i\hat{\mathbf{H}}_{\text{phys}}\psi(x,t), \tag{3.30}$$

which describes ballistic advection along the lattice direction accompanied by local symmetric mixing in the internal state space. In the massless limit $\bar{\theta} = 0$, the Hamiltonian generates purely linear propagation of the left and right moving components with dispersion $\omega = \pm k$, while the lazy channel $|S\rangle$ remains dynamically decoupled. The resulting dynamics are formally analogous to the one dimensional propagation of a free, massless spin-1 field, with the two propagating components playing the role of helicity eigenmodes. In this sense, the unitary $\mathcal{SU}(3)$ walk provides an analogue of photon-like propagation in one spatial dimension [Bialynicki-Birula, 1996]. It is important to note, however, that this analogy is kinematic rather than exact. Unlike the electromagnetic field, the present model does not possess gauge invariance or an associated transversality constraint, and the additional internal component is therefore retained as a physical degree of freedom rather than being removed by a gauge redundancy. The lazy channel $|S\rangle$ thus acts as a longitudinal or rest mode that extends the minimal two component Dirac structure to a three component multiplet, rather than representing an unphysical gauge degree of freedom.

In terms of field theories, the parameter θ plays the role of an effective mass. This interpretation follows directly from the structure of the Hamiltonian in (3.29). The kinetic term $-iJ_z^{(3)}\partial_x$ generates advection of the left and right moving components, while the $\bar{\theta}, G$ term induces local mixing in coin space that couples these propagation channels. In two-component quantum walks, such mixing terms are responsible for opening an effective mass gap in the dispersion relation. The same mechanism appears here in its $\mathcal{SU}(3)$ generalisation, as the operator G couples the three internal states symmetrically and acts as the analogue of a mass generating interaction. Consequently, the parameter $\bar{\theta}$ controls the strength of internal-state coupling and determines the opening of an effective mass gap. Small values of $\bar{\theta}$ yield nearly massless, strongly dispersive behaviour with sharply separated ballistic branches, while larger θ enhances mixing and progressively suppresses coherent propagation, leading to a flattening of

the dispersion relation. This trend is visible from the $\gamma = 0$ lines in Fig. 3.3, where increasing the effective mass parameter $\bar{\theta}$ suppresses ballistic propagation and reduces the maximum distance reached at fixed N, whereas smaller $\bar{\theta}$ corresponds to lighter, more freely propagating dynamics. The Hamiltonian $\hat{\mathbf{H}}_{\text{phys}}$ therefore represents the minimal $\mathcal{SU}(3)$ embedding of a Dirac-type continuum structure, with $J_z^{(3)}$ supplying the kinetic generator and G providing the mass coupling.

This structure defines the universal unitary backbone of the model: all differences between the decoherence channels developed below arise exclusively from how their respective Lindblad projectors act on this shared $\mathcal{SU}(3)$ Hamiltonian, while the underlying effective field theory description of the unitary dynamics remains the same.

3.3.3 Including Decoherence

I now incorporate decoherence into the discrete time evolution. The dynamics are governed by the Lindblad CPTP map given in Eqn. (2.105), with jump operators defined in Eqn. (3.18). Applying the chosen scaling with respect to the small parameter ε, I obtain the discrete Lindblad update rule, which I will subsequently expand to derive its continuum limit:

$$\hat{\boldsymbol{\rho}}(t+\varepsilon) = \hat{\boldsymbol{\rho}}(t) - i\varepsilon[\hat{\mathbf{H}}, \rho(t)] + \varepsilon\bar{\gamma}\sum_j \Big(\mathbb{P}_j\tilde{\mathbf{U}}\hat{\boldsymbol{\rho}}(t)\tilde{\mathbf{U}}^\dagger\mathbb{P}_j^\dagger - \tfrac{1}{2}\big\{\tilde{\mathbf{U}}^\dagger\mathbb{P}_j^\dagger\mathbb{P}_j\tilde{\mathbf{U}}, \hat{\boldsymbol{\rho}}(t)\big\}\Big) \tag{3.31}$$

where the projectors $\mathbb{P}_j$ depend on whether decoherence acts on the coin or position subspace. The same unitary generator $\hat{\mathbf{H}}_{\text{phys}}$ derived above appears in both cases. The following subsections derive the continuum limit for each model, emphasising that only the Lindblad term differs.

Coin Decoherence

I begin by considering the case in which decoherence acts solely on the coin subspace. In this setting, the projectors that define the decoherence channel act within the internal coin basis:

$$\mathbb{P}_L = |L\rangle\langle L|, \qquad \mathbb{P}_S = |S\rangle\langle S|, \qquad \mathbb{P}_R = |R\rangle\langle R|. \tag{3.32}$$

These projectors are substituted into the discrete Lindblad equation and I expand in ε to first order. I take the continuum limit and and retain the leading terms in ε. The real space description is obtained by inverting the Fourier representation, and consequently the momentum is replaced by the differential operator $k \mapsto -i\,\partial_x$ in $\hat{\mathbf{H}}_{\text{phys}}$. For clarity, it is convenient to introduce the projector onto the diagonal in the coin basis

$$\Pi_c[\boldsymbol{\rho}] \;=\; \text{diag}(\boldsymbol{\rho}), \tag{3.33}$$

which keeps only the diagonal coin populations in the $\{|L\rangle, |S\rangle, |R\rangle\}$ basis. The dephasing channel can then be written as

$$D_c[\boldsymbol{\rho}] \;=\; \boldsymbol{\rho} - \Pi_c[\boldsymbol{\rho}], \tag{3.34}$$

so that $D_c[\boldsymbol{\rho}]$ vanishes on populations and acts solely on coherences between the internal states. With this notation, the full coarse grained continuum equation reads:

$$\partial_t \boldsymbol{\rho}(x,t) \;=\; [\boldsymbol{\rho}(x,t), \partial_x J_z^{(3)}] \;-\; i\,\bar{\theta}\,[G,\, \boldsymbol{\rho}(x,t)] \;-\; \bar{\gamma}\, D_c[\boldsymbol{\rho}(x,t)], \tag{3.35}$$

where the coherent part is expressed in the Gell Mann basis via the fixed $\mathcal{SU}(3)$ generators $J_z^{(3)}$ and G, given by equations (3.15) and (3.17) respectively. The first commutator in Eq. (3.35) represents the kinetic (streaming) contribution, with the spatial derivative acting on $\boldsymbol{\rho}$ and $J_z^{(3)} = \text{diag}(1, 0, -1)$ selecting the left, rest, and right propagation channels. The second term implements their symmetric local mixing. The final term, $-\bar{\gamma}\, D_c[\boldsymbol{\rho}]$, damps only the off diagonal coin coherences, leaving the coin populations unchanged at leading order.

Eqn. (3.35) is linear, first order in time, and local in x, with the spatial derivative entering only through the commutator with $J_z^{(3)}$. It preserves Hermiticity and trace, and for $\bar{\gamma} \geq 0$ generates a CPTP semigroup on the coin space at each spatial point. The unitary part retains the full $\mathcal{SU}(3)$ structure of the lazy walk, while the dephasing produces a uniform decay of the six off diagonal Gell–Mann components $\{\lambda_1, \lambda_2, \lambda_4, \lambda_5, \lambda_6, \lambda_7\}$ and leaves the diagonal sector (λ_3, λ_8) intact. Consequently, the continuum dynamics

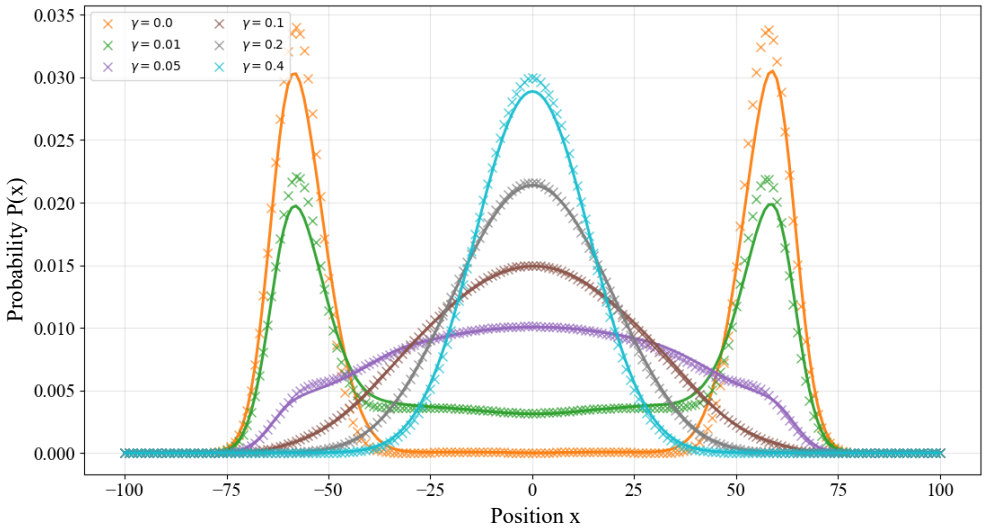

Figure 3.4: Discrete time quantum walk versus derived continuum master equation for coin dephasing. Final position probability $P(x,t)$ after $N = 100$ steps of the lazy three state quantum walk with $\theta = \pi/2$, shown for several decoherence rates γ. Solid lines show the solution of the corresponding derived continuum Lindblad equation. The initial state is a spatially delocalised Gaussian wavepacket of width $\sigma_0 = 5$ with a Fourier-symmetric coin superposition $(1, \omega, \omega^2)/\sqrt{3}$, where $\omega = e^{2\pi i/3}$.

preserve the underlying streaming structure while progressively reducing the interference between the left, rest, and right channels in space and time. The absence of diagonal contributions in D_c indicates that the decoherence channel does not directly modify the populations of the coin states at leading order, but instead damps their mutual phase coherences. Physically, this implies that probability weight is redistributed between internal levels solely through coherent dynamics, whereas interference between coin states is progressively lost. As a result, the coherent oscillations characteristic of the unitary walk are gradually attenuated, and the system evolves towards an effectively classical mixture in the coin subspace, providing a microscopic mechanism for the emergence of classical behaviour in the macroscopic limit.

Figure 3.4 shows the position probability distributions obtained from DTQW and from the corresponding derived continuum equation for several values of the coin dephasing parameter γ. The figure demonstrates very good agreement between the DTQW and the continuum Lindblad dynamics throughout the range of weak to moderate dephasing considered here ($0 \leq \gamma \leq 0.4$), which is the weak decoherence regime

in which the continuum master equation is derived. Within this range, the continuum solution accurately reproduces both the overall shape and the width of the position distribution produced by the discrete walk. In particular, the crossover from ballistic, multi peaked propagation at weak dephasing to a smoother, more centrally concentrated profile as the dephasing strength increases is captured faithfully. This agreement provides strong evidence that the DTQW with coin dephasing constitutes a consistent and controlled discretisation of the target continuum open quantum dynamics in its domain of validity.

In Figure 3.4, the continuum density matrix equations are solved numerically using a standard fourth order Runge Kutta method. In general, these equations do not admit simple closed form solutions suitable for direct evaluation, and numerical integration is therefore required to obtain the probability distributions shown. The continuum model is discretised on the same spatial grid as the discrete quantum walk, and full details of the numerical implementation and parameters are provided in Appendix A.2. The results are not sensitive to the specific numerical scheme, and equivalent behaviour is obtained using any standard differential equation solver. The observed differences between the discrete and continuum models are consistent with standard finite resolution effects. The relative error remains approximately uniform across the bulk, while larger absolute deviations occur in regions of steep gradients. In particular, the approximation is most accurate in the weak decoherence regime, as expected from the derivation of the continuum limit.

The discrete quantum walk is, by construction, a discrete approximation of the continuum model, and therefore a non-zero discrepancy is expected. In particular, the error is largest in regions where the probability distribution exhibits steep gradients. This is a standard feature of finite difference approximations. Sharper features are more difficult to resolve on a fixed lattice and therefore incur larger local errors. Conversely, in regions where the distribution varies slowly, the agreement is very good.

Figure 3.5 shows the absolute error $|P_{\text{cont}} - P_{\text{disc}}|$ for the coin dephasing case. As expected, the largest discrepancies occur in regions where the underlying distributions have the steepest gradients, while the error remains small in the bulk where the distri-

butions are smooth. I also consider the pointwise fractional error $\frac{|P_{\text{cont}}-P_{\text{disc}}|}{P_{\text{disc}}}$, shown in

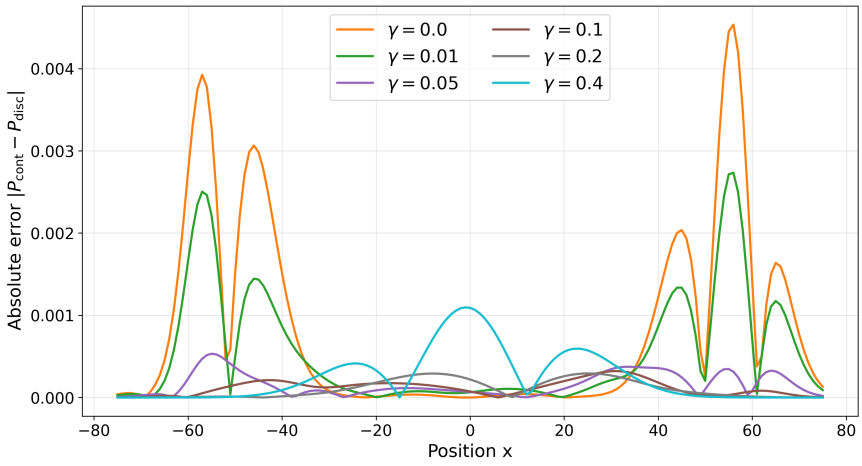

Figure 3.5: Absolute error $|P_{\text{cont}} - P_{\text{disc}}|$ for coin dephasing.

Fig. 3.6. Note that the scaling is slightly different on the fractional error graph. This is because the tails of the distribution are truncated: in these regions the probability density is extremely small, and the fractional error becomes artificially large due to division by values close to zero. These contributions are therefore not physically meaningful and are excluded for clarity. For all $\gamma > 0$, the fractional error remains small (typically below 1%) across the region where the probability density is appreciable. The apparent increase in fractional error at the edges of the domain, particularly for larger γ, is not physically significant. In these regions, both the discrete and continuum probabilities are extremely small, and the fractional error becomes artificially amplified due to division by a vanishingly small denominator. The corresponding absolute error remains negligible and does not affect any physically relevant observables.

For the $\gamma = 0$ case, two pronounced peaks are observed in the fractional error. These arise from small differences in the decay of the discrete and continuum distributions in low probability regions. Importantly, these peaks occur where the probability density is already very small and therefore do not reflect a significant discrepancy in the bulk behaviour. The larger absolute errors near the peaks reflect the greater magnitude of the distribution in these regions, while the fractional error remains approximately constant across the bulk . The largest spikes in fractional error instead occur where the

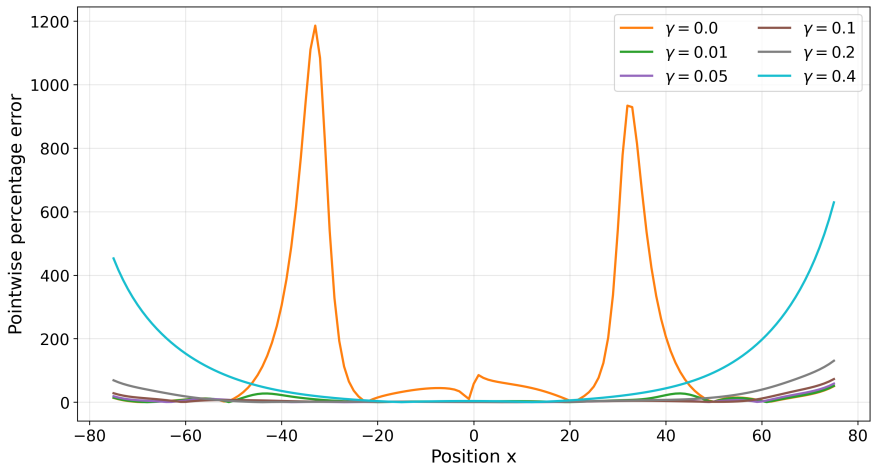

Figure 3.6: Coin dephasing fraction error for discrete vs continuum models.

probability density tends to zero, where division by very small values artificially inflates the error somewhat. The exact same behaviour is observed in the spatial dephasing case as well.

I emphasise that the overall behaviour is consistent with that of a controlled approximation. The discrepancy between the discrete and continuum models is governed by the local variation of the distribution and remains small in regions that contribute most to physical observables. The apparent non-monotonic dependence of the error on γ is therefore a consequence of how the shape of the distribution (and in particular its gradients) varies with γ, rather than indicating the existence of a special or optimal parameter.

Spatial Decoherence

I now consider the complementary case in which decoherence acts on the spatial subspace rather than on the coin. The discrete evolution is still governed by the Lindblad map in Eqn. (2.105) with jump operators of the form (3.18), but the projectors now resolve the position basis:

$$\mathbb{P}_x \;=\; |x\rangle\langle x| \otimes \mathbb{I}_c, \tag{3.36}$$

so that each Lindblad (or Kraus) operator isolates a single lattice site while leaving the internal (coin) degree of freedom untouched,

$$L_x \;=\; \mathbb{P}_x\,\tilde{\mathbf{U}}. \tag{3.37}$$

I assume the same weak coupling continuum scaling as in the coin-decoherence case, given by equations (3.20) and (3.21). I then expand the discrete Lindblad step to first order in ε. The discrete Lindblad map is again expanded to first order in ε and the continuous spacetime limit is taken. Again, the real space description is obtained by inverting the Fourier representation, and consequently the momentum is replaced by the differential operator $k \mapsto -\,i\,\partial_x$ in $\hat{\mathbf{H}}_{\text{phys}}$. To make the structure of the dissipator explicit, it is convenient to view $\boldsymbol{\rho}$ in the position representation, with matrix elements

$$\boldsymbol{\rho}_{cc'}(x,x',t) \;=\; \langle x,c|\;\boldsymbol{\rho}(t)\;|x',c'\rangle \tag{3.38}$$

I introduce the projector onto the diagonal in position space,

$$\big(\Pi_x[\boldsymbol{\rho}]\big)_{cc'}(x,x',t) \;=\; \delta_{x,x'}\,\boldsymbol{\rho}_{cc'}(x,x',t), \tag{3.39}$$

which removes all spatial off-diagonal components while preserving the full internal (coin) structure at each site. In analogy with the coin dephasing map (3.34), I define

$$D_x[\boldsymbol{\rho}] \;=\; \boldsymbol{\rho} - \Pi_x[\boldsymbol{\rho}], \tag{3.40}$$

so that $D_x[\boldsymbol{\rho}]$ vanishes on the spatial populations (the blocks with $x = x'$) and acts only on coherences between distinct lattice sites. The dissipator in Eqn. (3.41) then has matrix elements

$$\big(-\,\bar{\gamma}\,D_x[\boldsymbol{\rho}]\big)_{cc'}(x,x',t) = \begin{cases} 0, & x = x', \\ -\bar{\gamma}\,\boldsymbol{\rho}_{cc'}(x,x',t), & x \neq x'. \end{cases}$$

With this notation, the full coarse grained continuum equation reads:

$$\partial_t \boldsymbol{\rho}(x,x',t) \;=\; [\boldsymbol{\rho}(x,x',t), \partial_x J_z^{(3)}] \;-\; i\,\bar{\theta}\,[G,\, \boldsymbol{\rho}(x,x',t)] \;-\; \bar{\gamma}\, D_x[\boldsymbol{\rho}(x,x',t)], \quad (3.41)$$

where $\hat{\mathbf{H}}_{\text{phys}}$ is given in Eqn. (3.29) and the spatial dephasing channel D_x. Here, the first two terms correspond to the density matrix form of equation (3.30), and the final term gives the spatial dissipator. Eqn. (3.41) is again linear, first order in time, and local in t, with the coherent part identical to the coin-decoherence case and the dissipative part now acting purely in the spatial sector. It preserves Hermiticity and trace and, for $\bar{\gamma} \geq 0$, generates a CPTP semigroup on the full coin–position Hilbert space. The Hamiltonian term $-i[\hat{\mathbf{H}}_{\text{phys}}, \boldsymbol{\rho}]$ produces directional streaming and local coin mixing and is the density matrix form of equation (3.30). The dephasing channel $-\bar{\gamma}\, D_x[\boldsymbol{\rho}]$ leaves all on site density matrices $\boldsymbol{\rho}(x,x,t)$ invariant and drives an exponential decay of the off-diagonal blocks $\boldsymbol{\rho}(x,x',t)$ with $x \neq x'$.

Physically, this implies that the decoherence channel acts locally in position space and does not directly transfer probability between lattice sites or modify the internal coin populations at a given site, but instead suppresses coherent superpositions between different positions. As a consequence, although the walk retains its local streaming dynamics, the loss of long-range spatial coherence alters the interference mechanisms responsible for ballistic transport. The resulting evolution therefore differs qualitatively from the fully unitary walk, leading to a redistribution of probability in space. In the long-time limit, the state approaches a classical mixture in the position basis, providing a microscopic mechanism for the emergence of spatially classical transport while still resolving the internal coin degrees of freedom at each site. This makes clear why spatial decoherence is typically undesirable when quantum properties are required: dephasing in the position basis destroys the long-range interference that enables quantum advantage, leaving only the coin subspace quantum but removing the mechanism that produces ballistic transport.

Figure 3.7 shows the final position probability distribution obtained from the DTQW with spatial dephasing (markers) compared to the corresponding continuum Lindblad equation (solid lines), for several decoherence strengths $0 \leq \gamma \leq 0.4$. Again, I restrict

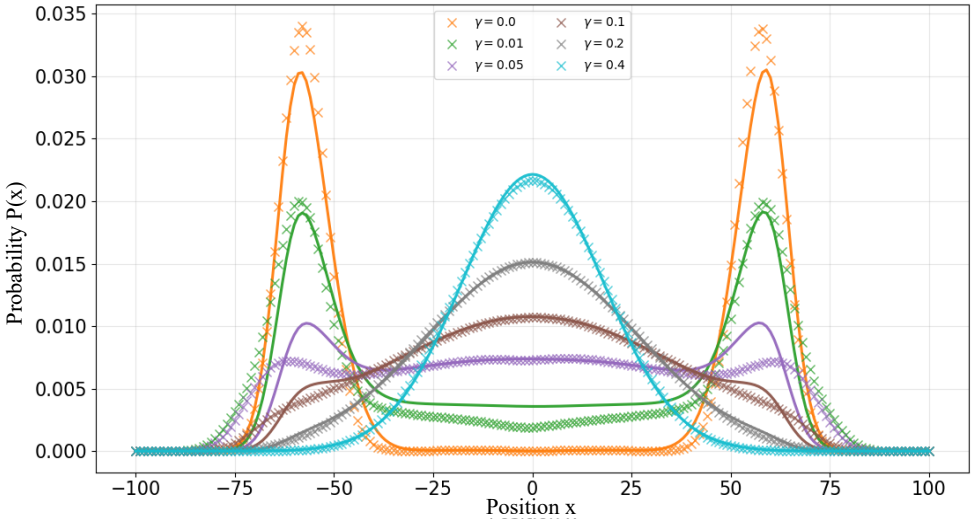

Figure 3.7: DTQWs derived continuum master equation for spatial dephasing. Final position probability $P(x,t)$ after $N = 100$ steps of the lazy three state quantum walk with $\theta = \pi/2$, shown for several decoherence rates γ. Solid lines show the solution of the corresponding derived continuum Lindblad equation with spatial dephasing. The initial state is a spatially delocalised Gaussian wavepacket of width $\sigma_0 = 5$ with a Fourier-symmetric coin superposition $(1, \omega, \omega^2)/\sqrt{3}$, where $\omega = e^{2\pi i/3}$.

the range of gamma between these values as the continuum equation is derived under the weak decoherence assumption. As the dephasing rate increases, the dynamics smoothly interpolate from a ballistic, interference dominated regime with pronounced side peaks to a broader, more diffusive profile concentrated around the origin. The continuum model accurately reproduces this overall transition, including the suppression of the ballistic fronts and the emergence of a central peak at larger γ values. Small quantitative differences are visible at low decoherence, particularly in the outer regions of the distribution, where lattice scale interference effects remain more pronounced in the discrete walk than in the continuum description. This slightly reduced agreement compared to the coin-dephasing case is expected, since spatial dephasing acts directly on nonlocal coherences between lattice sites, which are more sensitive to coarse graining in the continuum limit. Overall, the plot confirms that the continuum equation captures the dominant transport behaviour induced by spatial decoherence, while smoothing out fine interference features of the underlying discrete dynamics.

3.4 Discussion

I have derived the continuum limit of the one dimensional lazy open quantum walk for both coin and spatial decoherence models. The unitary part of the evolution generates advection and local mixing governed by the coin, while the dissipative part captures the distinct physical effects of each noise channel. In the continuum limit, the unitary structure remains identical and only the dephasing operator changes form, revealing how different microscopic decoherence mechanisms manifest at the macroscopic level.

The coin dephasing Lindblad term removes off diagonal elements in the coin subspace while leaving spatial coherences between lattice sites intact. Physically, this results in pure coin dephasing. Local internal superpositions decay, but population transport remains coherent until diffusion dominates at high γ values. This behaviour is in good qualitative agreement with early numerical studies of decoherence in the two state DTQW, which showed that coin noise suppresses interference without immediately destroying ballistic transport [López and Paz, 2003, Brun et al., 2003b, Love and Boghosian, 2005, Alberti et al., 2014]. By contrast, spatial decoherence mechanisms were found to rapidly suppress long-range coherence and induce diffusive behaviour [Romanelli et al., 2005], a qualitative distinction that is also observed here. Crucially, the numerical plots demonstrate good agreement between the DTQW and the derived continuum dynamics across the explored parameter regimes. This agreement confirms that the DTQW provides a faithful discretisation of the target continuum quantum field theory. The continuum equation represents the macroscopic model of interest, while the underlying discrete dynamics form a controlled numerical scheme for its simulation. Establishing this correspondence justifies the use of discrete quantum walks as simulators of continuum quantum transport equations and motivates the systematic derivation of the continuum limit [di Molfetta and Debbasch, 2012]. Spatial dephasing, by contrast, suppresses coherences between different lattice sites while leaving on-site populations unaffected. As a result, phase correlations decay exponentially in space, progressively eliminating long-range interference effects and driving the dynamics toward an incoherent, effectively classical transport regime. This behaviour aligns

with previous numerical observations of spatially dephased quantum walks Tude and de Oliveira [2022].

Together, these results show that coin and spatial dephasing impose qualitatively different transport characteristics. Coin dephasing preserves coherent advection with local phase damping, while spatial dephasing removes all spatial coherence and yields incoherent propagation that quickly reaches a classical, diffusive regime. The continuum framework developed here thus bridges microscopic Lindblad dynamics with the macroscopic behaviour observed in open quantum walk simulations. I emphasise that the continuum equations derived in this work are not intended to be solved analytically in general. Instead, they play a dual role. Firstly, as a validation tool confirming that the discrete quantum walk reproduces the correct macroscopic behaviour, and second, as a target description guiding the design of quantum walk based simulation schemes. This perspective shifts the focus from analytical solvability to constructive simulation. In particular, even in cases where the resulting PDEs are analytically intractable, the associated discrete quantum walk provides a natural numerical method for exploring their dynamics. The framework developed here therefore establishes a foundation for using quantum walks as flexible tools for PDE simulation, which will be extended in future work to more complex models and interactions.

The analytical structure of these limits may also be useful for future hydrodynamic applications. The recent work of Au-Yeung et al. [2025] illustrates how quantum walks can emulate nonlinear fluid dynamics by reconstructing hydrodynamic quantities from the collective behaviour of quantum walkers in a smoothed particle framework. This provides a clear example of how quantum walk dynamics can generate effective nonlinearities at the macroscopic level. While the present model captures only advective motion with local phase damping, it establishes the basic ingredients, unitary transport combined with controlled decoherence, upon which genuinely hydrodynamic behaviour could be built. Additional components would be required to produce viscosity, diffusion, or emergent hydrodynamic equations. Nevertheless, the lazy open quantum walk limit explored here provides a clean analytical foundation for such extensions and could form the basis for future quantum fluid algorithms. Assuming continued progress

along the UK's national quantum roadmap, such methods could be deployable on fault-tolerant quantum computers by around 2030 [UK Department for Science, Innovation and Technology, 2025, 2023]. Recent work by Foulds and Kendon [2025] has also presented gate model circuit implementations of the lazy discrete time quantum walk, suggesting compatibility with Rydberg atom architectures that naturally support multiqubit interactions. In order to introduce decoherence into the coin subsystem, the coin qubit(s) can be entangled with one or more ancillary qubits via a controlled-k-phase (CkZ) gate. The resulting degree of entanglement, tuneable through the state preparation of the ancilla qubit(s), directly controls the effective strength of decoherence acting on the coin. This procedure enlarges the overall Hilbert space and introduces additional junk components into the final state vector, but enables the extraction of the decohered dynamics through suitable post-processing of the ancilla degrees of freedom.

More broadly, the continuum limits derived here now provide a systematic route for connecting open quantum walks to macroscopic transport equations, and will serve as a useful starting point for applications in kinetic and hydrodynamic settings in the following chapters. The continuum limit obtained in the present work yields a linear quantum field theory featuring coherent advection together with simple dephasing induced dissipation. While this captures the leading transport behaviour, more elaborate quantum walk constructions, such as multistep schemes or different coins, are known to generate richer structures in the continuum limit, including dispersive terms [Jolly and Di Molfetta, 2023]. Such models may therefore offer a pathway towards more faithful quantum analogues of lattice kinetic theories. This chapter lays the foundation for the following work by showing how continuous spacetime limits can be taken for DTQW models with projective decoherence. The resulting master equations provide the mathematical and analytic framework used in the next chapter, where these expansion techniques are developed further in the context of LBM inspired quantum walk models.

Chapter 4

Decohered Quantum Walks as Kinetic Models for Diffusion

This chapter builds on the continuous spacetime limits derived in the previous chapter in order to develop quantum walk models for diffusive transport. The previous chapter introduced lazy three state quantum walks and derived their continuum limits in the presence of decoherence. Both the resulting limits and the method used to obtain them provide the starting point for the present analysis. Here, the same expansion idea is extended into a Chapman-Enskog style multiscale framework, so that the density matrix and time derivative are expanded over multiple scales. This allows the orders at which relaxation, coherent transport, and diffusive corrections appear to be identified systematically.

The focus is first placed on diffusion rather than full advection-diffusion, since a model should reproduce diffusive spreading reliably before more complex transport behaviour is considered. In particular, this chapter investigates the role of coin decoherence as a relaxation mechanism, the effect of adding a rest state through a lazy walk, and the possible use of plastic scaling to reorder the continuum hierarchy. Together, these tools are used to assess whether decohered DTQWs can provide a useful lattice Boltzmann inspired framework for simulating diffusive dynamics.

I first motivate the use of the rest state and coin decoherence, before introducing a general Chapman-Enskog style framework for DTQW models with decoherence and

outlining the hierarchy I would ideally like to obtain. I then apply this framework to two state quantum walk models and their lazy three state counterparts. In each case, I assess the models both analytically and numerically, showing that the addition of a rest state improves the diffusive behaviour and produces more useful large scale dynamics. Finally, I discuss what these models still lack and what further structures may be needed to achieve the desired hierarchy. This motivates the development of the model in the following chapter.

4.1 Motivation and Connection to classical LBM

In the classical LBM, the inclusion of a rest state is desirable for recovering the correct macroscopic fluid behaviour. In particular, the rest population plays an important role in ensuring the symmetry and isotropy of the discrete velocity set, preserving thermodynamic consistency, and allowing the model to recover PDEs such as the Navier Stokes equations in the multiscale continuum limit [Gebäck T, 2018]. One key reason for this is the requirement of isotropy. Physical fluids propagate equally in all directions, and in classical kinetic theory this symmetry is guaranteed through continuous velocity integrals. On a discrete lattice, however, isotropy must instead emerge from the structure of the discrete velocity set itself [Rubinstein and Luo, 2008]. The inclusion of a zero velocity rest population provides a central reference point for the lattice velocities, as seen in standard schemes such as D2Q9 and D3Q19, and helps cancel directional lattice artefacts that would otherwise appear.

This motivates the exploration of rest states in the DTQW setting when using quantum walks to simulate diffusive transport processes. In the previous chapter, the continuum limits of lazy DTQWs with decoherence were derived. These models provide the foundation for the present chapter, where those continuum limits are used to construct full CE expansions in order to determine which DTQW models are most suitable for simulating diffusion dynamics. In this framework, decoherence plays the role of an effective relaxation or collision mechanism, analogous to the collision operator in classical LBM, as discussed further in Section 4.3.2.

For the remainder of this work, I restrict attention to coin decoherence. This choice

is motivated by the analogy with the lattice Boltzmann method. In classical LBM, the streaming step transports the populations along the lattice velocities, while the collision step locally relaxes the populations towards equilibrium [Krüger et al., 2016]. To mirror this structure in the quantum walk setting, the shift/streaming part of the walk is kept unitary, and decoherence is applied only to the internal coin degree of freedom. Thus, the coherent unitary walk provides the analogue of streaming, while the coin dephasing channel provides an effective relaxation mechanism. There is also a physical motivation for focusing on coin decoherence rather than spatial decoherence. Previous studies have shown that coin dephasing preserves spatial coherence while suppressing only internal superpositions, making it more useful for modelling controlled decoherence in quantum transport and algorithmic settings [López and Paz, 2003, Love and Boghosian, 2005, Kendon and Tregenna, 2003]. Coin decoherence therefore allows dissipative and potentially diffusive behaviour to emerge while retaining more of the underlying quantum walk structure.

4.2 General Model

In this chapter, I build on the DTQW with decoherence model introduced in the previous chapter. The models considered here are again described by the evolution equation:

$$\hat{\boldsymbol{\rho}}(t+\Delta_t) = (1-\gamma)\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}(t)\hat{\mathbf{U}}^\dagger + \gamma\sum_j \mathbb{P}_j\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}(t)\hat{\mathbf{U}}^\dagger\mathbb{P}_j^\dagger. \tag{4.1}$$

Here, $\hat{\boldsymbol{\rho}}(t)$ is the density matrix, $\hat{\mathbf{U}}$ is the unitary quantum walk operator, γ is the decoherence strength, and the $\mathbb{P}_j$ are projectors onto the chosen decoherence basis. Again, this equation may be written in the Lindblad form, which is more convenient to deal with mathematically. The exact Lindblad CPTP map is given by:

$$\hat{\boldsymbol{\rho}}(t+\Delta_t) = \hat{\boldsymbol{\rho}}(t) - i[\hat{\mathbf{H}},\rho(t)] + \gamma\big(L_j\hat{\boldsymbol{\rho}}(t)L_j^\dagger - \frac{1}{2}\big\{L_j^\dagger L_j,\hat{\boldsymbol{\rho}}(t)\big\}\big). \tag{4.2}$$

Throughout this chapter, Eqn. (4.2) will be used as the common open walk model starting point, while different choices of $\hat{\mathbf{U}}$ will define the different quantum walk dynamics

considered. In the above equation, the Lindblad jump operators are

$$L_j = \mathbb{P}_j \hat{\mathbf{U}}. \tag{4.3}$$

Throughout this chapter, all of the two state quantum walk models have the coin decoherence projectors

$$\mathbb{P}_L = |L\rangle \langle L|, \qquad \mathbb{P}_R = |R\rangle \langle R|. \tag{4.4}$$

For the lazy three state models, the corresponding projectors are

$$\mathbb{P}_L = |L\rangle \langle L|, \qquad \mathbb{P}_S = |S\rangle \langle S|, \qquad \mathbb{P}_R = |R\rangle \langle R|. \tag{4.5}$$

These projectors define the basis in which coin coherences are dephased. The two state case is analogous to a D1Q2 velocity set with left and right moving components, while the lazy three state case is analogous to a D1Q3 velocity set with left, rest, and right components. In each model below, the decoherence basis is fixed in this way, while the unitary operator $\hat{\mathbf{U}}$ is changed to investigate different walk structures.

The next section outlines the general Chapman-Enskog style expansion used throughout the chapter. This separates the model independent multiscale procedure from the model specific algebra. Once the general framework has been established, it will be applied to each DTQW model by substituting the relevant unitary operator, scaling choice, and decoherence projectors. Then, the resulting hierarchy from the expansion can be analysed.

4.3 A General Quantum CE Style Framework

I now present a general Chapman-Enskog style expansion for a DTQW with projective decoherence. The same expansion can then be applied to several different quantum walks without repeating the full derivation each time. For each model, the only operators that need to be substituted are the corresponding Fourier space walk operator, the scaling of the coin and shift parameters, and the chosen decoherence projectors. This

provides a common framework for the rest of the chapter. Rather than writing out the full algebra for every model, I use the general expansion developed here to identify the order at which relaxation, coherent transport, and k^2 dependent diffusive terms appear. This makes it possible to compare different DTQW models directly and to determine which structures are most suitable for simulating diffusion.

The method used here is closely related to the continuum limit calculations in the previous chapter. There, the discrete walk was expanded in a single small parameter in order to identify the leading PDE. Here, I use the same expansion idea, but in a Chapman-Enskog style multiscale setting. The density matrix and time derivative are now expanded order by order, so that relaxation, coherent transport, and diffusive corrections can be tracked separately, analogously to the classical LBM Chapman-Enskog expansion.

Before attempting to model the full advection-diffusion equation, I first focus on the simpler diffusive limit. The aim of this chapter is therefore to identify the properties of walk structures which can produce diffusive spreading, and only then to consider how these models might be extended to include advective transport.

The algebra generated by these expansions can be lengthy, particularly for the lazy three state models. For this reason, the model specific calculations are implemented symbolically in code. Namely, the following mathematical process is implemented using the SymPy Python package [Meurer et al., 2017]. The code constructs the discrete open walk channel, expands the unitary operator, decoherence channel, density matrix, and time shifted left hand side in powers of the continuum parameter, and then extracts the equations order by order. The symbolic implementation therefore acts as a consistency check on the Chapman-Enskog hierarchy. It ensures that each term appearing in the written expansion is generated from the same discrete update rule, scaling assumptions, and decoherence projectors used in the model. In the sections below, I present the resulting order equations and use them to interpret the continuum behaviour, while the full symbolic algebra is handled computationally.

I begin with the discrete evolution equation, Eqn. (4.1), in the CPTP Lindblad form given by Eqn. (4.2). The shift and coin operators, and therefore the unitary operator,

of each specific model is transformed into Fourier representation. In the quantum walk models considered here, the jump operators are chosen to contain the unitary walk operator,

$$L_j = \mathbb{P}_j \tilde{\mathbf{U}}(k), \tag{4.6}$$

where $\mathbb{P}_j$ are projectors onto the chosen coin decoherence basis, given by Eqns (4.4) and (4.5), and $\tilde{\mathbf{U}}(k)$ is the Fourier space walk operator for the specific model.

Substituting Eqn. (4.6) into the dissipative part of Eqn. (4.2) gives

$$\sum_j L_j \hat{\boldsymbol{\rho}} L_j^\dagger = \Pi_c \left[\tilde{\mathbf{U}} \hat{\boldsymbol{\rho}} \tilde{\mathbf{U}}^\dagger \right], \tag{4.7}$$

where Π_c is the coin diagonal projection defined in Eqn. (3.33). Since the projectors resolve the identity in the coin space,

$$\sum_j L_j^\dagger L_j = \tilde{\mathbf{U}}^\dagger \left(\sum_j \mathbb{P}_j^\dagger \mathbb{P}_j \right) \tilde{\mathbf{U}} = \mathbb{I}, \tag{4.8}$$

the anticommutator term reduces to $-\hat{\boldsymbol{\rho}}$. The unitary operator is expanded in terms of ε. Therefore the one step open walk channel may be written as

$$\hat{\boldsymbol{\rho}}(t+\varepsilon) = \mathcal{L}_\varepsilon[\hat{\boldsymbol{\rho}}(t)], \tag{4.9}$$

with

$$\mathcal{L}_\varepsilon[\hat{\boldsymbol{\rho}}] = \tilde{\mathbf{U}}(k,\varepsilon)\hat{\boldsymbol{\rho}}\tilde{\mathbf{U}}^\dagger(k,\varepsilon) + \varepsilon\bar{\gamma}\left(\Pi_c \left[\tilde{\mathbf{U}}(k,\varepsilon)\hat{\boldsymbol{\rho}}\tilde{\mathbf{U}}^\dagger(k,\varepsilon) \right] - \hat{\boldsymbol{\rho}} \right), \tag{4.10}$$

where the first term describes the coherent unitary part of the evolution. In the continuum expansion, this term generates the Hamiltonian transport contributions. The second term describes the non-unitary coin dephasing channel and acts as the relaxation mechanism in the Chapman-Enskog hierarchy. Here I have used the weak decoherence scaling

$$\gamma = \varepsilon\bar{\gamma}. \tag{4.11}$$

This scaling is natural in the present chapter because decoherence is applied once per

timestep, so the dephasing strength is taken to scale with the physical timestep $\Delta_t = \varepsilon$.

The advantage of Eqn. (4.10) is that all model specific information is contained inside $\mathcal{L}_\varepsilon$. The choice of $\tilde{\mathbf{U}}(k,\varepsilon)$ determines the coherent walk dynamics, while the choice of projectors $\mathbb{P}_j$ determines the decoherence basis. Once these choices are made, the full open channel can be expanded in powers of ε.

I assume that the one step channel admits an expansion of the form

$$\mathcal{L}_\varepsilon = \mathcal{L}^{(0)} + \varepsilon\mathcal{L}^{(1)} + \varepsilon^2\mathcal{L}^{(2)} + \varepsilon^3\mathcal{L}^{(3)} + \mathcal{O}(\varepsilon^4). \tag{4.12}$$

The Chapman-Enskog expansion introduces multiple time scales,

$$\partial_t = \partial_{t_0} + \varepsilon\partial_{t_1} + \varepsilon^2\partial_{t_2} + \varepsilon^3\partial_{t_3} + \mathcal{O}(\varepsilon^4), \tag{4.13}$$

and expands the density matrix as

$$\hat{\boldsymbol{\rho}} = \hat{\boldsymbol{\rho}}_0 + \varepsilon\hat{\boldsymbol{\rho}}_1 + \varepsilon^2\hat{\boldsymbol{\rho}}_2 + \varepsilon^3\hat{\boldsymbol{\rho}}_3 + \mathcal{O}(\varepsilon^4). \tag{4.14}$$

The left hand side of the discrete update is expanded in the physical timestep as

$$\hat{\boldsymbol{\rho}}(t+\varepsilon) = \hat{\boldsymbol{\rho}} + \varepsilon\partial_t\hat{\boldsymbol{\rho}} + \frac{\varepsilon^2}{2}\partial_t^2\hat{\boldsymbol{\rho}} + \frac{\varepsilon^3}{6}\partial_t^3\hat{\boldsymbol{\rho}} + \mathcal{O}(\varepsilon^4). \tag{4.15}$$

Substituting Eqns. (4.12), (4.13), (4.14), and (4.15) into Eqn. (4.9), and matching equal powers of ε, gives a hierarchy of equations.

Before interpreting this hierarchy, the residual equation is divided by the timestep $\Delta_t = \varepsilon$. Therefore the raw $\mathcal{O}(1)$ coefficient corresponds to the scaled $\mathcal{O}(\varepsilon^{-1})$ equation, the raw $\mathcal{O}(\varepsilon)$ coefficient corresponds to the scaled $\mathcal{O}(1)$ equation, and so on. At raw order $\mathcal{O}(1)$, equivalently scaled order $\mathcal{O}(\varepsilon^{-1})$, one obtains

$$\hat{\boldsymbol{\rho}}_0 = \mathcal{L}^{(0)}[\hat{\boldsymbol{\rho}}_0], \tag{4.16}$$

or equivalently

$$\left(\mathcal{L}^{(0)} - \mathbb{I}\right)[\hat{\boldsymbol{\rho}}_0] = 0. \tag{4.17}$$

There is, however, an important distinction between the weakly decohered quantum walk limit used here and the standard classical LBM Chapman-Enskog expansion. In the classical BGK setting, the leading order kinetic equation usually imposes the local equilibrium condition $f^{(0)} = f^{\mathrm{eq}}$. Thus the collision operator already selects the slow equilibrium sector at leading order. In the weakly scaled quantum walk setting considered here, the situation is different. For a regular continuous time limit, the full one step channel must be close to the identity as $\varepsilon \to 0$. That is, the regularity condition is not only a condition on the unitary part of the walk, but on the complete open channel:

$$\mathcal{L}^{(0)} = \mathbb{I}. \tag{4.18}$$

Equivalently, for every admissible leading order state,

$$\mathcal{L}^{(0)}[\hat{\boldsymbol{\rho}}_0] = \hat{\boldsymbol{\rho}}_0. \tag{4.19}$$

This condition is required because the residual equation is divided by the timestep $\Delta_t = \varepsilon$. If the leading channel produced a finite change at order $\mathcal{O}(1)$, then the continuum generator would contain a term of size $\varepsilon^{-1}(\mathcal{L}^{(0)} - \mathbb{I})[\hat{\boldsymbol{\rho}}_0]$, which would diverge as $\varepsilon \to 0$. Therefore, in the weak decoherence continuum limit, the raw $\mathcal{O}(1)$ equation gives only the trivial identity condition

$$\hat{\boldsymbol{\rho}}_0 - \hat{\boldsymbol{\rho}}_0 = 0. \tag{4.20}$$

This condition is necessary for the continuum limit to be regular, but it does not by itself impose a local equilibrium condition on $\hat{\boldsymbol{\rho}}_0$.

The equilibrium selection must therefore occur at the first non-trivial order of the expansion. This is precisely why the weak decoherence scaling

$$\gamma = \varepsilon\bar{\gamma} \tag{4.21}$$

is used. With this scaling, the dephasing channel does not appear in $\mathcal{L}^{(0)}$, so the leading channel remains the identity. The dephasing term instead enters at the next order,

where it can act as a finite relaxation mechanism in the Chapman-Enskog hierarchy. If a strong dephasing scaling were chosen instead, with $\gamma = \bar{\gamma}$, then the leading channel would already contain a finite dephasing operation. Even if the leading unitary part satisfied $U^{(0)} = \mathbb{I}$, the full open channel would be

$$\mathcal{L}^{(0)}[\hat{\boldsymbol{\rho}}] = (1 - \bar{\gamma})\hat{\boldsymbol{\rho}} + \bar{\gamma}\Pi_c[\hat{\boldsymbol{\rho}}], \tag{4.22}$$

which is not equal to the identity channel unless $\hat{\boldsymbol{\rho}}$ is already coin diagonal. After division by the timestep, this would generate an $\mathcal{O}(\varepsilon^{-1})$ fast relaxation term. Such a strongly dephased limit would also cause the whole system to become classical very quickly, thus loosing any potential quantum advantage.

At raw order $\mathcal{O}(\varepsilon)$, equivalently scaled order $\mathcal{O}(1)$, the hierarchy gives

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_0 = \left(\mathcal{L}^{(0)} - \mathbb{I}\right)[\hat{\boldsymbol{\rho}}_1] + \mathcal{L}^{(1)}[\hat{\boldsymbol{\rho}}_0]. \tag{4.23}$$

At raw order $\mathcal{O}(\varepsilon^2)$, equivalently scaled order $\mathcal{O}(\varepsilon)$, one obtains

$$\partial_{t_1}\hat{\boldsymbol{\rho}}_0 + \partial_{t_0}\hat{\boldsymbol{\rho}}_1 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_0 = \left(\mathcal{L}^{(0)} - \mathbb{I}\right)[\hat{\boldsymbol{\rho}}_2] + \mathcal{L}^{(1)}[\hat{\boldsymbol{\rho}}_1] + \mathcal{L}^{(2)}[\hat{\boldsymbol{\rho}}_0]. \tag{4.24}$$

Higher orders are obtained in the same way.

Equations (4.16)-(4.24) define the general Chapman-Enskog hierarchy for the open DTQW channel. The important point is that this hierarchy does not automatically produce a closed scalar macroscopic PDE. It first produces a sequence of density matrix equations. These equations may contain diagonal populations, off diagonal coherences, and higher order corrections $\hat{\boldsymbol{\rho}}_1, \hat{\boldsymbol{\rho}}_2, \ldots$. To recover something closer to the classical LBM Chapman-Enskog structure, an additional closure step is needed. The higher order corrections must be solved for, projected out, or written in terms of the lower order slow variables. This is only possible in a clean way if the expansion first identifies a suitable local equilibrium sector, so that the non-equilibrium corrections can be treated as perturbations around it.

4.3.1 Closure Step

The Chapman-Enskog expansion produces a hierarchy of density matrix equations for

$$\hat{\boldsymbol{\rho}} = \rho^{(0)} + \varepsilon\rho^{(1)} + \varepsilon^2\rho^{(2)} + \cdots . \tag{4.25}$$

However, this hierarchy is not automatically a closed macroscopic equation. A further closure step is required if one wants to reduce the density matrix system to an equation for a smaller set of macroscopic variables, such as the total probability density

$$n(x,t) = \mathrm{Tr}_c\left[\rho^{(0)}(x,t)\right]. \tag{4.26}$$

Here Tr_c denotes the trace over the coin space. In this context, closure means that the higher order corrections, such as $\rho^{(1)}$ and $\rho^{(2)}$, can be solved for, projected out, or expressed in terms of the slow variable $n(x,t)$. This is the analogue of the classical LBM Chapman-Enskog procedure, where the non-equilibrium corrections to the distribution function are written in terms of gradients of the hydrodynamic variables. For the quantum walk hierarchy, the corresponding question is whether the coin populations and coherences can be eliminated so that the final equation involves only $n(x,t)$.

This distinction is important throughout the chapter. The Chapman-Enskog expansion determines the continuum density matrix hierarchy generated by the quantum walk, but closure asks the further question of whether that hierarchy can be reduced to a scalar macroscopic PDE. A model may therefore have a valid continuum limit while still failing to close to a scalar diffusion equation. In that case, the observed spreading is produced by a coupled quantum kinetic system rather than by an isolated macroscopic diffusion equation for $n(x,t)$.

Because the open quantum walk channel is trace preserving, the total probability is conserved automatically. This means that

$$\int n(x,t)\,dx \tag{4.27}$$

is constant in time, or equivalently that the zero Fourier mode $n(k=0,t)$ is conserved.

However, this does not mean that the local density $n(x,t)$ is constant. Probability can still be redistributed in space, so the non-zero Fourier modes may decay in time.

A closed diffusive limit would require the Chapman-Enskog hierarchy to reduce to an equation of the form

$$\partial_t n(x,t) = D\partial_x^2 n(x,t), \tag{4.28}$$

or, in Fourier space,

$$\partial_t n(k,t) = -Dk^2 n(k,t). \tag{4.29}$$

The k^2 term therefore does not violate probability conservation. It vanishes at $k = 0$, so the total probability is conserved, while the non-zero Fourier modes decay and the spatial profile spreads diffusively.

Thus, when k^2 terms appear in the diagonal population equations, the key question is not whether they break conservation. Rather, the question is whether, after tracing over the coin space and eliminating the remaining coin variables and coherences, those terms produce a closed contribution proportional to $-k^2 n(k,t)$. If this can be done, the model has an n closure to a scalar diffusion equation. If it cannot be done, then the model may still display diffusive spreading, but the continuum description remains a coupled density matrix kinetic hierarchy.

The model sections below, and in the following chapter, therefore use the Chapman-Enskog hierarchy in two steps. First, the expansion is used to identify which dephasing, transport, and k^2 terms appear at each order. Second, the structure of these terms is assessed to determine whether they can be closed into an effective macroscopic diffusion equation for $n(x,t)$. This closure step relies on having a well defined local equilibrium sector. In the classical LBM, the leading-order collision equation selects $f^{(0)} = f^{\mathrm{eq}}$, so the higher order non-equilibrium corrections can be computed as perturbations around this equilibrium. The same idea is required here. If the lowest non trivial order of the quantum walk does not first identify a slow equilibrium sector, for example a coin diagonal density matrix selected by dephasing, then the higher order corrections are not uniquely tied to the macroscopic density $n(x,t)$. The coherences and additional coin populations remain independent dynamical variables, making it difficult to eliminate

them and obtain a closed scalar equation. The exact form of this equilibrium is further discussed in the following subsection.

4.3.2 Equilibrium and Desired Hierarchy

The Chapman-Enskog expansion used in this chapter is intended to identify whether a decohered DTQW can produce an LBM inspired hierarchy of relaxation, transport, and diffusion. A key requirement is that the first non-trivial order should select a suitable local equilibrium sector before the coherent transport terms appear. For the coin dephasing channel used here, this would mean obtaining an equation of the form

$$\partial_{t_0}\rho^{(0)} = -\bar{\gamma}D_c[\rho^{(0)}]. \tag{4.30}$$

The dephasing operator is defined by

$$D_c[\boldsymbol{\rho}] = \boldsymbol{\rho} - \Pi_c[\boldsymbol{\rho}], \tag{4.31}$$

where Π_c keeps only the diagonal coin populations. Therefore D_c keeps only the off-diagonal coin coherences, and the dephasing term damps entries of the form

$$-\gamma\rho_{ij}, \qquad i \neq j. \tag{4.32}$$

The diagonal entries, corresponding to the coin populations, are not damped by this term. If this dephasing equation appears before the coherent transport and k^2 dependent terms, then the off diagonal coherences relax on the fast t_0 scale. After this relaxation, the leading order density matrix lies in the kernel of the dephasing operator,

$$D_c[\rho^{(0)}] = 0 \qquad \Longleftrightarrow \qquad \rho^{(0)} = \Pi_c[\rho^{(0)}]. \tag{4.33}$$

For a three state walk with coin basis $|L, S, R\rangle$, this gives

$$\rho^{(0)} = \begin{pmatrix} \rho_{00} & 0 & 0 \\ 0 & \rho_{11} & 0 \\ 0 & 0 & \rho_{22} \end{pmatrix}. \tag{4.34}$$

This diagonal form is not imposed by hand. It is selected dynamically by the lowest non-trivial dephasing terms in the continuum expansion. The diagonal populations then act as the slow variables, while the coherences are treated as non-equilibrium corrections.

The dephasing is fast only relative to the chosen continuum scaling. The claim is not that physical decoherence is instantaneous. Rather, the scaling is arranged so that the dephasing terms enter earlier in the hierarchy than the transport and diffusion terms. The coin coherences then relax on the fast kinetic scale, while the macroscopic density evolves on the slower hydrodynamic scale. This is the sense in which coin dephasing can define a local equilibrium sector. This does not mean that the full quantum walk is assumed to be classical, or that the density matrix is exactly diagonal at every order. The full expansion still has the form

$$\hat{\boldsymbol{\rho}} = \rho^{(0)} + \varepsilon\rho^{(1)} + \varepsilon^2\rho^{(2)} + \cdots . \tag{4.35}$$

Although $\rho^{(0)}$ may be coin diagonal, the higher order corrections $\rho^{(1)}, \rho^{(2)}, \ldots$ contain off-diagonal coherences generated by the unitary part of the walk. These coherences are not removed manually. Instead, they are damped by the dephasing channel and, in a successful Chapman-Enskog closure, become slaved to the slowly evolving diagonal populations. The leading sector is therefore classical looking, but the transport mechanism is still quantum.

This local equilibrium should not be confused with a final static probability distribution over the whole lattice. On an infinite line, or on an open lattice, a decohered quantum walk may continue to spread diffusively, often approaching a classical Gaussian profile in position. Thus the spatial distribution does not necessarily converge to

a single time independent equilibrium state. Instead, the equilibrium relevant to the Chapman-Enskog construction is local and coin based.

This interpretation is also consistent with the behaviour of the reduced coin state. The reduced coin state is obtained by tracing over position,

$$\rho_{\text{coin}} = \text{Tr}_x[\hat{\boldsymbol{\rho}}]. \tag{4.36}$$

For example, in an unbiased three state walk, one expects the reduced coin state to approach

$$\rho_{\text{coin}} \longrightarrow \begin{pmatrix} 1/3 & 0 & 0 \\ 0 & 1/3 & 0 \\ 0 & 0 & 1/3 \end{pmatrix}. \tag{4.37}$$

This state is diagonal and maximally mixed, so it is isotropic in the coin space [Romanelli et al., 2005]. However, this statement applies after tracing over position. The full density matrix can still contain position-coin correlations and higher order non-equilibrium corrections. Therefore, the use of a coin diagonal leading equilibrium is not an arbitrary classical assumption. It reflects the natural tendency of coin decoherence to suppress coherences and drive the coin sector towards a mixed local state, while the full open walk continues to evolve spatially.

Although total probability is conserved, this does not automatically give a closed scalar equation for $n(x,t)$. Closure requires the remaining coin variables and coherences to be either removed by projection or expressed in terms of the slow density. This is why the ordering of the hierarchy matters. If dephasing selects a clear equilibrium sector at the lowest non-trivial order, then the higher order equations can be used to solve for the non-equilibrium corrections and obtain an effective macroscopic equation. If instead the first non-trivial order contains dephasing and coherent unitary terms at the same time, then the situation is less clean. The leading equation no longer simply relaxes the density matrix towards the diagonal coin sector. Instead, the unitary terms continuously generate coherences while the dephasing term damps them at the same rate. The off-diagonal entries are therefore not purely fast variables relaxing around a

fixed local equilibrium; they are part of the leading order dynamics. In that case, the equilibrium sector is not clearly selected before transport begins, and the higher order corrections cannot be straightforwardly written only in terms of $n(x,t)$. The model may still show diffusive spreading numerically, but the continuum description remains a coupled density matrix hierarchy rather than a clean scalar diffusion equation.

The desired outcome of the quantum Chapman-Enskog style expansion is therefore an ordered hierarchy in which the different physical effects appear on separated scales. The raw leading order must be trivial, ensuring that the continuum limit exists. The first non-trivial order should then contain the dephasing terms, so that the diagonal equilibrium manifold is selected before transport effects dominate. The next order should contain the coherent transport terms, proportional to k and to the coin angle parameters, which generate the first non-equilibrium coherence corrections. The following order should then contain the feedback of these coherences onto the slow variables, together with the explicit k^2 terms responsible for diffusion. Schematically, the target ordering is

$$\text{raw } \mathcal{O}(1): \quad \text{identity condition for a valid continuum limit,} \tag{4.38}$$

$$\mathcal{O}(\varepsilon): \quad \text{dephasing and equilibrium selection,} \tag{4.39}$$

$$\mathcal{O}(\varepsilon^2): \quad \text{coherent transport,} \tag{4.40}$$

$$\mathcal{O}(\varepsilon^3): \quad \text{coherence feedback and } k^2 \text{ diffusion.} \tag{4.41}$$

This ordering is not guaranteed for an arbitrary quantum walk. Across the models considered in this chapter, the first non-trivial orders may contain different combinations of coin dephasing, coherent advective terms proportional to k, and higher order terms proportional to k^2. The aim is therefore to identify walk structures and scalings which arrange these terms in a useful order.

This construction is closely related to the role of equilibrium in a classical lattice Boltzmann method. In LBM, the collision step relaxes the distribution functions towards a chosen local equilibrium distribution. The Chapman-Enskog expansion then uses this relaxation to separate the slow hydrodynamic variables from the fast non-

equilibrium corrections. Here, the coin dephasing channel plays an analogous role. It acts as a quantum collision operator and damps the off-diagonal coin coherences and defines a local equilibrium sector of admissible density matrices. The diagonal coin populations then act as the leading slow variables. In this sense, the coin diagonal state can be interpreted as the natural local equilibrium, or local steady sector, selected by the open quantum walk. This is directly analogous to the role of equilibrium in LBM: the collision step does not make the whole fluid static, but relaxes the local velocity populations while the conserved quantities continue to evolve on larger space-time scales. Similarly, the dephasing channel does not make the full quantum walk static. It organises the leading hydrodynamic sector around the fixed points of the dissipative relaxation, while higher order terms may still contain coherences generated by the unitary dynamics.

It is worth noting that other decoherence choices are also possible. For example, one could dephase one coin state more strongly than another, or dephase only between a chosen subspace and its complement. Such choices would change the equilibrium sector selected by the dissipator and could allow a wider range of quantum equilibria. However, the resulting channel must remain completely positive and trace preserving, and asymmetric dephasing can break left/right symmetry. This may generate biased or advective contributions and can make the closure problem more complicated.

4.4 Grover Quantum Walk Model

The previous chapters introduced the $\mathcal{SU}(2)$ quantum walk and its lazy $\mathcal{SU}(3)$ extension. We begin with these models because they provide the simplest setting in which to test whether a quantum walk with coin decoherence can reproduce diffusive behaviour. In the classical LBM, the inclusion of a rest state plays an important role in constructing velocity sets with the correct symmetry and relaxation properties. Motivated by this, I ask whether adding a rest state to the quantum walk can improve the simulation of diffusive behaviour. The following subsections therefore revisit both the two state and three state lazy Grover walks, apply the Chapman-Enskog style expansion developed above, and compare the resulting continuum hierarchies. I then support this analysis

with numerical simulations, using the observed spreading and density profiles to assess whether either model provides a useful route to diffusive dynamics.

4.4.1 The two state Grover quantum walk model

The two state Grover quantum walk model was introduced in Eqn. (2.78). I begin with the coin decoherence projectors defined in Eqn. (4.4). The shift operator is given by Eqn. (2.65) and the coin as Eqn. (2.78), and the corresponding Fourier space unitary operator is given by Eqn. (2.84). For this model I use the standard continuum scaling

$$\Delta_x = \varepsilon, \qquad \Delta_t = \varepsilon, \qquad \theta = \varepsilon\bar{\theta}. \tag{4.42}$$

The Fourier space unitary operator, the coin decoherence projectors, and the scaled decoherence strength are then substituted into the discrete Lindblad equation, Eqn. (4.2). This gives the open quantum walk channel used as the starting point for the Chapman-Enskog expansion described above.

The above general framework for the Chapman-Enskog expansion is applied to this model. The first non-trivial scaled equation appears at order $\mathcal{O}(\varepsilon^0)$. The Chapman-Enskog hierarchy gives

$$\mathcal{O}(\varepsilon^{(0)}) = \partial_t \rho^{(0)} + i\left[k\sigma_z + \bar{\theta}\sigma_x, \rho^{(0)}\right] + \bar{\gamma} D_c[\rho^{(0)}] \tag{4.43}$$

This order contains coherent transport, coin mixing, and coin dephasing simultaneously. The term proportional to $k\sigma_z$ gives the left/right transport contribution, the term proportional to $\bar{\theta}\sigma_x$ mixes the two coin states, and the dephasing term damps the off-diagonal coin coherences. These two terms derive from the unitary Hamiltonian, and are give the Dirac Hamiltonian terms. The dephasing coin term, where D_c is given by Eqn. (3.34), gives the coin dephasing. Therefore, the model does not first relax to a purely diagonal equilibrium sector before transport appears.

At the next order, ε^1, the hierarchy becomes

$$\mathcal{O}(\varepsilon^1) = \partial_t \rho^{(1)} + \frac{1}{2}\partial_t^2 \rho^{(0)} \boldsymbol{\mathcal{U}}_2[\rho^{(0)}, \rho^{(1)}] \bar{\gamma}\left(\Pi_c\left[\boldsymbol{\mathcal{U}}_1[\rho^{(0)}, \rho^{(1)}]\right]\rho^{(1)}\right) \tag{4.44}$$

Here,

$$\boldsymbol{\mathcal{U}}_1[\rho^{(0)}, \rho^{(1)}] = \rho^{(1)} i \left[k\sigma_z + \bar{\theta}\sigma_x, \rho^{(0)} \right], \tag{4.45}$$

while $\boldsymbol{\mathcal{U}}_2$ contains the second order unitary corrections. These include the first explicit k^2 dependent terms. These terms arise partly from the second order expansion of the Fourier space shift,

$$e^{-i\varepsilon k\sigma_z} = \mathbb{I} - i\varepsilon k\sigma_z - \frac{\varepsilon^2 k^2}{2}\sigma_z^2 + \mathcal{O}(\varepsilon^3). \tag{4.46}$$

Since

$$\sigma_z^2 = \mathbb{I}, \tag{4.47}$$

the second order shift contribution is proportional to the identity in coin space. Thus the k^2 dependent terms do appear, but they do not distinguish between different velocity sectors beyond the left/right pair. Hence the contributions associated with diffusive spreading enter at this order, but they remain coupled to the coherence sector and to the first correction $\rho^{(1)}$, rather than producing a closed scalar diffusion equation for $n(x,t)$.

Unlike in the classical LBM Chapman-Enskog expansion, the leading order equation here does not define a clear local equilibrium density matrix depending only on the total probability density $n(x,t) = \mathrm{Tr}_c[\rho(x,t)]$. Therefore the higher order corrections, such as $\rho^{(1)}$, cannot be systematically expressed solely in terms of $n(x,t)$, and the hierarchy does not close to a scalar diffusion equation for n alone.

At order ε^2, the hierarchy involves powers such as k^3, $k^2\bar{\theta}$, $k\bar{\theta}^2$, and $\bar{\theta}^3$. It therefore represents higher order feedback between transport, coin mixing, coherence corrections, and dephasing. In the two state walk, the k^2 remain coupled to the coin populations, coherences, and the correction $\rho^{(1)}$. Thus the model gives a valid trace preserving quantum kinetic hierarchy, but it does not close to a scalar diffusion equation for the total density.

Ideally, for an LBM inspired hierarchy, the first non-trivial order would contain only the dephasing contribution, so that the diagonal coin sector is selected as the local equilibrium before transport appears. The next order would then contain the coherent Dirac Hamiltonian terms proportional to k and $\bar{\theta}$, generating the first non-equilibrium

corrections, while the higher orders would contain the k^2 terms that feed back onto the slow variables and produce diffusion. In the present two state model this ordering is not obtained: dephasing, transport, and coin mixing all enter at the same first non-trivial order, so the diffusive terms that appear later cannot be cleanly separated into a closed scalar equation for $n(x,t)$. I now extend this model to include a rest state, in order to determine whether the additional stationary component makes the hierarchy more supportive of the structures required to simulate diffusion.

4.4.2 The three state lazy Grover quantum walk model

The 3 state lazy grover walk with coin decoherence was explored in in Section 3.3.3 and its continuum limit was derived. This model will now be used in the CE framework above and the limit will be expanded to a multiscale expansion. We begin with the evolution equation, again given by Eqn. (4.2). The decoherence projectors are given by Eqn. (4.5). The standard, non plastic scalings are given by Eqn. (3.20), therefore the shift and coin and unitary evolution operator in Fourier space are given by Eqns (3.24),(3.23) and (3.22). For ease of notation, I drop the (3) superscript and from now on denote $J_z^{(3)} = J_z$. The decoherence projectors and unitary operators is inserted into the discrete evolution equation and the CE style expansion as described above is applied. The results are as follows.

At scaled order $\mathcal{O}(1)$, the first non-trivial order, the Chapman-Enskog type equation is

$$\mathcal{O}(1) = \partial_t \hat{\rho}^{(0)} + i[\hat{H}_1, \hat{\rho}^{(0)}] + \bar{\gamma} D_c[\rho^{(0)}] = 0. \tag{4.48}$$

The first order Hamiltonian was already determined in the previous chapter, given in Eqn. (3.26), therefore:

$$\hat{H}_1 = k\hat{J}_z + \bar{\theta}\hat{G}. \tag{4.49}$$

Now, Eqn. (4.48) may be written as

$$\partial_t \hat{\rho}^{(0)} + ik[\hat{J}_z, \hat{\rho}^{(0)}] + i\bar{\theta}[\hat{G}, \hat{\rho}^{(0)}] + \bar{\gamma} D_c[\rho^{(0)}] = 0. \tag{4.50}$$

This is essentially a three component Dirac-type Hamiltonian with coin dephasing. As in the two state case, the coherent unitary terms and the dephasing term appear at the same first non-trivial order. The dephasing therefore does not act first to select a diagonal local equilibrium sector before transport begins. Instead, relaxation, transport, and coin mixing are coupled immediately, so the hierarchy is not ordered in the desired LBM inspired form.

At order ε, the second order hierarchy becomes

$$\mathcal{O}(\varepsilon) = \partial_t \hat{\rho}^{(1)} + \frac{1}{2}\partial_t^2 \hat{\rho}^{(0)} - \boldsymbol{\mathcal{U}}_2[\hat{\rho}^{(0)}, \hat{\rho}^{(1)}] - \bar{\gamma}\left(\mathbf{\Pi}_c[\boldsymbol{\mathcal{U}}_1[\hat{\rho}^{(0)}, \hat{\rho}^{(1)}]] - \hat{\rho}^{(1)}\right) = 0. \tag{4.51}$$

Here

$$\boldsymbol{\mathcal{U}}_1[\hat{\rho}^{(0)}, \hat{\rho}^{(1)}] = \hat{\rho}^{(1)} - i[\hat{H}_1, \hat{\rho}^{(0)}], \tag{4.52}$$

and $\mathcal{U}_2$ contains the second order unitary corrections from

$$\hat{U}\hat{\rho}\hat{U}^\dagger. \tag{4.53}$$

These terms include contributions proportional to $k^2, k\bar{\theta}$ and $\bar{\theta}^2$. In particular, the k^2 dependent terms arise from the second order expansion of the shift operator and from the second order unitary product $\hat{U}\hat{\rho}\hat{U}^\dagger$. Since the lazy shift is generated by $\hat{J}_z = \mathrm{diag}(1, 0, -1)$, the Fourier space shift contains the expansion

$$e^{-i\varepsilon k \hat{J}_z} = \mathbb{I} - i\varepsilon k \hat{J}_z - \frac{\varepsilon^2 k^2}{2}\hat{J}_z^2 + \mathcal{O}(\varepsilon^3). \tag{4.54}$$

Thus the second order hierarchy naturally contains terms involving $\hat{J}_z^2$. For the lazy walk,

$$\hat{J}_z^2 = \mathrm{diag}(1, 0, 1), \tag{4.55}$$

so these second order terms distinguish the left and right moving states from the stationary rest state.

The standard lazy Grover walk therefore produces a Chapman-Enskog style hierarchy, but the leading order equation is a quantum kinetic equation rather than a

scalar diffusion equation. The rest state is included through the D1Q3 shift generator $\hat{J}_z = \mathrm{diag}(1, 0, -1)$, while the Grover coin couples the three velocity sectors. At order $\mathcal{O}(1)$, the equation contains coherent transport terms generated by $k\hat{J}_z$, coin mixing terms generated by $\bar{\theta}\hat{G}$, and weak dephasing of the off diagonal coin coherences. The trace projection gives conservation of the total density. At order $\mathcal{O}(\varepsilon)$, explicit k^2 dependent terms enter the density matrix hierarchy. These terms describe spreading corrections, but they do not appear as a closed scalar diffusion equation because the population dynamics remains coupled to the coherence sectors. Thus the normal lazy Grover walk gives a useful reference hierarchy. It conserves total density and contains both coherent transport and dephasing induced relaxation, while its diffusive behaviour remains embedded in the full matrix valued kinetic system.

4.4.3 Comparison of the two state and three state Grover walks

Analytically, the lazy three state walk has a structural advantage over the two state walk because it contains a genuine rest population. This difference first becomes important in the second order terms of the Chapman-Enskog hierarchy, where the expansion of the Fourier space shift produces factors involving k^2. In the two state case, the k^2 term is proportional to σ_z, and therefore

$$\sigma_z^2 = \mathbb{I}. \tag{4.56}$$

Both coin states therefore have the same squared velocity and contribute equally to the second order shift term. By contrast, in the lazy three state walk,

$$J_z = \mathrm{diag}(1, 0, -1), \qquad J_z^2 = \mathrm{diag}(1, 0, 1) \neq \mathbb{I}. \tag{4.57}$$

The important point is note that the k^2 dependent terms don't distinguish left moving and right moving directions, since both have the same squared velocity. Rather, they distinguish the moving states from the rest state. This is closer in spirit to the classical D1Q3 lattice Boltzmann velocity set, where the rest population contributes to the density but not to the second moment. The rest state therefore gives an additional

degree of freedom for tuning diffusive spreading and makes the lazy three state walk more closer in analogy to classical LBM than the two state walk.

However, this should only be taken as a structural indication and not as a proof of scalar diffusion closure. The standard lazy Grover walk still produces a coupled density matrix hierarchy, with coherent transport, coin mixing, and dephasing appearing together at the leading non-trivial order. Thus the rest state improves the diffusive structure of the model, but it does not by itself reproduce the clean LBM ordering in which collision first selects a local equilibrium and the higher order equations close to a scalar diffusion equation.

In both cases, the leading order channel does not select a non-trivial slow sector. In particular, the leading order equation does not impose

$$\hat{\boldsymbol{\rho}}_0 = \Pi_c[\hat{\boldsymbol{\rho}}_0], \tag{4.58}$$

and so it does not force $\hat{\boldsymbol{\rho}}_0$ to be coin diagonal. This means that the off-diagonal coherences are not removed at leading order and cannot yet be treated as fast variables slaved to the diagonal populations. Consequently, the lazy walk does not realise the desired LBM inspired ordering in which a local equilibrium sector is first selected, followed by higher order transport corrections. Instead, dephasing, coherent transport, and the first corrections to the density matrix enter at the same non-trivial order, leaving the continuum hierarchy as a coupled density matrix system rather than an immediately closed scalar diffusion equation.

The Grover comparison shows that adding a rest state improves the numerical diffusive behaviour. This is consistent with the LBM picture: in a D1Q3 model, the rest population contributes to the density without directly carrying flux, while the moving populations carry the left and right transport channels. Analytically, this distinction appears in the higher order spatial terms. In the two state walk, the k^2 contribution does not distinguish between the two moving states, whereas in the lazy three state walk the corresponding operator separates the moving sectors from the stationary sector. This gives the model a more useful second moment structure, closer to the role of the rest state in classical LBM diffusion.

However, the CE hierarchy also shows that the standard lazy Grover walk does not realise the desired LBM inspired ordering. The leading non-trivial order contains coherent transport, coin mixing, and dephasing together, rather than first selecting a clean local equilibrium sector. This motivates the search for a modified walk structure in which the lower order coherent terms cancel, allowing dephasing, coherent transport, and k^2 dependent spreading terms to appear in a more useful order.

Numerical Simulation

To compare how well the different DTQW models reproduce diffusion, I simulated both a two state Grover walk and a lazy three state Grover walk with coin dephasing, where the evolution in each case follows Eqn. (4.1). The aim is to test whether introducing a rest state improves the ability of the walk to reproduce a classical diffusive profile.

Both simulations were performed on a line of length $N = 101$ for $T = 150$ timesteps, with dephasing strength $\gamma = 0.2$ and initial Gaussian width $\sigma_0 = 10$. Periodic boundary conditions were used in all of the DTQW model simulations in this work. In each case, a matched classical D1Q2/3 LBM profile was constructed using the fitting procedure described in Appendix A.4. This makes the comparison as fair as possible. The dephasing strength, lattice size, simulation time, and initial spatial width are all fixed, so the main structural difference is the presence of the additional rest state in the lazy three state walk.

For the lazy three state walk, the initial coin state is taken to be

$$|\chi\rangle = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ i \\ 1 \end{pmatrix}, \tag{4.59}$$

written in the basis $(|L\rangle, |S\rangle, |R\rangle)$. For the two state walk, an unbiased initial coin state is used,

$$|\chi_2\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \tag{4.60}$$

in the basis $(|L\rangle, |R\rangle)$.

Figure 4.1 shows the resulting probability densities, variances and mean positions. The probability density of the three state walk is visibly closer to the fitted LBM curve than that of the two state walk. This is also reflected quantitatively in the final density error, measured using the root mean square difference between the DTQW and LBM density profiles at the final timestep. For the three state walk, this error is 8.11×10^{-4}, whereas for the two state walk it is 1.773×10^{-3}. Thus the three state model gives roughly a factor of two improvement in the final density match.

The clearest difference, however, appears in the variance plots. The variance of the two state walk is still noticeably curved and does not follow the fitted LBM variance particularly well, indicating that the spreading remains less diffusive. By contrast, the variance of the lazy three state walk is much closer to the LBM result and is strikingly more linear over the simulated time window. This is especially important, because approximately linear variance growth is one of the main signatures of diffusive behaviour. The mean position of both walks remain extremely small, of order 10^{-15}, so neither model exhibits any significant drift. These fluctuations are likely due to the numerical errors, rather than anything significant in the models.

Taken together, these results show that the lazy three state walk provides a significantly better setting for diffusion than the corresponding two state walk in the numerical setting. The dephasing strength is the same in both simulations, and the overall setup is otherwise matched, yet the three state walk gives both a better density fit and variance match. This suggests that the additional rest state plays a useful dynamical role. Although the two and three state Grover coins are closely related, the extra stay channel allows the walk to spread in a more diffusion-like manner. In this sense, the rest state appears to support diffusive transport and makes the lazy three state walk a more natural candidate for lattice Boltzmann inspired quantum diffusion models.

The previous section shows that the two state and three state walks contain several useful ingredients for diffusion. The continuum hierarchies include coin dephasing, coherent transport terms, coin mixing terms, and higher order k^2 contributions associated with diffusive spreading. However, these ingredients do not appear in the desired order.

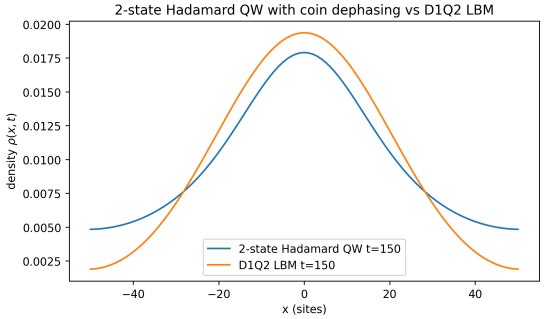

(a) Probability density: two state walk vs fitted LBM

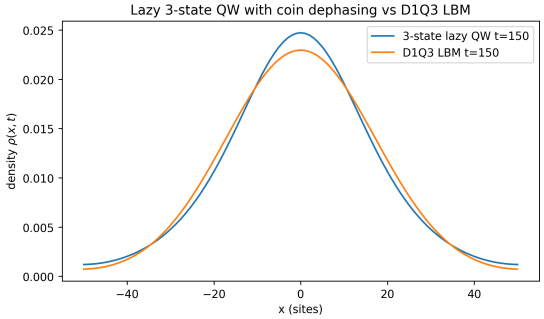

(b) Probability density: three state walk vs fitted LBM

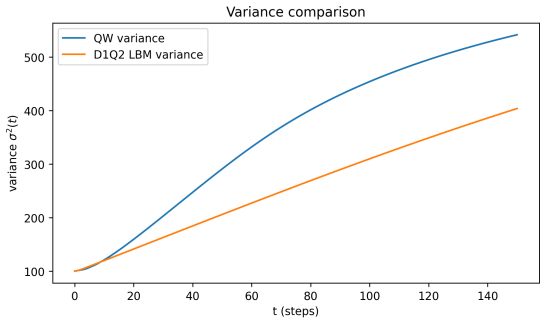

(c) Variance comparison for the two state walk

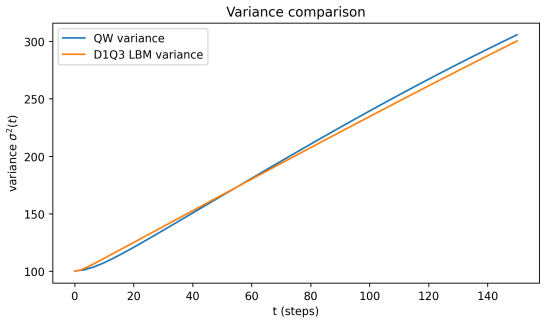

(d) Variance comparison for the three state walk

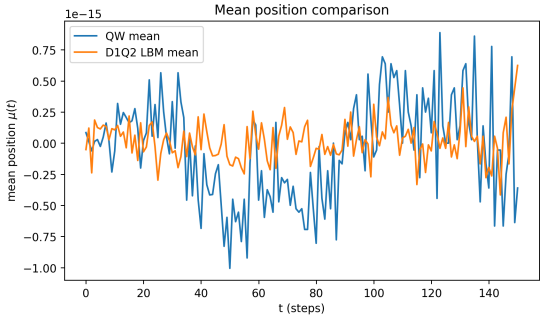

(e) Mean position for the two state walk

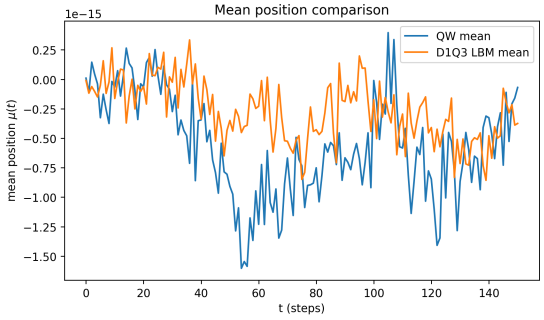

(f) Mean position for the three state walk

Figure 4.1: Comparison between the two state and lazy three state Grover walks with coin dephasing, together with their matched D1Q3 lattice Boltzmann diffusion profiles. Both models were simulated on a line of length $N = 101$ for $T = 150$ timesteps, with dephasing strength $\gamma = 0.2$ and initial Gaussian width $\sigma_0 = 10$.

In both cases, dephasing and coherent Hamiltonian evolution enter at the same first non-trivial order, so the diagonal equilibrium sector is not selected before transport begins. The higher order equations also become increasingly long and coupled, making it difficult to identify a clean closure to a scalar diffusion equation.

I therefore now ask whether the ordering and structure can be improved by modifying the walk structure. The aim is to construct a unitary walk in which the lowest order coherent transport terms are cancelled or delayed, while the dephasing contribution remains as the first non-trivial effect. If this can be achieved, the hierarchy may separate more naturally into equilibrium selection, coherent non-equilibrium corrections, and higher order diffusive feedback. In this way, a different walk structure may provide a cleaner route towards the LBM inspired hierarchy sought in this chapter.

4.5 Plastic Scaling and Regularity

The previous sections showed that the standard lazy Grover walk contains several useful diffusion friendly ingredients. In particular, the rest state gives the walk a D1Q3 structure, and the higher order Chapman-Enskog equations contain k^2 dependent terms associated with diffusive spreading. However, the ordering of the hierarchy is not ideal. In the standard non-plastic scaling, the first non-trivial equation contains coherent transport, coin mixing, and dephasing at the same order. Thus the dephasing does not first select a diagonal equilibrium sector before the unitary transport begins.

A natural question is therefore whether changing the continuum scaling can help reorder the hierarchy. In the previous general Chapman-Enskog framework, the timestep and lattice spacing were both ordered by ε. In the plastic scaling, instead one writes

$$\Delta_x = \epsilon^{1/2} = q, \qquad \Delta_t = \varepsilon = q^2, \qquad \gamma = \varepsilon\bar{\gamma} = q^2\bar{\gamma}. \tag{4.61}$$

The decoherence strength is still scaled with the physical timestep, because decoherence is applied once per timestep. What changes is the ordering of the spatial shift and coin angle terms. Since the shift depends on $\Delta_x = q$, the Fourier space factors expand as

$$e^{\pm ikq} = 1 \pm ikq - \frac{k^2q^2}{2} \mp \frac{ik^3q^3}{6} + \mathcal{O}(q^4), \tag{4.62}$$

with analogous expansions for any coin angles scaled in powers of q. Thus the right hand side of the discrete update may contain powers of q, while the left hand side

evolves over the longer timestep q^2.

This scaling is closely related to the usual diffusive scaling. Since

$$\Delta_x = q, \qquad \Delta_t = q^2, \tag{4.63}$$

one has

$$\Delta_t \propto \Delta_x^2. \tag{4.64}$$

Equivalently, space scales like the square root of time, which is the natural scaling for diffusion, where the width of a spreading profile grows like $\sqrt{t}$ [Schwarz, 2022]. For this reason, plastic scaling is a useful regime to explore when trying to obtain diffusive behaviour from a quantum walk. It changes the order at which spatial terms enter the Chapman-Enskog hierarchy, and potentially provides extra flexibility for arranging the dephasing, transport, and k^2 dependent terms into a more diffusion friendly structure.

However, plastic scaling also introduces an additional regularity condition. Since the residual equation is divided by the physical timestep q^2, any raw q^1 term would become a singular contribution of order q^{-1}. Therefore, if the walk operator expands as

$$\tilde{\mathbf{U}}(k, q) = \mathbb{I} + qU^{(1)} + q^2U^{(2)} + \mathcal{O}(q^3), \tag{4.65}$$

then a regular continuum limit on the timescale $\Delta_t = q^2$ requires

$$U^{(1)} = 0. \tag{4.66}$$

Otherwise, after division by the timestep, the term $qU^{(1)}$ produces

$$\frac{qU^{(1)}}{q^2} = q^{-1}U^{(1)}, \tag{4.67}$$

which diverges as $q \to 0$. Thus plastic scaling is not merely the substitution $q = \sqrt{\varepsilon}$ into the standard calculation. It changes the power counting and can reorder the hierarchy, but only for walk structures whose lower order terms are compatible with the q^2 timescale. In practice, this often requires a stroboscopic or otherwise symmetrised

construction. If the raw q terms do not cancel, the plastic continuum limit is singular and the walk is not a valid candidate for the desired diffusive hierarchy.

The standard lazy Grover walk cannot be taken directly into the plastic scaling regime as a way of reordering the Chapman-Enskog hierarchy. Consider the naive repeated construction

$$\tilde{\mathbf{U}} = \tilde{\mathbf{S}}\hat{\mathbf{C}}\tilde{\mathbf{S}}\hat{\mathbf{C}}, \tag{4.68}$$

where the lazy shift is generated by J_z and the Grover coin is generated by G. Note that now, the unitary operator must be considered over a two step unitary rather than over a single step. This is because the single step operator generally has a non-trivial zeroth order coin action and would not approach the identity as $q \to 0$. A regular continuum limit requires the full stroboscopic unitary. Using the plastic scaling, the first order expansion gives

$$\tilde{\mathbf{U}} = \mathbb{I}_3 - 2iq\left(kJ_z + \bar{\theta}G\right) + \mathcal{O}(q^2). \tag{4.69}$$

Hence

$$U^{(1)} = -2i\left(kJ_z + \bar{\theta}G\right) \neq 0. \tag{4.70}$$

The two substeps therefore add their first order generators rather than cancelling them. This still produces a non-zero raw q^1 term, which becomes an $\mathcal{O}(q^{-1})$ divergence after division by the plastic timestep q^2. This is consistent with the plasticity conditions discussed by Molfetta [2021], where the existence of a plastic limit requires the lower order terms in the expanded walk equation not to generate divergent contributions before the continuum time derivative appears. Thus, although the lazy Grover walk is a valid DTQW, it is not a regular strong plastic walk under the stroboscopic construction. To use a lazy walk in the plastic regime, the unitary must be modified so that the first order generator cancels. I therefore now consider a different stroboscopic model which is explicitly constructed to be regular under plastic scaling.

4.6 Twistless Quantum Walk Model

A useful model to explore is the previously discussed twisted quantum walk introduced by Jolly and Di Molfetta [2023]. This model is relevant because it is built from a plastic, stroboscopic structure in which the coin and shift operations are arranged so that the lowest order terms cancel. As a result, the walk has a regular plastic continuum limit, unlike the lazy Grover construction. In the unitary case, this continuum limit is particularly simple. When the twist and mass are removed, the model reduces to a clean effective transport equation with a single leading Hamiltonian term. This makes the twistless version a useful test case for the present work. It provides a simple, controlled walk structure in which to add coin dephasing and explore whether the resulting dynamics can support diffusive spreading. Rather than beginning with the full twisted and massive model, I first study this reduced case in order to isolate the basic cancellation mechanism and determine whether it gives a cleaner starting point for diffusive behaviour than the standard Grover walks.

The original twisted walk acts on a two dimensional coin space, with basis states denoted by $(|L\rangle, |R\rangle)$. The conditional shift is written as

$$\hat{\mathbf{S}} = \sum_x \left(|R\rangle \langle R| \otimes |x + \Delta_x\rangle \langle x| + |L\rangle \langle L| \otimes |x - \Delta_x\rangle \langle x| \right). \tag{4.71}$$

In Fourier space, as before, this becomes

$$\tilde{\mathbf{S}} = e^{i\Delta_x \sigma_z}. \tag{4.72}$$

The twisted walk is built from two substeps,

$$\hat{\mathbf{W}}_\alpha = \hat{\mathbf{S}}\hat{\mathbf{C}}_{\beta_1}, \qquad \hat{\mathbf{W}}_\beta = \hat{\mathbf{T}}^{-1}\hat{\mathbf{S}}_x\hat{\mathbf{T}}\hat{\mathbf{C}}_{\beta_2}, \tag{4.73}$$

where the twist is a rotation in coin space,

$$\hat{\mathbf{T}} = R_y(\theta) = e^{-i\theta\sigma_y/2}. \tag{4.74}$$

The one period operator is

$$\hat{\mathbf{G}} = \hat{\mathbf{W}}_{\beta_1}\hat{\mathbf{W}}_{\beta_2}, \tag{4.75}$$

and the full stroboscopic evolution is defined by

$$\hat{\mathbf{U}}_T = \hat{\mathbf{M}}\hat{\mathbf{G}}^2, \qquad \hat{\mathbf{M}} = e^{i\mu\sigma_z}. \tag{4.76}$$

Thus one timestep of the stroboscopic walk is

$$|\psi(t+\Delta_t)\rangle = \hat{\mathbf{U}}_T\,|\psi(t)\rangle\,. \tag{4.77}$$

Following Jolly and Di Molfetta [2023], the coin operators may be written as

$$\hat{\mathbf{C}}_{\beta_1} = \begin{pmatrix} \sin\alpha & -i\cos\alpha \\ i\cos\alpha & -\sin\alpha \end{pmatrix} = \sigma_y R_x(2\alpha), \tag{4.78}$$

and

$$\hat{\mathbf{C}}_{\beta_2} = \begin{pmatrix} -\cos\alpha & -i\sin\alpha \\ -i\sin\alpha & -\cos\alpha \end{pmatrix} = -R_x(-2\alpha). \tag{4.79}$$

The plastic scaling is then chosen as

$$\Delta_t = \varepsilon, \qquad \Delta_x = \sqrt{\varepsilon} = q, \qquad \alpha = \sqrt{\varepsilon}\bar{\alpha} = q\bar{\alpha}, \qquad \mu = m\varepsilon. \tag{4.80}$$

With this scaling, it is not the single operator $\hat{\mathbf{G}}$ that expands around the identity, but the stroboscopic operator $\hat{\mathbf{G}}^2$. In Fourier space, the full walk operator admits an expansion of the form

$$\tilde{\mathbf{U}}_T(k) = \mathbb{I}_2 - i\varepsilon\hat{\mathbf{H}}_T(k) + \mathcal{O}(\varepsilon^2), \tag{4.81}$$

where the continuum Hamiltonian is

$$\hat{\mathbf{H}}_T(k) = -\frac{m}{2}\sigma_z + k\left(-2\bar{\alpha} + k\sin\theta\right)\sigma_y. \tag{4.82}$$

The twist angle θ controls the quadratic k^2 term [Jolly and Di Molfetta, 2023]. Thus, for non-zero twist, the continuum limit contains both a Dirac linear contribution and

a dispersive second order contribution.

In this work, I choose to set the twist and mass to zero for simplicity,

$$\theta = 0, \qquad \hat{\mathbf{T}} = \mathbb{I}_2. \tag{4.83}$$

and

$$\mu = 0, \qquad \hat{\mathbf{M}} = \mathbb{I}_2. \tag{4.84}$$

This choice removes the additional dispersive and mass terms, leaving the simplest stroboscopic structure to analyse. The aim here is not to study the full twisted model, but to test whether the underlying plastic walk structure can improve the diffusive hierarchy once coin dephasing is added. Since the focus of this chapter is diffusion, rather than advection or more general dispersive behaviour, the mass and twist are set to zero so that the resulting continuum equations contain the fewest possible terms. These parameters could be reintroduced in later work to explore advection, mass effects, or more general transport. The shifted substep then reduces to

$$\hat{\mathbf{W}}_\beta = \hat{\mathbf{S}}_x \hat{\mathbf{C}}_{\beta_2}, \tag{4.85}$$

so that

$$\hat{\mathbf{G}} = \hat{\mathbf{S}}_x \hat{\mathbf{C}}_{\beta_2} \hat{\mathbf{S}}_x \hat{\mathbf{C}}_{\beta_1}. \tag{4.86}$$

The full twistless stroboscopic walk is therefore

$$\hat{\mathbf{U}} = \hat{\mathbf{G}}^2 = \left(\hat{\mathbf{S}}_x \hat{\mathbf{C}}_{\beta_2} \hat{\mathbf{S}}_x \hat{\mathbf{C}}_{\beta_1} \right)^2. \tag{4.87}$$

Equivalently, in Fourier space,

$$\tilde{\mathbf{U}}(k) = \left(\tilde{\mathbf{S}}_x(k) \hat{\mathbf{C}}_{\beta_2} \tilde{\mathbf{S}}_x(k) \hat{\mathbf{C}}_{\beta_1} \right)^2. \tag{4.88}$$

This twistless model still expands around the identity because of the stroboscopic square,

$$\tilde{\mathbf{U}}(k) = \mathbb{I}_2 - i\varepsilon \hat{\mathbf{H}}_0(k) + \mathcal{O}(\varepsilon^2). \tag{4.89}$$

However, setting $\theta = 0$ removes the quadratic dispersive term. The effective Hamiltonian becomes

$$\hat{\mathbf{H}}_0(k) = -2\bar{\alpha}k\,\sigma_y. \tag{4.90}$$

Thus the twistless model recovers the usual plastic Dirac type limit, rather than the generalised Dirac equation with dispersion. It remains useful here because it provides a clean two state stroboscopic walk whose leading continuum dynamics is linear in k. Decoherence can then be added to this twistless walk in order to investigate how relaxation modifies the spreading and whether diffusive behaviour can emerge from the open quantum walk dynamics.

The twisted construction is useful because the two coin operations are designed to partially undo one another over the stroboscopic step. This cancellation suppresses the leading ballistic contribution, so the walk does not spread as aggressively as a standard coherent quantum walk. The resulting dynamics are therefore closer to the slower, smoother spreading expected in a diffusive regime. This makes the model an interesting candidate to combine with controlled decoherence, since the unitary part already reduces the strongest ballistic effects, while dephasing can then act as an additional relaxation mechanism driving the system towards diffusive behaviour.

4.6.1 The 2 state twistless model

I now apply the Chapman-Enskog framework to the twistless two state stroboscopic walk with decoherence. The open evolution is again given by the discrete Lindblad equation, Eqn. (4.2) with the two state coin decoherence projectors defined in Eqn. (4.4). The unitary operator is the twistless stroboscopic operator given by Eqn. (4.88). The plastic scaling is used,

$$\Delta_t = \varepsilon = q^2, \qquad \Delta_x = q, \qquad \alpha = q\bar{\alpha}, \qquad \gamma = q^2\bar{\gamma}. \tag{4.91}$$

Since the physical timestep is q^2, the residual equation is divided by q^2. Therefore the raw q^0 and q^1 equations become the scaled orders q^{-2} and q^{-1}. For a regular plastic continuum limit, both of these orders must vanish.

For this stroboscopic walk, the unitary expansion gives

$$\tilde{\mathbf{U}}(k) = I + q^2 U^{(2)} + q^3 U(3) + q^4 U^{(4)} + \mathcal{O}(q^5), \tag{4.92}$$

with

$$U^{(0)} = I, \qquad U^{(1)} = 0. \tag{4.93}$$

Thus the dangerous raw q^1 term cancels. This shows that the twistless stroboscopic walk is regular in the plastic scaling. The first non-trivial unitary contribution appears at raw order q^2:

$$\mathbf{U}^{(2)} = \begin{pmatrix} 0 & -4\bar{\alpha}k \\ 4\bar{\alpha}k & 0 \end{pmatrix}. \tag{4.94}$$

Since

$$\tilde{\mathbf{U}}(k) = I - iq^2 \hat{H}_0 + \mathcal{O}(q^3). \tag{4.95}$$

and the corresponding effective Hamiltonian is given by Eqn. (4.90). At scaled order $\mathcal{O}(1)$, the first non-trivial Chapman-Enskog equation is therefore

$$\mathcal{O}(1) = \partial_t \rho^{(0)} + i\left[4\bar{\alpha}k\sigma_y, \rho^{(0)}\right] + \bar{\gamma} D_c[\rho^{(0)}] = 0. \tag{4.96}$$

This order contains coherent transport and coin dephasing simultaneously. The Hamiltonian term proportional to $4\bar{\alpha}k\sigma_y$ couples the diagonal populations to the off-diagonal coherences, while the dephasing term damps the off-diagonal entries. Thus, although the stroboscopic construction has removed the singular plastic terms, the first non-trivial equation is still not a dephasing-only equilibrium equation. The diagonal coin sector is therefore not selected before transport begins.

This equation is much cleaner than the corresponding lazy Grover hierarchy. The stroboscopic construction has successfully removed the singular plastic terms, and the leading coherent dynamics is described by a single effective Hamiltonian proportional to $k\sigma_y$. In this sense, the plastic stroboscopic walk gives a more controlled continuum expansion than the naive lazy walk. However, it still does not produce the desired LBM inspired ordering. The Hamiltonian transport term and the coin dephasing term

appear at the same first non-trivial order. This is because the dephasing is weakly scaled as $\gamma = q^2\bar{\gamma}$, so it enters the hierarchy at the same scaled order as the first non zero unitary contribution. Therefore the first non trivial equation is not a dephasing-only equilibrium equation. The diagonal coin sector is not selected before transport begins; instead, coherent dynamics and dephasing act simultaneously. Thus plasticity improves the regularity and cleanliness of the expansion, but it does not by itself separate equilibrium selection from unitary transport.

At the next scaled order, $\mathcal{O}(q)$, the hierarchy gives

$$\mathcal{O}(q) = \partial_t \rho^{(1)} + i[\hat{H}_0, \rho^{(1)}] + i[\hat{H}_1, \rho^{(0)}] + \bar{\gamma} D_c[\rho^{(1)}] = 0, \tag{4.97}$$

where the leading Hamiltonian is again Eqn. (4.90) and the next Hamiltonian correction comes from the raw q^3 unitary term,

$$U^{(3)} = \begin{pmatrix} 0 & 8i\bar{\alpha}k^2 \\ 8i\bar{\alpha}k^2 & 0 \end{pmatrix}. \tag{4.98}$$

Since the unitary expansion has the form

$$\tilde{U}(k) = \mathbb{I} - iq^2\hat{H}_0 - iq^3\hat{H}_1 + \mathcal{O}(q^4), \tag{4.99}$$

this gives

$$\hat{H}_1 = -8\bar{\alpha}k^2\sigma_x. \tag{4.100}$$

Therefore the $\mathcal{O}(q)$ equation can be written explicitly as

$$\partial_t \rho^{(1)} + i[4\bar{\alpha}k\sigma_y, \rho^{(1)}] - [8i\bar{\alpha}k^2\sigma_x, \rho^{(0)}] + \bar{\gamma} D_c[\rho^{(1)}] = 0. \tag{4.101}$$

The first explicit k^2 contribution therefore appears at this order. This is encouraging from the point of view of diffusion, because k^2 terms are the Fourier space signature of second spatial derivatives. The structure is also much cleaner than in the standard Grover walks: the leading unitary contribution is a single $k\sigma_y$ term, and the next correction is a single $k^2\sigma_x$ term, rather than a mixture of many contributions. Thus

the twistless walk gives a very simple continuum structure to analyse. However, the hierarchy is still not quite ordered in the desired way. The leading $\mathcal{O}(1)$ equation already contains both the coherent Hamiltonian term $4\bar{\alpha}k\sigma_y$ and the dephasing term $\bar{\gamma}D_c[\rho^{(0)}]$. Therefore dephasing does not first select a diagonal local equilibrium sector before unitary transport appears. The k^2 term is clean, but it enters inside the coupled density matrix hierarchy and remains linked to the zeroth order population/coherence dynamics. Thus the model is structurally promising and much easier to interpret than the Grover walks, but plasticity alone still does not produce the fully separated relaxation/transport/diffusion hierarchy. Consequently, the model gives a regular plastic quantum kinetic hierarchy with promising k^2 structure, but it does not by itself close to a scalar diffusion equation for the total density.

If the dephasing contribution could be made to appear one order earlier than the coherent Hamiltonian term, then this model would provide a very clean hierarchy: first relaxation to the diagonal coin sector, then coherent non-equilibrium corrections, and finally higher order k^2 diffusive feedback.

Overall, the twistless two state stroboscopic walk is a useful improvement over the naive plastic walk. Its main advantage is that the stroboscopic square cancels the singular lower order terms and produces clean, regular CE equations. The leading hierarchy is simple enough to interpret and contains a plausible diffusion mechanism through the damping of coherences. Its limitation is that the model still lacks the desired equilibrium structure. Dephasing does not select a slow diagonal manifold before coherent transport appears. Therefore, the model is diffusion friendly, but not yet a full lattice Boltzmann inspired quantum kinetic closure. I now generalise this model to the D1Q3 case.

4.6.2 The 3 state twistless model

I now generalise the twistless stroboscopic model to a lazy three state quantum walk with the same basic structure. The aim is to keep the useful plastic cancellation of the two state twistless walk, while adding a rest state so that the model has a D1Q3 structure. The open evolution again follows Eqn. (4.1), with coin dephasing projectors

given by Eqn. (4.5). The coin basis is ordered as $|L, S, R\rangle$, where $|S\rangle$ denotes the stationary rest state. The lazy shift operator is the same as the shift used in the lazy grover model, given by Eqn. (3.24). The coin structure is chosen to mimic the two state twistless construction, but embedded into the $|L, S, R\rangle$ space. The two twistless coins are first embedded into the moving $|L, R\rangle$ subspace while leaving the rest state unchanged:

$$\hat{\mathbf{C}}^{(3)}_{\beta_1} = \begin{pmatrix} \sin\alpha & 0 & -i\cos\alpha \\ 0 & 1 & 0 \\ i\cos\alpha & 0 & -\sin\alpha \end{pmatrix}, \tag{4.102}$$

and

$$\hat{\mathbf{C}}^{(3)}_{\beta_2} = \begin{pmatrix} -\cos\alpha & 0 & -i\sin\alpha \\ 0 & 1 & 0 \\ -i\sin\alpha & 0 & -\cos\alpha \end{pmatrix}. \tag{4.103}$$

These coins preserve the cancelling two state structure between the moving states, while the rest state is initially left untouched. To make the model genuinely lazy, I then introduce small rotations which mix the rest state with the moving states. The first coin mixes the $|L\rangle$ and $|S\rangle$ sectors, while the second coin mixes the $|S\rangle$ and $|R\rangle$ sectors:

$$\hat{\mathbf{R}}_{LS}(\eta) = \begin{pmatrix} \cos\eta & -\sin\eta & 0 \\ \sin\eta & \cos\eta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{4.104}$$

and

$$\hat{\mathbf{R}}_{SR}(\eta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\eta & -\sin\eta \\ 0 & \sin\eta & \cos\eta \end{pmatrix}. \tag{4.105}$$

The full coins are therefore

$$\hat{\mathbf{C}}_1 = \hat{\mathbf{R}}_{LS}(\eta)\hat{\mathbf{C}}^{(3)}_{\beta_1}, \qquad \hat{\mathbf{C}}_2 = \hat{\mathbf{R}}_{SR}(\eta)\hat{\mathbf{C}}^{(3)}_{\beta_2}. \tag{4.106}$$

This construction allows the original twistless cancellation to remain concentrated in the moving left/right subspace, while the additional rotations control how much amplitude

is exchanged with the stationary rest state. The R_{LS} rotation couples the left moving and rest sectors before one part of the walk, while R_{SR} couples the rest and right moving sectors before the other. In this way, the rest state is not simply passive. It can temporarily store amplitude and feed it back into the moving sectors, giving a three state walk with a D1Q3 velocity structure. The purpose is to retain the useful cancelling structure of the twistless walk, but allow the rest population to participate in the dynamics. It is important to note that this is not the only way of including a rest state in these coins, and there may be a more suitable parameterisation that has yet to be explored.

One period of the walk is then

$$\mathbf{W} = \tilde{\mathbf{S}}(k)\hat{\mathbf{C}}_1\tilde{\mathbf{S}}(k)\hat{\mathbf{C}}_2, \tag{4.107}$$

and the full stroboscopic unitary is

$$\tilde{\mathbf{U}} = \mathbf{W}^2 = \left(\tilde{\mathbf{S}}(k)\hat{\mathbf{C}}_1\tilde{\mathbf{S}}(k)\hat{\mathbf{C}}_2\right)^2. \tag{4.108}$$

The use of the square is again essential. As in the two state twistless walk, the stroboscopic structure is designed so that the dangerous first order plastic terms cancel. I use the same plastic scaling as in Section 4.5, together with

$$\eta = q^2\bar{\eta}. \tag{4.109}$$

The parameter $\bar{\eta}$ control how strongly the rest state mixes with the moving states. This makes the coupling between the stationary state and the moving states weak on the microscopic lattice scale, but finite on the continuum time scale. Since the same rest mixing scale is introduced consistently in both coins, the rest state participates in the stroboscopic structure without producing an additional raw q order term. In this way, the rest state can be included without spoiling the plastic cancellation inherited from the two state twistless walk. Thus this model provides a lazy, plastic, stroboscopic extension of the twistless walk: it keeps the cancellation mechanism needed for a regular

plastic limit, while adding a rest channel that may support more diffusion friendly behaviour.

The parameters α and η control different couplings within the walk. While $\alpha = q\bar{\alpha}$ governs the primary dynamics between the left and right moving states, $\eta = q^2\bar{\eta}$ introduces weaker coupling between the moving states and the stationary rest state through rotations in both coins. Physically, this allows the rest state to temporarily store and redistribute amplitude without dominating the underlying moving state dynamics, potentially reducing ballistic spreading and encouraging diffusive behaviour. The weaker scaling of η also ensures that these additional rotations do not introduce lower order terms that would disrupt the twistless cancellation, while still contributing to the continuum dynamics.

The Chapman-Enskog style expansion may now be applied to the lazy three state twistless walk. The first non-trivial equation appears at raw order q^2, equivalently scaled order $\mathcal{O}(1)$. It can be written compactly as

$$\mathcal{O}(1) = \partial_t \rho^{(0)} + i[H^{(2)}_{\mathrm{tw},3}, \rho^{(0)}] + \bar{\gamma} D_c[\rho^{(0)}] = 0. \tag{4.110}$$

Here $H^{(2)}_{\mathrm{tw},3}$ is the first effective Hamiltonian generated by the lazy three state twistless unitary. Using the Gell Mann matrices, this Hamiltonian may be written as

$$H^{(2)}_{\mathrm{tw},3} = 4\bar{\alpha}k\lambda_5 + \bar{\eta}\left(\lambda_1 + \lambda_2 - \lambda_6 - \lambda_7\right). \tag{4.111}$$

This form makes the structure of the leading equation much clearer. The term $4\bar{\alpha}k\lambda_5$ gives the leading left and right transport contributions. It couples the left and right sectors directly. The terms proportional to $\bar{\eta}$ mix the rest state with the moving states. However, these additional $\bar{\eta}$ terms may induce extra advection in the resulting dynamics, unlike in the 2 state model. This is important to bear in mind when attempting to simulate pure diffusion and may be a detriment to the model. The final term in Eqn. (4.110), $\bar{\gamma} D_c[\rho^{(0)}]$, damps the off diagonal coherences. Therefore, the first non-trivial scaled equation is a three component quantum kinetic equation with transport,

rest state mixing, and coin dephasing:

$$\partial_t \rho^{(0)} + i\left[4\bar{\alpha}k\lambda_5 + \bar{\eta}\left(\lambda_1 + \lambda_2 - \lambda_6 - \lambda_7\right), \rho^{(0)}\right] + \bar{\gamma} D_c[\rho^{(0)}] = 0. \tag{4.112}$$

The plastic regularity conditions are satisfied, however, it is still not the ideal LBM inspired hierarchy. The first non-trivial order is not only a dephasing equation. Instead, dephasing appears at the same order as the coherent Hamiltonian terms. The diagonal local equilibrium sector is therefore not selected before transport and rest state mixing begin. Therefore this construction still does not achieve the desired hierarchy.

The Hamiltonian can also be expressed using spin-1 operators, defined in Eqn. (3.28), as:

$$\hat{\mathbf{H}}_{\text{phys}}(k) = kJ_z + \frac{\bar{\theta}}{3}\left(\sqrt{2}J_x + J_x^2 - J_y^2\right). \tag{4.113}$$

Here, J_z describes directional propagation, while the linear and quadratic spin terms describe coherent coupling between the coin states.

At the next scaled order, $\mathcal{O}(q)$, the hierarchy has the schematic form

$$\mathcal{O}(q) = \partial_t \rho^{(1)} + i[H^{(2)}_{\text{tw},3}, \rho^{(1)}] + i[\hat{H}^{(3)}_{\text{tw},3}, \rho^{(0)}] + \bar{\gamma} D_c[\rho^{(1)}] = 0. \tag{4.114}$$

Here, $\hat{H}^{(3)}_{\text{tw},3}$ contains the next coherent corrections from the stroboscopic unitary. These include explicit $k^2 J_z^2$ dependent terms. Such terms are important because k^2 is the Fourier space signature of second spatial derivatives and therefore of diffusive spreading. However, at this order the k^2 terms remain coupled to the full density matrix, including the coin populations, coherences, and the first correction $\rho^{(1)}$. They do not immediately give a closed scalar contribution of the form

$$\partial_t n(k,t) = -Dk^2 n(k,t). \tag{4.115}$$

The lazy three state twistless model therefore sits between the standard lazy Grover walk and the desired LBM inspired hierarchy. Nevertheless, the first non-trivial equation still contains both coherent Hamiltonian dynamics and dephasing. Thus the model contains useful ingredients for diffusion, including a rest state, a regular plastic scaling,

and higher order k^2 terms, but it still does not fully separate equilibrium selection, transport, and diffusion into the desired Chapman-Enskog ordering.

4.6.3 Comparison of the two state and three state twistless models

Mathematically, the two state and three state twistless models are both cleaner than the corresponding Grover walks. In the Grover cases, the leading Chapman-Enskog equations contain several coupled contributions from the shift, the coin, and the dephasing channel. This makes the hierarchy algebraically complicated, with transport and diffusion terms not cleanly separated. By contrast, the twistless stroboscopic construction cancels the dangerous lower order plastic terms and produces a much simpler effective Hamiltonian structure. For the two state twistless walk, the first non-trivial equation is given by Eq. (4.6.1). The leading Hamiltonian is a single $k\sigma_y$ term, while the next correction, given by Eq. (4.6.1), contains a clean $k^2\sigma_x$ contribution. This makes the hierarchy much easier to interpret than the Grover case. However, the two state twistless model still does not reproduce the lattice Boltzmann inspired Chapman-Enskog ordering. In the two state twistless walk, the coherent Hamiltonian term and the dephasing term appear at the same first non-trivial order. Therefore, the diagonal coin sector is not selected before transport begins. The model has a clean k^2 structure, but it remains a coupled density matrix hierarchy rather than a closed scalar diffusion equation.

The three state twistless model improves the velocity structure by introducing a rest state. This makes it more diffusion friendly in a structural sense, because it contains a D1Q3 velocity set with left moving, right moving, and stationary components. The rest state contributes to the total density without directly contributing to the lattice flux, giving the model an additional local population through which probability can be stored and redistributed. This is closer to the role of the rest population in classical D1Q3 lattice Boltzmann diffusion models, where the stationary population provides additional freedom in the equilibrium distribution and in the second moment. Nevertheless, the presence of a rest state does not by itself prove scalar diffusion closure; it only provides a more suitable kinetic structure from which such a closure might emerge.

There is also a trade off in how the rest state is included. To preserve the plastic cancellation, the rest mixing angle must be scaled as

$$\eta = q^2\bar{\eta}. \tag{4.116}$$

This makes the rest coupling weak on the microscopic lattice scale, but finite on the continuum time scale. It keeps the model regular, but it also means that the rest state enters as a slow continuum order coupling rather than as a strong microscopic lazy component. Furthermore, the particular rest mixing construction used here may introduce unwanted drift or advection into the simulations. This could be useful in an advection diffusion setting, but it is not ideal when the aim is to simulate pure diffusion. This does not mean that adding a rest state is a bad idea, rather, it suggests that this particular rest state coupling may not be the optimal one.

Compared with LBM, the essential difference is therefore still the ordering of the hierarchy. In the twistless quantum walks, the k and k^2 structures are present and are much cleaner than in the Grover walks. However, they appear inside a coupled density matrix evolution. The off-diagonal coherences are damped by dephasing, but they are also generated by the Hamiltonian at the same order. Therefore, they cannot simply be eliminated in the same way as non-equilibrium LBM populations without an additional closure argument.

Overall, the twistless models represent a clear mathematical improvement over the Grover models. The two state twistless walk gives a clean D1Q2 hierarchy with a simple leading $k\sigma_y$ Hamiltonian and a clear higher order k^2 correction. The three state twistless walk extends this to a D1Q3 setting by adding a rest state and rest mixing terms, making it structurally more diffusion friendly. Nevertheless, both models still fall short of the ideal LBM inspired ordering. Plasticity removes the singular lower order terms and makes the expansion much cleaner, but it does not reorder the hierarchy in the way intended: equilibrium selection, transport, and diffusion are still not separated into distinct Chapman-Enskog orders.

Numerical Simulation Comparison

To further assess the models and the role of laziness, I again compare the decohered twistless two state model and the corresponding lazy three state construction and classical LBM. Both simulations were performed on a line of length $N = 101$ for $T = 150$ timesteps, with dephasing strength $\gamma = 0.2$, coin angles $\alpha, \eta = 0.15$, and initial Gaussian width $\sigma_0 = 10$. In each case, a matched classical lattice Boltzmann reference profile was constructed using the fitting procedure described in Appendix A.4. Again, since the two state and three state walks have different underlying kinetic structures, the two state model is compared with a fitted D1Q2 LBM, whereas the lazy three state model is compared with a fitted D1Q3 LBM. The initial states for the two and three state models are given by Eqs. (4.60) and (4.59), respectively.

The numerical comparison, shown in Figure 4.2 again shows that the lazy model provides the better diffusive approximation. For the two state twistless walk, the final density error at $t = 150$ is 0.001236, and the relative variance error is 0.177605. For the three state twistless walk, these give 0.000507, and 0.004114 respectively. Thus the three state model gives both a noticeably closer final density profile and a more accurate variance growth. In particular, the variance error drops from around 1.8×10^{-1} in the two state case to around 4.1×10^{-3} in the three state case, which is a substantial improvement.

For the two state walk, the final density profile is reasonably close to the fitted LBM curve, but the agreement is less precise than in the lazy case. This is a significant improvement from the two state Grover model, where the variance is still somewhat ballistic. The variance curve departs somewhat from the fitted classical line, indicating that the spreading is not as completely diffusive. By contrast, the lazy three state walk gives a final density profile that is much closer to the fitted D1Q3 LBM profile on the scale of the plot, and its variance growth is very nearly linear and closely tracks the classical benchmark over the full time window.

However, the numerical results for the lazy three state twistless walk show a more noticeable drift in mean position than the corresponding two state twistless model. This is interesting because it is also reflected in the Chapman-Enskog expansion. In

the two state twistless walk, the first non-trivial order is controlled only by the k dependent coherent transport terms. In the three state construction, however, the same order contains both the k dependent transport terms and the additional $\bar{\eta}$ dependent rest coupling terms. These terms mix the rest state into the moving sectors and may therefore introduce an effective transport bias if the coupling is not exactly left/right symmetric. This provides a useful link between the analytical hierarchy and the simulations. The increased drift seen numerically is not just an unexplained artefact of the plots; it is consistent with the structure of the expansion, where the rest state mixing appears already at the leading transport order. In this sense, the Chapman-Enskog calculation is informative: it shows that the extra three state coupling changes the effective transport structure of the walk, rather than simply adding a passive rest population. This should not be interpreted as a failure of the D1Q3 velocity structure itself. The rest state remains useful, especially because it gives the model an LBM-like second moment structure through $J_z^2 = \mathrm{diag}(1, 0, 1)$. Rather, the drift is likely caused by the particular way the rest state is coupled into the two coin twistless construction. The rest rotations used here may break the left/right parity symmetry of the effective walk, thereby introducing an advective bias even when the target behaviour is pure diffusion. A more symmetric rest state coupling could therefore retain the useful D1Q3 structure while reducing the unwanted drift.

Overall, this comparison supports the same conclusion as in the Grover case. Even when the underlying two state and three state twistless walks are not microscopically identical, introducing laziness through an additional rest channel leads to a significantly better approximation of classical diffusion. In this sense, the rest state again appears to stabilise the spreading dynamics and to produce a kinetic structure that is more naturally matched to an LBM diffusion model.

4.7 Discussion

The aim of this chapter was to identify the structural ingredients required for a decohered DTQW to reproduce diffusive transport, motivated by the LBM. In LBM, streaming transports populations along discrete velocities, while collisions provide lo-

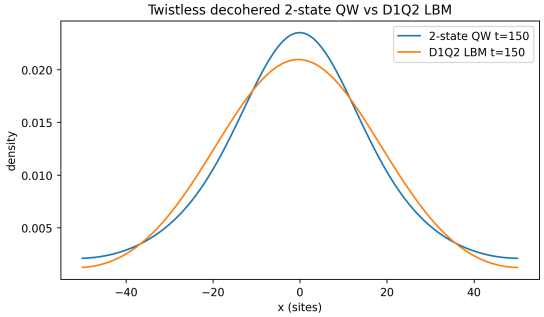

(a) Probability density: two state walk vs fitted D1Q2 LBM

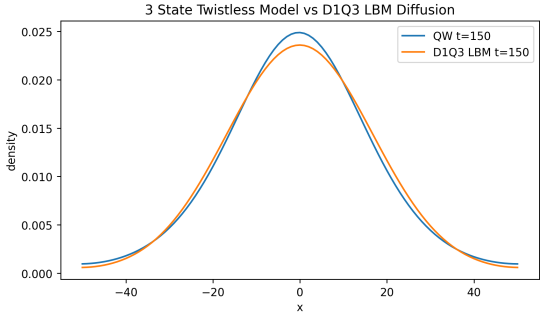

(b) Probability density: three state walk vs fitted D1Q3 LBM

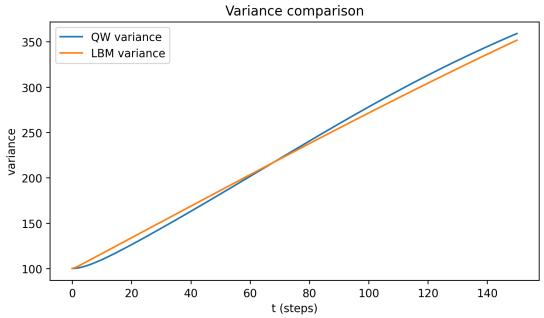

(c) Variance comparison for the two state walk

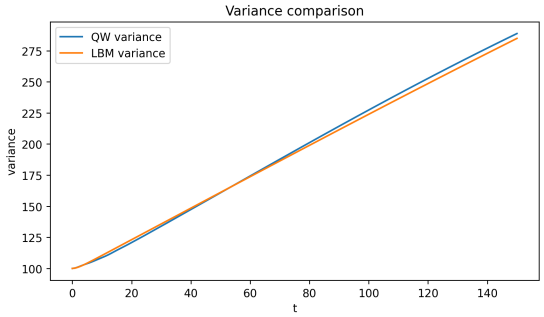

(d) Variance comparison for the three state walk

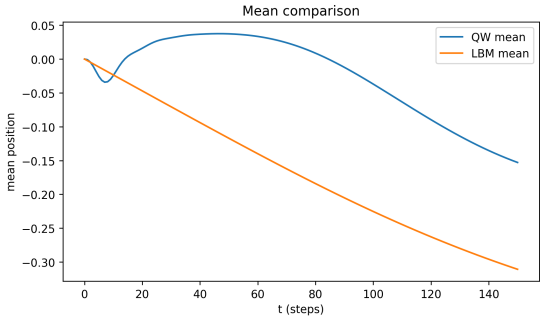

(e) Mean position for the two state walk

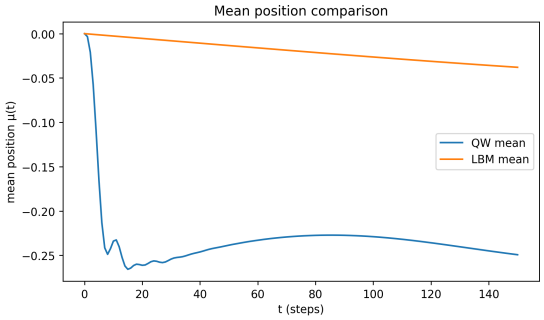

(f) Mean position for the three state walk

Figure 4.2: Comparison of decohered twistless two state and lazy three state quantum walks against fitted classical lattice Boltzmann models. The two state model is compared with a D1Q2 LBM, while the lazy three state model is compared with a D1Q3 LBM. In both cases the walks are evolved on a lattice of size $N = 101$ for $T = 150$ timesteps with dephasing strength $\gamma = 0.2$ and initial Gaussian width $\sigma_0 = 10$.

cal relaxation towards equilibrium. Analogously, coherent shift and coin operations govern quantum transport, while coin decoherence can act as a local relaxation mechanism. Rather than simply suppressing interference to produce more classical spreading,

the objective was to establish whether decoherence could play a kinetic role within a CE hierarchy.

For an LBM inspired diffusion mechanism, the desired ordering is for dephasing to first select the coin-diagonal sector as the local equilibrium, followed by coherent generation of off-diagonal non-equilibrium corrections. At higher order, these corrections should feed back into the slow population sector, producing k^2, or second spatial derivative, diffusive contributions. This parallels classical CE analysis, where diffusion emerges through the substitution of lower order non-equilibrium corrections into higher order moment equations. The models investigated here test whether this ordering can be achieved through coin decoherence, rest states, and plastic scaling.

The first important finding is that the addition of a rest state provides a more suitable kinetic structure for diffusion. The lazy three-state walk has a D1Q3 velocity set, allowing probability to be stored locally without directly contributing to directional flux. This also introduces a distinction between moving and stationary populations in the squared velocity operator, which is absent in the corresponding two state model. Numerically, the three state Grover walk exhibits variance growth closer to linearity and better agreement with fitted LBM diffusion profiles than its two state counterpart. Thus, the rest state improves the diffusive character of the walk, although it does not independently guarantee a closed macroscopic diffusion equation.

The main limitation of the Grover constructions is their CE ordering. In both the two and three state models, coherent transport, coin mixing, and dephasing appear at the same first non-trivial order. Consequently, dephasing cannot independently select the diagonal equilibrium sector, since coherences are generated and damped simultaneously. Although higher order equations contain k^2 terms associated with diffusive spreading, these remain coupled to the full density matrix hierarchy, preventing a straightforward scalar closure. This distinction is important: smooth probability distributions, approximately linear variance growth, and agreement with LBM profiles provide evidence of diffusion-like behaviour, but do not establish a closed macroscopic PDE. A stronger result requires a kinetic hierarchy with separated relaxation and transport contributions, followed by a closure for the conserved density.

The plastic twistless models were introduced to improve this ordering through stroboscopic cancellation. With the timestep scaling as q^2, any uncancelled $O(q)$ unitary contribution produces a singular term in the continuum limit. Plasticity therefore imposes a regularity condition requiring the lowest order coherent terms to cancel, demonstrating that scaling alone is insufficient without an appropriate microscopic walk structure. The twistless construction satisfies this condition, producing a considerably cleaner hierarchy than the Grover model.

In the two state twistless walk, the leading unitary contribution contains a single $k\sigma_y$ Hamiltonian term, while the next correction includes a clear $k^2\sigma_x$ contribution. Numerically, its variance growth agrees substantially better with the matched LBM simulation than the two-state Grover walk at the same decoherence strength. This supports the improved analytical structure, although a complete scalar diffusion closure is still not obtained.

Extending the twistless construction to three states introduces a D1Q3 velocity structure and additional rest state mixing. However, the CE expansion shows that these mixing terms enter at the same order as the leading coherent transport, modifying the effective dynamics and potentially breaking the left/right symmetry required for unbiased diffusion. This is reflected numerically in the more pronounced drift of the three state model compared with its two state counterpart. The agreement between the analytical hierarchy and numerical behaviour demonstrates that the CE expansion captures genuine transport properties of the discrete walk. While this drift is undesirable for pure diffusion, it could be useful for advection-diffusion, where directional transport is required. The results therefore suggest that future rest-state couplings should incorporate explicit parity symmetry, or isotropy in higher dimensions, to control unwanted advective contributions.

Despite their improved regularity, neither twistless model achieves the desired equilibrium-transport-diffusion ordering. With weak dephasing scaled as $\gamma = q^2\bar{\gamma}$, relaxation enters at the same first non-trivial scaled order as coherent transport. The leading equation therefore does not independently select the diagonal slow sector, and populations and coherences remain coupled. Nevertheless, the resulting density matrix equations are

valid quantum kinetic hierarchies containing coherent transport, coherence damping, and higher order spreading terms, which may be useful for studying open quantum transport beyond classical diffusion. What remains missing for an LBM-inspired model is a closure in which the coherences become slaved to the slow density variables, yielding a scalar diffusion or advection-diffusion equation.

Overall, this chapter identifies several important design principles. Coin decoherence provides a local relaxation mechanism, a rest state introduces a useful D1Q3 kinetic structure and improves numerical spreading, and plastic stroboscopic cancellation removes singular lower order contributions. However, these ingredients do not automatically produce the required CE ordering. The central remaining challenge is to suppress coherent transport at sufficiently low order that dephasing can first select the equilibrium sector, allowing higher order coherent corrections to generate diffusion.

Chapter 5

The Parity Echo Model

The previous chapter explored the use of coin decoherence and plasticity to obtain the correct hierarchy when performing a Chapman-Enskog style expansion to simulate classical diffusion. The aim was to find models that could be used to simulate diffusion by arranging the continuum limit so that relaxation, coherent transport, and diffusive terms appear at analogous orders with respect to classical LBM simulations. This chapter follows a similar direction, but instead of using plastic scaling to repair the hierarchy, it focuses on a single engineered model in which the unitary itself is designed to suppress the leading ballistic contribution.

The model introduced in this chapter is a parity echo lazy quantum walk. It is inspired by the cancellation structure of the twisted quantum walk model discussed in the previous chapter. In that model, the two coin rotations are arranged so that, when combined in the full step, their leading contributions partially undo one another. This cancellation suppresses the leading ballistic contribution, so the walk does not spread as aggressively as a standard coherent quantum walk. Here, I use the same idea. The two coin operations are chosen with opposite rotation signs and are connected by a parity swap, so that the leading directional part of the motion is cancelled by the second half of the step. The aim is to test whether this cancellation, built directly into the unitary dynamics, produces a cleaner Chapman-Enskog hierarchy for diffusion, and whether the analytic prediction is visible in direct numerical simulations.

5.1 The Parity Echo Model

This section introduces the parity echo model, which is later used as the main example for the diffusion analysis. The model is a lazy three state DTQW with a split step structure, so that one full unitary step is built from a shift-coin-shift-coin sequence. The coin space has basis states corresponding to left moving, rest, and right moving components, as before. The coins couple the left and right moving states to the rest state, so the rest state acts as a symmetric mediator between the two propagating directions. The key feature of the construction is the use of parity operators between opposite sign coin rotations. The parity operator swaps the left and right moving coin labels, while leaving the rest state fixed. This means that the second part of the unitary step partially reverses the directional structure generated by the first part. As a result, the dynamics are not simply a direct left/right propagation, but are mediated through coherent mixing with the rest state. This construction is useful for diffusion because a typical coherent DTQW has a strong ballistic component. In the parity echo model, the leading ballistic contribution is suppressed by the unitary sequence itself. The walk therefore has less variance than a typically ballistic DTQW, making it a better starting point for a diffusion construction. Instead of relying on plastic scaling to push terms into the desired orders, the parity echo model uses the structure of the coin sequence to produce a cleaner continuum hierarchy.

Physically, the walk first mixes the coin states in one direction, while the parity transformed part of the step reverses the leading directional bias. The result is that the lowest order ballistic drift cancels. What remains is a higher order transport mechanism mediated by the rest state and by the coherences generated during the walk. This cancellation pushes the unitary advection terms and the diffusive k^2 terms to higher orders, where they appear at the appropriate orders in the multiscale Chapman-Enskog expansion. In the previous chapter, plastic scaling was used as an attempt to force different components of the expansion to appear at the desired orders. Here, by contrast, the coin sequence itself is engineered so that plasticity is not required.

In the Chapman-Enskog analysis below, I focus first on the diffusive case. This

means that the walk is studied in a symmetric setting, without adding an explicit left/right bias to the initial coin state or to the coin parameters. This is mainly a simplifying choice as it keeps the hierarchy cleaner and allows the diffusive mechanism to be isolated before adding drift. In this setting, the aim is to understand if the dephasing term selects the diagonal slow sector, how the unitary walk generates coherence corrections, and how these corrections feed back into the populations to produce diffusive k^2 behaviour. This diffusion focused calculation is then tested using direct numerical simulations of the discrete walk with decoherence. These simulations are used to check whether the analytic predictions from the CE hierarchy are actually visible in the full discrete dynamics. The same parity echo model can also be used to produce advection-diffusion behaviour. Rather than including this bias in the CE derivation, I introduce it later at the level of the simulations by choosing an initial coin state with a larger right moving component. This generates a nonzero drift velocity while keeping the underlying parity echo walk structure the same. The advection-diffusion simulations are therefore discussed separately in Section 5.5, after the diffusion mechanism has been isolated.

5.1.1 The unitary parity echo model

For clarity and ease to understand how the parity echo model works, it is useful to understand the dynamics of the unitary walk first. Here, the unitary model is introduced, before including decoherence in subsequent sections. I consider a lazy DTQW on the line with Hilbert space $\mathcal{H}_{xc} = \mathcal{H}_x \otimes \mathcal{H}_c$. The position space $\mathcal{H}_x$ is spanned by $\{|x\rangle : x \in \Delta_x \mathbb{Z}\}$ with lattice spacing $\Delta_x > 0$. The coin Hilbert space is three dimensional, $\mathcal{H}_c = \text{span}\{|L\rangle, |S\rangle, |R\rangle\}$, corresponding to "left", "stay", and "right" moves. The global state evolves by a two step parity echo unitary,

$$|\psi(t+1)\rangle = \hat{\mathbf{U}}\,|\psi(t)\rangle = \hat{\mathbf{S}}\,(\mathbb{I} \otimes \hat{\mathbf{C}}_1)\,\hat{\mathbf{S}}\,(\mathbb{I} \otimes \hat{\mathbf{C}}_2)\,|\psi(t)\rangle, \tag{5.1}$$

where $\hat{\mathbf{C}}_1, \hat{\mathbf{C}}_2 \in \mathcal{U}(3)$ are coin operators, $\mathbb{I}$ is the identity on the position space, $\hat{\mathbf{S}}$ is the conditional shift, and $\hat{\mathbf{U}}$ is the full unitary evolution operator. The conditional shift

operator moves the particle left, right, or leaves it in place:

$$\hat{\mathbf{S}} = \sum_x \Big(|x+\Delta_x\rangle\langle x| \otimes |R\rangle\langle R| + |x\rangle\langle x| \otimes |S\rangle\langle S| + |x-\Delta_x\rangle\langle x| \otimes |L\rangle\langle L| \Big). \tag{5.2}$$

The coin operators are chosen to form a parity echo sequence. Let

$$\hat{\mathbf{P}} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \tag{5.3}$$

which swaps the left and right moving coin states while leaving the rest state fixed. I also define the symmetric, rest mixing generator

$$\hat{\mathbf{H}}_{\text{rest}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}. \tag{5.4}$$

The two coins are then given by

$$\hat{\mathbf{C}}_1 = \hat{\mathbf{P}} \exp(-i\theta \hat{\mathbf{H}}_{\text{rest}}), \qquad \hat{\mathbf{C}}_2 = \exp(+i\theta \hat{\mathbf{H}}_{\text{rest}}) \hat{\mathbf{P}}, \tag{5.5}$$

These coins may also be written in matrix form. The exponential of the rest mixing generator can be written exactly as

$$\exp(\mp i\theta \hat{H}_{\text{rest}}) = \mathbb{I} \mp i \sin\theta \hat{H}_{\text{rest}} + (\cos\theta - 1)\hat{H}^2_{\text{rest}}. \tag{5.6}$$

Therefore, writing $c_\theta = \cos\theta$ and $s_\theta = \sin\theta$, the two parity echo coins are

$$\hat{\mathbf{C}}_1 = \begin{pmatrix} \frac{c_\theta - 1}{2} & -\frac{is_\theta}{\sqrt{2}} & \frac{1 + c_\theta}{2} \\ -\frac{is_\theta}{\sqrt{2}} & c_\theta & -\frac{is_\theta}{\sqrt{2}} \\ \frac{1 + c_\theta}{2} & -\frac{is_\theta}{\sqrt{2}} & \frac{c_\theta - 1}{2} \end{pmatrix}, \tag{5.7}$$

and

$$\hat{\mathbf{C}}_2 = \begin{pmatrix} \frac{c_\theta - 1}{2} & \frac{is_\theta}{\sqrt{2}} & \frac{1 + c_\theta}{2} \\ \frac{is_\theta}{\sqrt{2}} & c_\theta & \frac{is_\theta}{\sqrt{2}} \\ \frac{1 + c_\theta}{2} & \frac{is_\theta}{\sqrt{2}} & \frac{c_\theta - 1}{2} \end{pmatrix}. \tag{5.8}$$

The two coins are adjoints of one another,

$$\hat{\mathbf{C}}_2 = \hat{\mathbf{C}}_1^\dagger, \tag{5.9}$$

which is the echo structure responsible for cancelling the lowest order coherent contribution in the stroboscopic walk. Now, the full unitary step is

$$\hat{\mathbf{U}} = \hat{\mathbf{S}}(\mathbb{I}_3 \otimes \hat{\mathbf{C}}_1)\hat{\mathbf{S}}(\mathbb{I}_3 \otimes \hat{\mathbf{C}}_2). \tag{5.10}$$

The parameter θ controls the strength of the mixing between the propagating states $|L\rangle, |R\rangle$ and the rest state $|S\rangle$. The operator $\hat{\mathbf{H}}_{\text{rest}}$ couples $|S\rangle$ symmetrically to $|L\rangle$ and $|R\rangle$, while the parity operator $\hat{\mathbf{P}}$ reverses the direction labels between the two substeps.

This parity echo structure is chosen so that the leading directional contributions from the two substeps partially cancel. In a standard lazy walk with a single coin, the shift generates first order spatial terms associated with the operator distinguishing left and right moving modes. Under a continuum expansion, these terms typically appear as odd powers of the wave number k, corresponding to directed ballistic transport. In the parity echo construction, the first substep is followed by a second substep in which the left and right coin labels are exchanged. Thus, a contribution which is left moving in one part of the update is mirrored by a corresponding right moving contribution in the other part. This is analogous to an echo sequence. The second half of the evolution reverses part of the directional structure generated by the first half.

The aim is not to remove all motion, but to suppress the strongest ballistic terms which dominate the continuum limit of a purely unitary quantum walk. By arranging the coins in this form, the two coin rotations are complex conjugate in sign and are separated by parity swaps. This means that terms which are odd under $L \leftrightarrow R$ are

naturally paired with terms of opposite sign. Consequently, the low order expansion of the full step operator can have reduced or cancelled half order and first order directional contributions, leaving higher order terms to control the effective macroscopic behaviour. This makes the walk a useful candidate for constructing a diffusive scaling limit, since diffusion arises from even spatial derivatives, such as k^2, rather than from dominant first order advection terms. DTQW models with coin dephasing often produce dephasing and unitary Hamiltonian terms at the same order in the continuum expansion. In that case, the dephasing term cannot be cleanly interpreted as an equilibrium constraint, since it is already coupled to transport at leading order. In the previous chapter, this ordering issue was attempted to be addressed using tuneable plastic scaling, however this was unsuccessful. The parity echo model provides a different solution. The coin sequence itself is engineered so that the leading ballistic contribution cancels. As a result, the first non trivial order may be governed by dephasing, which selects the diagonal slow manifold, while the unitary Hamiltonian and diffusive k^2 terms appear only at higher orders. This produces a cleaner Chapman-Enskog hierarchy without requiring plastic scaling.

Physically, this also makes the model better suited to diffusive simulation, because the parity echo suppresses the large ballistic spreading typical of coherent quantum walks. The probability distribution therefore remains more localised, and the observed broadening is dominated by the slower, decoherence mediated mechanism, rather than by rapid ballistic propagation. The parity echo walk can therefore be viewed as a modified lazy DTQW designed to reduce the intrinsic ballistic spreading of the standard unitary walk. The inclusion of the rest state provides the three velocities needed for a lazy or D1Q3 type transport structure, while the parity echo suppresses the leading left/right asymmetric contributions. This makes the model a starting point for studying how controlled decoherence or continuum scaling can produce effective diffusive dynamics from an underlying quantum walk.

In Fig. 5.1, I compare the unitary parity echo walk and the standard lazy Grover walk using the complex symmetric coin state coin state $\chi_{\text{cs}} = (1, 1, i)^T/\sqrt{3}$. This initial condition introduces a relative phase between the coin components due to the

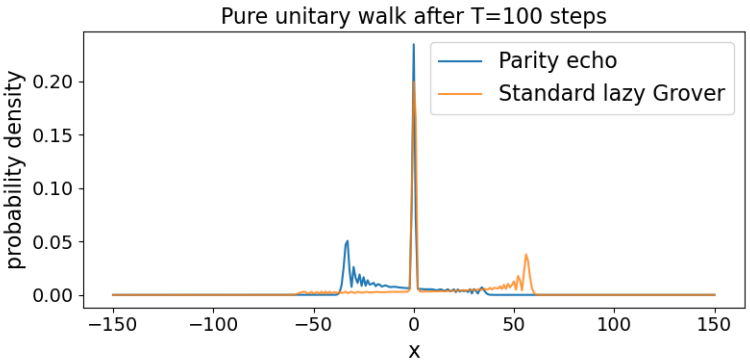

(a) Final probability density for the parity echo and standard lazy Grover walks.

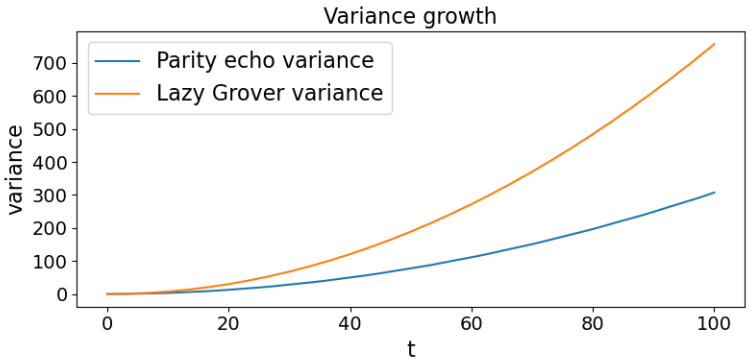

(b) Variance growth for the parity echo and standard lazy Grover walks.

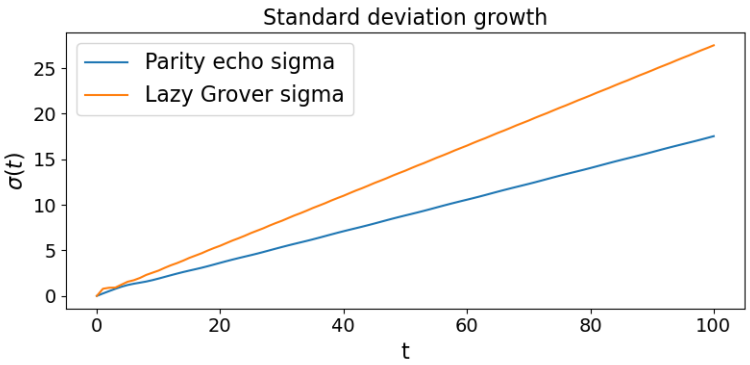

(c) Standard deviation growth for the parity echo and standard lazy Grover walks.

Figure 5.1: Comparison of the unitary parity echo lazy walk and the standard lazy Grover walk for the complex symmetric coin state coin state $\chi_{\text{cs}} = (1, 1, i)^T/\sqrt{3}$. The three panels show the final probability density, the variance growth, and the standard deviation growth.

imaginary part of the initial coin state, while keeping equal initial weights in the left and right moving states. It therefore tests the walks' response to initial coin coherence and demonstrates that both models produce a bias drift when an bias initial state is used. In Fig. 5.2, I repeat the comparison using the unbiased real coin state $\chi_{\mathrm{sym}} = (1, 1, 1)^T/\sqrt{3}$, which provides a baseline case where no explicit phase bias is inserted through the initial state. Comparing both cases is useful because it separates effects caused by the initial coin bias from effects intrinsic to the unitary update rule. In both the complex and real symmetric simulations, the parity echo model spreads less than the standard lazy Grover walk, as seen both in the probability density profiles and in the smaller variance and standard deviation. This reduced spreading is encouraging for the intended diffusion construction, since the parity echo walk has a weaker ballistic component and therefore provides a more controlled starting point before introducing decoherence or taking a continuum scaling limit.

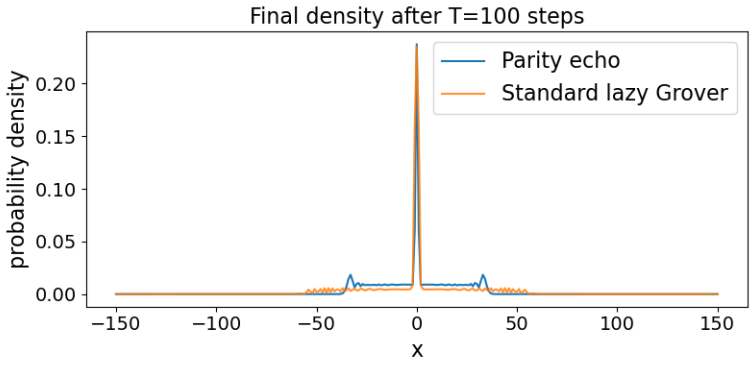

(a) Final probability density for the unbiased initial coin state $\chi_{\text{sym}} = (1,1,1)^T/\sqrt{3}$.

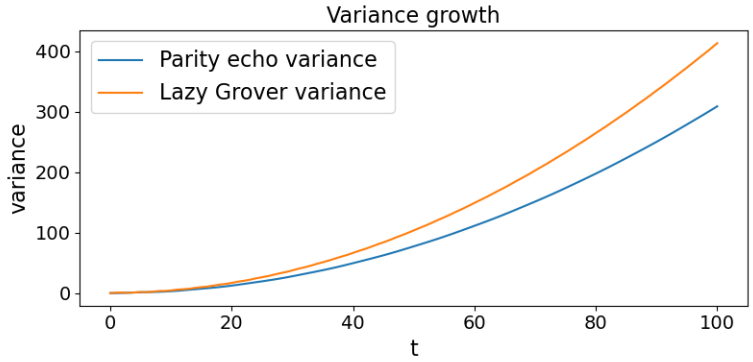

(b) Variance growth for the unbiased initial coin state $\chi_{\text{bias}} = (1,1,1)^T/\sqrt{3}$.

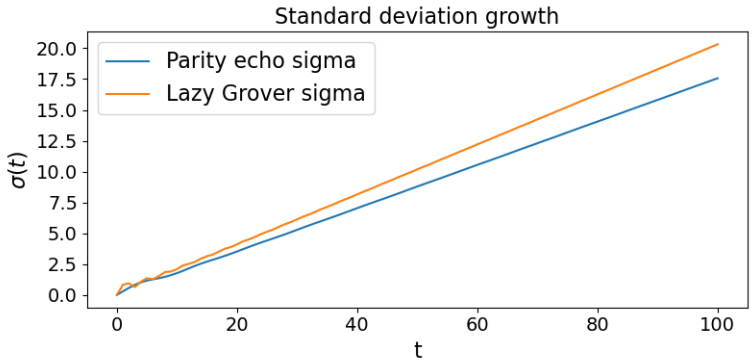

(c) Standard deviation growth for the parity echo and standard lazy Grover walks with unbiased initial state.

Figure 5.2: Comparison of the unitary parity echo lazy walk and the standard lazy Grover walk for the unbiased initial coin state $\chi_{\text{sym}} = (1,1,1)^T/\sqrt{3}$. The three panels show the final probability density, the variance growth, and the standard deviation growth.

5.2 The Chapman-Enskog Expansion

We now carry out the full CE style expansion of the parity echo model with decoherence to show how diffusive terms may arise. Directional bias could be introduced either through a biased initial coin state, leading to advection-diffusion behaviour. Here, however, I focus on the unbiased case in order to isolate the diffusive mechanism, keeping keeps the derivation clearer. Advection-diffusion will be addressed in later sections. Again, the model follows the CE expansion introduced in Section 4.3, however it will be detailed more explicitly for this model in order for the closure calculation to be clear.

5.2.1 Decoherence Map and Scalings

I begin with Eqn. (4.9), with the unitary operator being given by Eqn. (5.10). I now express both the shift and coin in a separable form in Fourier space, as following the previous chapters. The shift operator may again be written as:

$$\tilde{\mathbf{S}}(k) = \exp\big[-i\,\Delta_x k\,(\mathbb{I}_x \otimes J_z^{(3)})\big] \tag{5.11}$$

which is the standard D1Q3 shift operator used throughout this work. The scaling here is for the non plastic case, is given by

$$\Delta_t = \varepsilon \tag{5.12}$$

$$\Delta_x = \varepsilon \tag{5.13}$$

$$\theta = \varepsilon\bar{\theta} \tag{5.14}$$

$$\gamma = \varepsilon\bar{\gamma}. \tag{5.15}$$

The coin and shift operators are expanded as:

$$\tilde{\mathbf{S}}(\varepsilon, k) = \mathbb{I}_3 - i\varepsilon k J_z - \frac{\varepsilon^2 k^2}{2} J_z^2 + \frac{i\varepsilon^3 k^3}{6} J_z^3 + .. \tag{5.16}$$

The two coin operators are expanded as:

$$\hat{\mathbf{C}}_1(\varepsilon,k)=\hat{\mathbf{P}}(\mathbb{I}_3-i\varepsilon\bar{\theta}\hat{\mathbf{H}}_{rest}-\frac{\varepsilon^2\bar{\theta}^2}{2}\hat{\mathbf{H}}^2_{rest}+\frac{i\varepsilon^3\bar{\theta}^3}{6}\hat{\mathbf{H}}^3_{rest})+... \tag{5.17}$$

$$\hat{\mathbf{C}}_2(\varepsilon,k)=(\mathbb{I}_3+i\varepsilon\bar{\theta}\hat{\mathbf{H}}_{rest}-\frac{\varepsilon^2\bar{\theta}^2}{2}\hat{\mathbf{H}}^2_{rest}-\frac{i\varepsilon^3\bar{\theta}^3}{6}\hat{\mathbf{H}}^3_{rest})\hat{\mathbf{P}}+... \tag{5.18}$$

Now the structure of the DTQW with decoherence has been specified, and the operators have been expanded around ε in the non plastic scaling regime, I may move forward and show the CE hierarchy for this model.

Chapman-Enskog hierarchy for the parity echo model

I now apply the general Chapman-Enskog framework outlined in Section 4.3 to the parity echo model. Since this model gives the clearest closure of the density matrix hierarchy, I present the expansion in more detail here. This makes it possible to see explicitly how the dephasing term selects the diagonal slow sector, and how the higher order coherent corrections then generate the diffusive contribution.

The density matrix and time scales are expanded following Section 4.3. For the parity echo model, the unitary expansion takes the form

$$\tilde{\mathbf{U}}(k)=\mathbb{I}_3+\varepsilon^2U_2(k)+\varepsilon^3U_3(k)+\mathcal{O}(\varepsilon^4). \tag{5.19}$$

with

$$U_1(k)=0. \tag{5.20}$$

Thus the parity echo structure removes the first order coherent contribution. The second order unitary coefficient may be written as

$$U_2(k)=k\bar{\theta}A_2, \tag{5.21}$$

where

$$A_2=\begin{pmatrix}0 & -\frac{\sqrt{2}}{2} & 0\\ \frac{\sqrt{2}}{2} & 0 & -\frac{\sqrt{2}}{2}\\ 0 & \frac{\sqrt{2}}{2} & 0\end{pmatrix}. \tag{5.22}$$

The third order unitary coefficient separates into a part linear in k and a part quadratic in k:

$$U_3(k) = ik\bar{\theta}^2 A_3 + ik^2\bar{\theta}B_3, \tag{5.23}$$

with

$$A_3 = \begin{pmatrix} -\frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix}, \qquad B_3 = \begin{pmatrix} 0 & \frac{\sqrt{2}}{4} & 0 \\ \frac{\sqrt{2}}{4} & 0 & \frac{\sqrt{2}}{4} \\ 0 & \frac{\sqrt{2}}{4} & 0 \end{pmatrix}. \tag{5.24}$$

Here A_2, A_3, and B_3 are coin space matrix coefficients obtained from the expansion of the unitary operator. The corresponding superoperators acting on density matrices are defined by

$$\mathcal{A}_2[\hat{\boldsymbol{\rho}}] = A_2\hat{\boldsymbol{\rho}} + \hat{\boldsymbol{\rho}}A_2^\dagger, \tag{5.25}$$

$$\mathcal{A}_3[\hat{\boldsymbol{\rho}}] = A_3\hat{\boldsymbol{\rho}} + \hat{\boldsymbol{\rho}}A_3^\dagger, \tag{5.26}$$

and

$$\mathcal{B}_3[\hat{\boldsymbol{\rho}}] = B_3\hat{\boldsymbol{\rho}} + \hat{\boldsymbol{\rho}}B_3^\dagger. \tag{5.27}$$

The left hand side of the evolution equation is Taylor expanded as

$$\hat{\boldsymbol{\rho}}(t+\varepsilon) = \hat{\boldsymbol{\rho}} + \varepsilon\partial_t\hat{\boldsymbol{\rho}} + \frac{\varepsilon^2}{2}\partial_t^2\hat{\boldsymbol{\rho}} + \frac{\varepsilon^3}{6}\partial_t^3\hat{\boldsymbol{\rho}} + \mathcal{O}(\varepsilon^4). \tag{5.28}$$

Substituting the CE expansion of $\hat{\boldsymbol{\rho}}$ and ∂_t, the left hand side coefficients are

$$\mathrm{LHS}_0 = \hat{\boldsymbol{\rho}}_0, \tag{5.29}$$

$$\mathrm{LHS}_1 = \hat{\boldsymbol{\rho}}_1 + \partial_{t_0}\hat{\boldsymbol{\rho}}_0, \tag{5.30}$$

$$\mathrm{LHS}_2 = \hat{\boldsymbol{\rho}}_2 + \partial_{t_0}\hat{\boldsymbol{\rho}}_1 + \partial_{t_1}\hat{\boldsymbol{\rho}}_0 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_0, \tag{5.31}$$

and

$$\begin{aligned} \mathrm{LHS}_3 =& \hat{\boldsymbol{\rho}}_3 + \partial_{t_0}\hat{\boldsymbol{\rho}}_2 + \partial_{t_1}\hat{\boldsymbol{\rho}}_1 + \partial_{t_2}\hat{\boldsymbol{\rho}}_0 \\ &+ \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_1 + \partial_{t_0}\partial_{t_1}\hat{\boldsymbol{\rho}}_0 + \frac{1}{6}\partial_{t_0}^3\hat{\boldsymbol{\rho}}_0. \end{aligned} \tag{5.32}$$

The unitary part is expanded as

$$\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}\hat{\mathbf{U}}^\dagger = \mathcal{U}_0 + \varepsilon\mathcal{U}_1 + \varepsilon^2\mathcal{U}_2 + \varepsilon^3\mathcal{U}_3 + \mathcal{O}(\varepsilon^4). \tag{5.33}$$

Here $\mathcal{U}_n$ denotes the order ε^n contribution to the unitary sandwich $\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}\hat{\mathbf{U}}^\dagger$. It should not be confused with the unitary expansion coefficients U_n. The coefficients $\mathcal{U}_n$ are obtained by substituting both the CE expansion of $\hat{\boldsymbol{\rho}}$ and the unitary expansion of $\tilde{\mathbf{U}}(k)$ into the sandwich $\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}\hat{\mathbf{U}}^\dagger$, and then collecting powers of ε. Thus the matrices $U_2(k)$ and $U_3(k)$ from the expansion of the unitary above generate the superoperator terms appearing in $\mathcal{U}_2$ and $\mathcal{U}_3$. Since $\tilde{\mathbf{U}} = \mathbb{I}_3 + \varepsilon^2 U_2 + \varepsilon^3 U_3 + \mathcal{O}(\varepsilon^4)$, the first coefficients are

$$\mathcal{U}_0 = \hat{\boldsymbol{\rho}}_0, \tag{5.34}$$

$$\mathcal{U}_1 = \hat{\boldsymbol{\rho}}_1, \tag{5.35}$$

$$\mathcal{U}_2 = \hat{\boldsymbol{\rho}}_2 + k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0], \tag{5.36}$$

and

$$\mathcal{U}_3 = \hat{\boldsymbol{\rho}}_3 + k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_1] + ik\bar{\theta}^2\mathcal{A}_3[\hat{\boldsymbol{\rho}}_0] + ik^2\bar{\theta}\mathcal{B}_3[\hat{\boldsymbol{\rho}}_0]. \tag{5.37}$$

The dissipative contribution is

$$\varepsilon\bar{\gamma}\left(\Pi_c[\hat{\mathbf{U}}\hat{\boldsymbol{\rho}}\hat{\mathbf{U}}^\dagger] - \hat{\boldsymbol{\rho}}\right). \tag{5.38}$$

Therefore, for $n \geq 1$, the order ε^n dissipative coefficient is

$$\mathcal{D}_n = \bar{\gamma}\left(\Pi_c[\mathcal{U}_{n-1}] - \hat{\boldsymbol{\rho}}_{n-1}\right). \tag{5.39}$$

In particular,

$$\mathcal{D}_1 = \bar{\gamma}\left(\Pi_c[\hat{\boldsymbol{\rho}}_0] - \hat{\boldsymbol{\rho}}_0\right) = -\bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0], \tag{5.40}$$

$$\mathcal{D}_2 = \bar{\gamma}\left(\Pi_c[\hat{\boldsymbol{\rho}}_1] - \hat{\boldsymbol{\rho}}_1\right) = -\bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1], \tag{5.41}$$

and

$$\mathcal{D}_3 = \bar{\gamma}\left(\Pi_c[\hat{\boldsymbol{\rho}}_2 + k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]] - \hat{\boldsymbol{\rho}}_2\right). \tag{5.42}$$

Equivalently,

$$\mathcal{D}_3 = -\bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_2] + \bar{\gamma}\Pi_c[k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]]. \tag{5.43}$$

I now equate the left and right hand side coefficients at each order. At $\mathcal{O}(\varepsilon^0)$, both sides give $\hat{\boldsymbol{\rho}}_0$, so as expected

$$\mathrm{CE}_0: \qquad 0 = 0. \tag{5.44}$$

At $\mathcal{O}(\varepsilon)$, equating LHS_1 with $\mathcal{U}_1 + \mathcal{D}_1$ gives

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_0 + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0] = 0. \tag{5.45}$$

This is the first non trivial equation in the hierarchy. Since $U_1 = 0$, the leading order equation contains only the dephasing operator and no coherent transport term. Therefore the leading relaxation selects the diagonal slow sector. I impose

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_0 = 0, \tag{5.46}$$

because the slow sector is defined as the nullspace of the leading fast relaxation operator. After the t_0 scale dephasing transient has decayed, the leading order state no longer evolves on the fast time scale. This then gives

$$\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0] = 0, \tag{5.47}$$

Thus the diagonal leading order state is not imposed by hand. It is the nullspace selected by the leading dephasing dynamics after the fast t_0 scale transient has decayed. Hence

$$\hat{\boldsymbol{\rho}}_0 = \Pi_c[\hat{\boldsymbol{\rho}}_0] = \begin{pmatrix} \rho_{00} & 0 & 0 \\ 0 & \rho_{11} & 0 \\ 0 & 0 & \rho_{22} \end{pmatrix}. \tag{5.48}$$

This is the local equilibrium condition selected by the leading order dephasing channel.

Again, I remind the reader that it does not mean that the full density matrix is classical or exactly diagonal. Higher order corrections $\hat{\boldsymbol{\rho}}_1, \hat{\boldsymbol{\rho}}_2, \ldots$ still contain coherences generated by the coherent unitary step. Instead, this shows how dephasing in a DTQW can play a role analogous to relaxation in LBM. It selects a local equilibrium sector that the leading order dynamics relaxes towards, while higher order non-equilibrium corrections mediate transport.

At $\mathcal{O}(\varepsilon^2)$, equating LHS_2 with $\mathcal{U}_2 + \mathcal{D}_2$ gives

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_1 + \partial_{t_1}\hat{\boldsymbol{\rho}}_0 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_0 - k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0] + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1] = 0. \tag{5.49}$$

Using the leading slow sector condition

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_0 = 0, \tag{5.50}$$

this simplifies to

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_1 + \partial_{t_1}\hat{\boldsymbol{\rho}}_0 - k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0] + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1] = 0. \tag{5.51}$$

Projecting Eq. (5.51) onto the diagonal sector gives

$$\partial_{t_1}\hat{\boldsymbol{\rho}}_0 = 0. \tag{5.52}$$

This follows because $\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]$ is off-diagonal when $\hat{\boldsymbol{\rho}}_0$ is diagonal, and $\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1]$ is also off-diagonal. Thus there is no first order population transport in the unbiased parity echo model. The off-diagonal part of Eq. (5.51) gives the damped coherence equation

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} + \bar{\gamma}\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} = k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]. \tag{5.53}$$

Thus the first correction is not zero. The unitary walk regenerates coherences, while the dephasing channel damps them.

At $\mathcal{O}(\varepsilon^3)$, equating LHS_3 with $\mathcal{U}_3 + \mathcal{D}_3$ gives

$$\begin{aligned} 0 =& \partial_{t_0}\hat{\boldsymbol{\rho}}_2 + \partial_{t_1}\hat{\boldsymbol{\rho}}_1 + \partial_{t_2}\hat{\boldsymbol{\rho}}_0 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_1 + \partial_{t_0}\partial_{t_1}\hat{\boldsymbol{\rho}}_0 + \frac{1}{6}\partial_{t_0}^3\hat{\boldsymbol{\rho}}_0 \\ & - k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_1] - ik\bar{\theta}^2\mathcal{A}_3[\hat{\boldsymbol{\rho}}_0] - ik^2\bar{\theta}\bar{\mathcal{B}}_3[\hat{\boldsymbol{\rho}}_0] \\ & + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_2] - \bar{\gamma}\Pi_c[k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]]. \end{aligned} \tag{5.54}$$

Using

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_0 = 0, \qquad \partial_{t_1}\hat{\boldsymbol{\rho}}_0 = 0, \tag{5.55}$$

this may be written in the slightly cleaner form

$$\begin{aligned} 0 =& \partial_{t_0}\hat{\boldsymbol{\rho}}_2 + \partial_{t_1}\hat{\boldsymbol{\rho}}_1 + \partial_{t_2}\hat{\boldsymbol{\rho}}_0 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_1 \\ & - k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_1] - ik\bar{\theta}^2\mathcal{A}_3[\hat{\boldsymbol{\rho}}_0] - ik^2\bar{\theta}\bar{\mathcal{B}}_3[\hat{\boldsymbol{\rho}}_0] \\ & + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_2] - \bar{\gamma}\Pi_c[k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]]. \end{aligned} \tag{5.56}$$

This is the end of the unclosed CE hierarchy. It already contains explicit k^2-structure through the $ik^2\bar{\theta}\bar{\mathcal{B}}_3[\hat{\boldsymbol{\rho}}_0]$ term, and it also contains the coherence term $k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_1]$. However, because the equation still contains the unknown correction $\hat{\boldsymbol{\rho}}_1$, it is not yet a closed equation for the diagonal populations. To make the hierarchy more compact, I may define

$$\mathcal{T}_2[\hat{\boldsymbol{\rho}}] = k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}], \tag{5.57}$$

and

$$\mathcal{T}_3[\hat{\boldsymbol{\rho}}] = ik\bar{\theta}^2\mathcal{A}_3[\hat{\boldsymbol{\rho}}] + ik^2\bar{\theta}\bar{\mathcal{B}}_3[\hat{\boldsymbol{\rho}}]. \tag{5.58}$$

Here $\mathcal{T}_2$ is the first coherent transport superoperator, and $\mathcal{T}_3$ contains the third order coherent corrections, including the explicit k^2 term. With this notation, the three non trivial CE equations can be written as

$$\mathcal{O}(\varepsilon): \qquad \partial_{t_0}\hat{\boldsymbol{\rho}}_0 + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0] = 0, \tag{5.59}$$

$$\mathcal{O}(\varepsilon^2): \qquad \partial_{t_0}\hat{\boldsymbol{\rho}}_1 + \partial_{t_1}\hat{\boldsymbol{\rho}}_0 - \mathcal{T}_2[\hat{\boldsymbol{\rho}}_0] + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1] = 0, \tag{5.60}$$

and

$$\begin{aligned}\mathcal{O}(\varepsilon^3): \qquad 0 =& \partial_{t_0}\hat{\boldsymbol{\rho}}_2 + \partial_{t_1}\hat{\boldsymbol{\rho}}_1 + \partial_{t_2}\hat{\boldsymbol{\rho}}_0 + \frac{1}{2}\partial_{t_0}^2\hat{\boldsymbol{\rho}}_1 \\ & - \mathcal{T}_2[\hat{\boldsymbol{\rho}}_1] - \mathcal{T}_3[\hat{\boldsymbol{\rho}}_0] + \bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_2] - \bar{\gamma}\Pi_c[\mathcal{T}_2[\hat{\boldsymbol{\rho}}_0]].\end{aligned} \tag{5.61}$$

This more compact form is analogous in spirit to the way CE hierarchies are written in LBM as the relaxation operator, transport operator, and higher order correction operators are separated. It keeps the order by order structure clear while making the appearance of the k^2-term transparent through the definition of $\mathcal{T}_3$.

5.2.2 Overdamped coherence closure

The CE hierarchy previously derived gives a density matrix form of the hierarchy. In this form, the continuum limit is not yet a closed diffusion equation for the diagonal populations alone. This is not a failure of the expansion. Rather, it shows that the natural non-equilibrium variables of the open quantum walk are the off-diagonal coherence corrections. These coherences are generated by the unitary parity echo step and damped by the dephasing channel, and they mediate the feedback into the population sector at the next order. To make contact with a classical diffusion equation, an additional closure step is needed. This is directly analogous to the Chapman-Enskog derivation in LBM. In LBM, the diffusive term is not obtained by setting the first non-equilibrium correction to zero. Instead, the first correction is solved from the lower order kinetic equation and then substituted into the next order moment equation. Here, the analogous object is the first off-diagonal coherence correction, $\hat{\boldsymbol{\rho}}_1^{\text{off}}$. The closure therefore consists of solving, or approximating, this coherence correction in terms of the leading order diagonal populations $\hat{\boldsymbol{\rho}}_0$. In the following, I use an overdamped coherence closure. This assumes that the off-diagonal coherence sector relaxes quickly compared with the slow population dynamics, so that after a short transient it is approximately slaved to $\hat{\boldsymbol{\rho}}_0$. Substituting this slaved coherence back into the population equation gives an effective k^2 response in the diagonal sector, and hence a diffusive equation. This closure is therefore not imposed at the start of the CE hierarchy. It is a reduction of the coupled population-coherence system obtained from the previous order by order expansion.

The diagonal population sector evolves according to

$$\partial_{t_2}\hat{\boldsymbol{\rho}}_0 = \Pi_c\left[k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_1^{\mathrm{off}}]\right], \tag{5.62}$$

while the first order coherence correction satisfies

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} + \bar{\gamma}\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} = k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]. \tag{5.63}$$

This is the retained coherence form of the continuum limit. It does not assume that the first order coherences vanish. Instead, $\hat{\boldsymbol{\rho}}_1^{\mathrm{off}}$ is a damped coherence field generated by the unitary parity echo step and damped by coin dephasing.

To make contact with classical diffusion and with the lattice Boltzmann method, I now consider the overdamped coherence regime. Here, "overdamped" means that the off-diagonal coherence correction does not behave as an independent oscillatory mode. Instead, it relaxes rapidly under coin dephasing and reaches a balance between coherent generation by the walk and damping by the dephasing channel. This assumption is consistent with the weak decoherence scaling used above: the coherent correction and the dephasing term both enter the continuum hierarchy at finite order, so neither is removed from the limiting dynamics. The closure is therefore obtained by balancing these two effects, rather than by neglecting either the unitary walk or the dissipative relaxation. In Eq. (5.63), this corresponds to neglecting the remaining fast time derivative,

$$\partial_{t_0}\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} \simeq 0. \tag{5.64}$$

This is not an identity of the discrete walk, but an asymptotic closure valid when the coherence sector relaxes faster than the diagonal population sector evolves. Under this overdamped approximation, Eq. (5.63) gives

$$\hat{\boldsymbol{\rho}}_1^{\mathrm{off}} \simeq \frac{k\bar{\theta}}{\bar{\gamma}}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]. \tag{5.65}$$

Thus the first order coherences are not set to zero. Rather, they are slaved to the leading order diagonal populations through the balance between coherent generation

and dephasing.

This time scale separation is analytically supported by the structure of the dephasing operator. The diagonal population sector lies in the nullspace of $\mathcal{D}_c$, while the off-diagonal coherence sector is damped at rate $\bar{\gamma}$. The coherent driving term has size $k\bar{\theta}$. Therefore the overdamped coherence regime corresponds to

$$\frac{k\bar{\theta}}{\bar{\gamma}} \ll 1. \tag{5.66}$$

This condition can be satisfied in two ways. One possibility is to make the dephasing rate large compared with the coherent driving. However, this is not the regime I want to rely on here, since very strong dephasing can suppress transport and lead to a undesired slowing of the walk. Instead, the intended regime is a weak decoherence continuum regime, where the microscopic dephasing strength scales as $\gamma = \varepsilon\bar{\gamma}$, but the relevant Fourier modes are long wavelength. In this case k is small, so the ratio $k\bar{\theta}/\bar{\gamma}$ can be small even when the microscopic decoherence is weak. Thus the closure should be understood as a long wavelength, weak decoherence approximation. The dephasing must be strong enough, in the rescaled sense, to relax the coherence correction relative to the coherent driving $k\bar{\theta}$, but not so strong that it freezes the transport. This is why the approximation is expected to work best for broad, smooth initial states and weak to moderate dephasing. There, the small k modes dominate, coherent oscillations are sufficiently damped, and the transport is not yet strongly suppressed.

For broad, smooth initial probability distributions, the dominant Fourier modes are long wavelength modes. If the initial packet has width σ_0, then a typical wavenumber is

$$k_{\text{typ}} \sim \frac{1}{\sigma_0}. \tag{5.67}$$

The coherence slaving condition can therefore be estimated as

$$\frac{k_{\text{typ}}\bar{\theta}}{\bar{\gamma}} \sim \frac{\bar{\theta}}{\bar{\gamma}\sigma_0} \ll 1. \tag{5.68}$$

Thus broad initial states suppress large k modes and make the continuum diffusion

closure more reliable. Substituting Eq. (5.65) into Eq. (5.62) gives

$$\partial_{t_2}\hat{\boldsymbol{\rho}}_0 = \frac{k^2\bar{\theta}^2}{\bar{\gamma}}\Pi_c\left[\mathcal{A}_2[\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]]\right]. \tag{5.69}$$

This is the same structural role played by the first non-equilibrium correction in LBM. In the LBM Chapman-Enskog derivation, the diffusive or viscous term is not obtained by setting $f_i^{(1)} = 0$. Instead, $f_i^{(1)}$ is solved from the lower order kinetic equation and substituted into the next order moment equation. Here, the analogous non-equilibrium variable is the coherence correction $\hat{\boldsymbol{\rho}}_1^{\text{off}}$. Writing

$$\hat{\boldsymbol{\rho}}_0 = \begin{pmatrix} \rho_{00} & 0 & 0 \\ 0 & \rho_{11} & 0 \\ 0 & 0 & \rho_{22} \end{pmatrix}, \tag{5.70}$$

one obtains

$$\Pi_c\left[\mathcal{A}_2[\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]]\right] = -\begin{pmatrix} \rho_{00}-\rho_{11} & 0 & 0 \\ 0 & -\rho_{00}+2\rho_{11}-\rho_{22} & 0 \\ 0 & 0 & \rho_{22}-\rho_{11} \end{pmatrix}. \tag{5.71}$$

Therefore, defining

$$\mathbf{f} = \begin{pmatrix} \rho_{00} \\ \rho_{11} \\ \rho_{22} \end{pmatrix}, \tag{5.72}$$

the overdamped population equation becomes

$$\partial_{t_2}\mathbf{f} = -\frac{\bar{\theta}^2}{\bar{\gamma}}k^2\begin{pmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{pmatrix}\mathbf{f}. \tag{5.73}$$

In position space, using $k^2 \to -i\partial_x^2$, this has the form of a diffusive population equation.

The following section interprets the CE hierarchy and the closure derivation developed above. The subsequent sections then test these analytic results using numerical

simulations of the discrete parity echo model. These simulations show that the behaviour predicted by the CE analysis is visible in the full dynamics, providing evidence that the expansion is capturing the relevant transport mechanism rather than only giving a formal asymptotic description.

5.3 Interpretation of the hierarchy

The compact hierarchy, given by Eqns 5.59 5.60 and 5.61, separates the different physical roles of the terms in the expansion. The leading relaxation term, $\bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0]$, selects the local equilibrium sector by damping the coin basis coherences of $\hat{\boldsymbol{\rho}}_0$. Its nullspace is the diagonal coin sector, and so this term is responsible for the condition $\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0] = 0$. In other words, at leading order the dephasing channel selects the diagonal population sector as the slow manifold. This is a natural equilibrium condition for an open quantum walk with coin dephasing. The dephasing channel suppresses the coin basis coherences, so even though the coherent unitary step may continually regenerate off-diagonal terms at higher orders, the leading relaxation mechanism drives the system towards a classical looking diagonal sector. In this sense, the diagonal population sector plays the same role as a local equilibrium manifold in LBM.

The next term, $\mathcal{T}_2[\hat{\boldsymbol{\rho}}_0] = k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]$, is the first coherent transport contribution. Since A_2 is off-diagonal in the coin basis, this term does not directly produce population transport when $\hat{\boldsymbol{\rho}}_0$ is diagonal. Instead, it generates the first off-diagonal coherence correction. The term $\bar{\gamma}\mathcal{D}_c[\hat{\boldsymbol{\rho}}_1]$ then damps these generated coherences. Thus the second order equation describes a balance between coherent generation and dephasing induced relaxation of $\hat{\boldsymbol{\rho}}_1^{\text{off}}$.

At third order, $\mathcal{T}_3[\hat{\boldsymbol{\rho}}_0] = ik\bar{\theta}^2\mathcal{A}_3[\hat{\boldsymbol{\rho}}_0] + ik^2\bar{\theta}\mathcal{B}_3[\hat{\boldsymbol{\rho}}_0]$ contains the higher order coherent corrections, including the explicit k^2 structure required for a second spatial derivative. The term $\mathcal{T}_2[\hat{\boldsymbol{\rho}}_1]$ describes the feedback of the first coherence correction into the next order of the hierarchy. This is the mechanism by which coherences generated at $\mathcal{O}(\varepsilon^2)$ can affect the slow diagonal sector at $\mathcal{O}(\varepsilon^3)$.

The hierarchy therefore has the same broad structure as a kinetic CE expansion. First, the leading equilibrium sector is selected. Then non-equilibrium corrections are

generated. Finally, those corrections feed back into the slow variables. At $\mathcal{O}(\varepsilon)$, the parity echo model gives a pure dephasing equation. In the slow sector this selects $\mathcal{D}_c[\hat{\boldsymbol{\rho}}_0] = 0$, so the leading order state is diagonal in the coin basis. This is the local equilibrium condition for the open quantum walk. It is physically natural. Under coin dephasing, the fast coin basis coherences are damped, while the diagonal populations remain as slow variables. This does not mean that the full walk has become classical. It only means that the leading order equilibrium manifold is classical looking, in the same sense that the local equilibrium in LBM is a population distribution.

The next orders show why the model is still genuinely quantum. At $\mathcal{O}(\varepsilon^2)$, the coherent parity echo term generates off-diagonal corrections from the diagonal leading state. Thus the first correction $\hat{\boldsymbol{\rho}}_1$ plays the role of a non-equilibrium coherence field. These coherences are damped by dephasing, but they are not set to zero. They are the quantities through which the coherent walk mediates transport.

At $\mathcal{O}(\varepsilon^3)$, the CE hierarchy contains the first explicit k^2 structure. This is the order at which second spatial derivative terms enter the continuum expansion. The important point is that the hierarchy contains two ingredients relevant for diffusion. The explicit k^2 terms from the third order unitary coefficient, and the feedback terms in which the $\mathcal{O}(\varepsilon^2)$ coherences feed back into the diagonal sector. Therefore the CE expansion itself shows that the parity echo model has the second derivative structure required for diffusive transport.

To obtain a closed classical diffusive equation, one must then make an additional closure step. This is analogous to LBM, where the diffusive term is not obtained by setting the first non-equilibrium correction to zero, but by solving for that correction from the lower order kinetic equation and substituting it into the next order equation. Here, the analogous non-equilibrium correction is the coherence field $\hat{\boldsymbol{\rho}}_1^{\text{off}}$. In the overdamped coherence regime, these coherences are assumed to relax rapidly compared with the slow population dynamics, so they may be slaved to $\hat{\boldsymbol{\rho}}_0$. This gives an effective k^2 response in the diagonal population sector.

The closure should therefore be interpreted as a reduction of the CE hierarchy, rather than the hierarchy itself. This approximation is expected to be most accurate in

the small k regime, corresponding physically to long wavelength variations where the probability distribution changes slowly over many lattice sites. In the numerical simulations, this regime is favoured by choosing a broad initial Gaussian distribution with width $\sigma_0 = 10$. Its Fourier representation is concentrated around small wavenumbers, with a characteristic scale $k_{\text{typ}} \sim 1/\sigma_0 \approx 0.1$ in lattice units. Although the simulations evolve the full range of lattice wavenumbers, the initial state predominantly excites long wavelength modes, for which the coherence slaving condition $\bar{\theta}/(\bar{\gamma}\sigma_0) \ll 1$ is more likely to hold. The approximation may therefore become less accurate for narrower initial distributions, weak dephasing, or stronger coherent mixing. The approximately linear variance growth observed in the following simulations supports the emergence of diffusive behaviour in this long wavelength regime.

There are, however, some limitations to this form of the hierarchy. Although the parity echo construction removes the leading ballistic term and gives a cleaner ordering than the generic walk, the third order equation is still fairly complicated. It contains several different contributions. Explicit k^2 terms, higher order coherent corrections, dephasing terms, and feedback from the first coherence correction $\hat{\boldsymbol{\rho}}_1^{\text{off}}$. Therefore the CE expansion does not automatically produce a closed scalar diffusion equation. Instead, it produces a hierarchy which must be reduced by an additional closure assumption.

5.4 Diffusion Simulations

The parity echo model will now be studied numerically to investigate how it behaves as a transport model and how it may be used to simulate both diffusion and later, advection-diffusion. The following subsections examine the effect of the dephasing strength γ, the coin angle θ. Particular attention is paid to how these parameters influence the effective diffusion coefficient, the shape of the probability density, and the quality of the agreement with the corresponding LBM solutions. These numerical results also provide an important validation of the CE analysis developed above, allowing the continuum predictions to be compared directly with the full discrete dynamics. The initial coin state used for the diffusion simulations is chosen to have equal weight in the left moving, rest, and right moving coin states, with a relative phase on the rest

component. Explicitly, in the $|L, S, R\rangle$ basis,

$$\chi_{\text{diff}} = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ i \\ 1 \end{pmatrix}. \tag{5.74}$$

5.4.1 Decoherence Strength

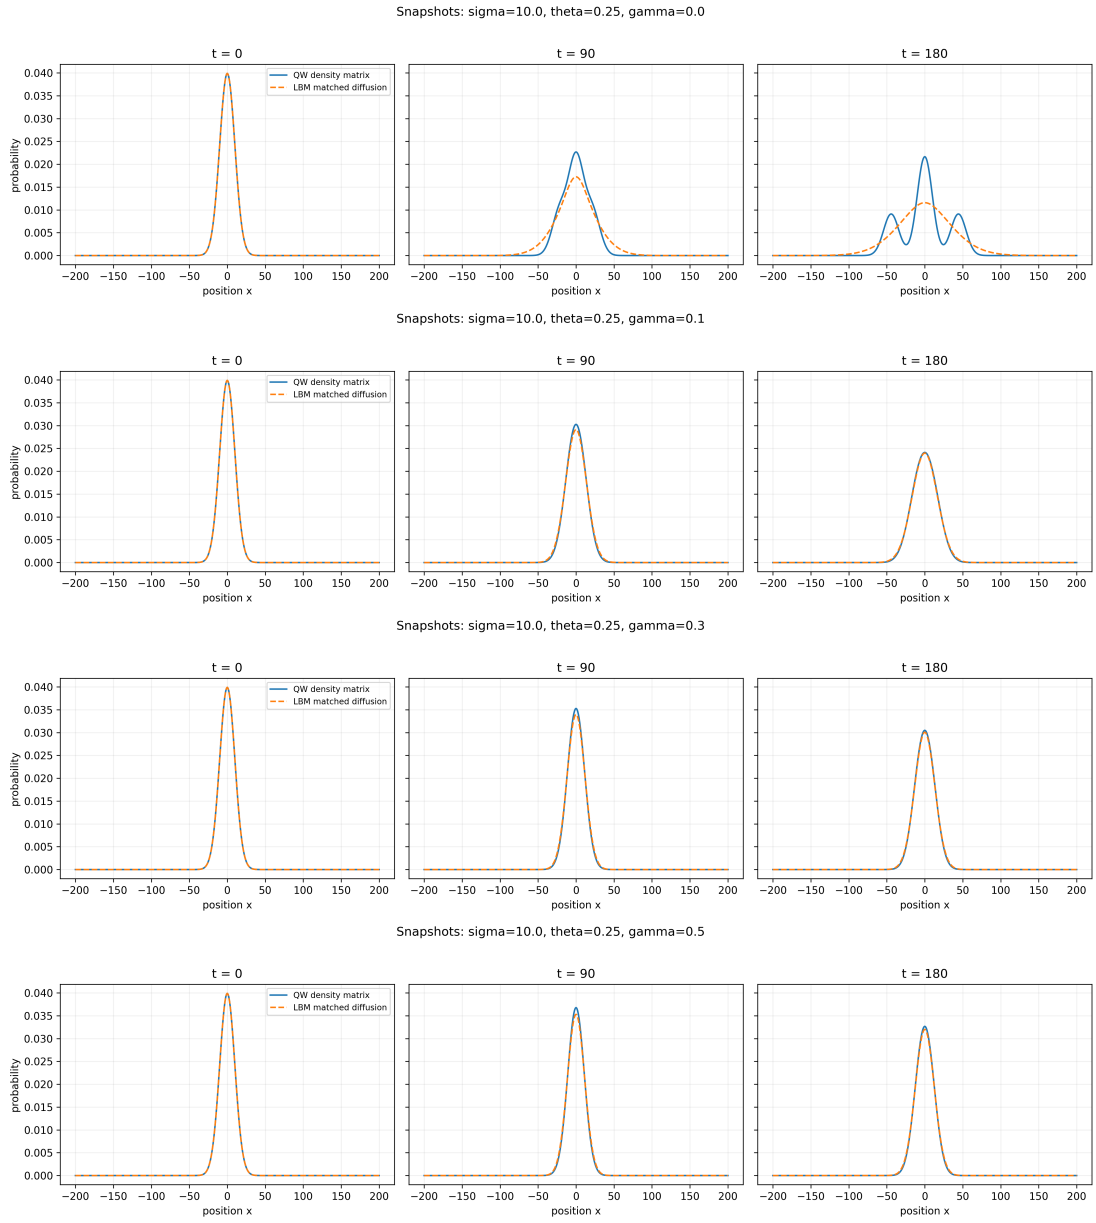

Figure 5.3: Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths are $\gamma = 0$, $\gamma = 0.1$, $\gamma = 0.3$, and $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$, while the initial coin state is defined in Eq. (5.74). In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve.

Figure 5.3 shows snapshots of the probability distributions of the parity echo quantum walk with decoherence for several different dephasing strengths γ. In each case, the quantum walk density is compared with a matched classical D1Q3 lattice Boltzmann solution, following the procedure described in Appendix A.4. The coin angle is kept fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian packet with width $\sigma_0 = 10$. For each value of γ, an effective diffusion coefficient is fitted from the quantum walk variance growth using $\sigma^2(t) \simeq \sigma_0^2 + 2D_{\text{eff}}t$. This fitted value is then used to run the corresponding D1Q3 LBM diffusion model, allowing the full probability density of the quantum walk to be compared with a classical diffusion profile. A representative subset of the γ sweep is shown here in the main text. The full set of probability density snapshots and variance curves for all simulated dephasing strengths is included in Appendix A.5. These additional plots are used to check that the trends discussed here are not specific to the selected examples.

The $\gamma = 0$ case shows the coherent unitary behaviour of the parity echo walk without dephasing. In this case there is visible ballistic structure, and the probability density separates into components associated with the left moving, rest, and right moving sectors. As dephasing is introduced, these coherent features are progressively damped and the density becomes smoother and more diffusive. However, increasing γ does not simply make the packet spread more. At larger values, such as $\gamma = 0.5$, the distribution spreads less than in the weaker dephasing cases because the coherences that mediate transport are damped more strongly. This behaviour is consistent with both the analytic closure derived above and the known behaviour of decoherence in DTQWs. In a fully coherent walk, interference produces ballistic spreading. Adding decoherence suppresses these interference effects and can produce a more classical, diffusive profile. However, too much decoherence can also suppress the coherent transport mechanism itself, reducing the effective spreading rate. The useful regime is therefore an intermediate one: there must be enough dephasing to damp the unwanted ballistic structure, but not so much that the coherences generated by the unitary parity echo step are destroyed before they can mediate transport. This is precisely the balance assumed in

the Chapman-Enskog closure, where the off-diagonal coherences are generated by the coherent walk and then relaxed by the dephasing channel.

Figure 5.4 shows the fitted diffusion coefficient as a function of the dephasing strength. The fitted diffusion coefficient decreases as γ is increased. This trend is consistent with the overdamped coherence closure, which predicts $D_{\text{eff}} \sim \bar{\theta}^2/\bar{\gamma}$. In this picture, diffusion is mediated by coherences generated by the unitary parity echo step and then damped by the dephasing channel. Increasing γ damps these coherences more rapidly, reducing their feedback into the diagonal population sector and therefore lowering the effective diffusion coefficient. The dependence of the fitted diffusion coefficient on the dephasing strength supports the overdamped coherence interpretation. The CE closure predicts that the first order coherence correction is approximately slaved to the population sector as $\hat{\boldsymbol{\rho}}_1^{\text{off}} \sim (k\bar{\theta}/\bar{\gamma})\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]$, so that the effective diffusive response scales approximately as $D_{\text{eff}} \propto \bar{\theta}^2/\bar{\gamma}$. To test this numerically, the fitted diffusion coefficients were compared with an inverse γ fit of the form $D_{\text{eff}}(\gamma) = A/\gamma + B$. Excluding the singular $\gamma = 0$ case and fitting over the range $0.02 \leq \gamma \leq 0.5$ gives $D_{\text{eff}}(\gamma) \approx 0.03596/\gamma + 0.11087$, with $R^2 \simeq 0.995$. This high R^2 value shows that the inverse γ form explains the measured diffusion coefficients very well over this dephasing range. This numerical trend agrees with the analytic picture obtained from the CE hierarchy. In the overdamped closure, stronger dephasing reduces the size of the coherence correction $\hat{\boldsymbol{\rho}}_1^{\text{off}}$, and therefore reduces its feedback into the diagonal population sector. The observed inverse relationship between D_{eff} and γ is therefore not only a numerical feature of the simulations, but directly supports the analytic mechanism identified in the CE expansion. This provides evidence that the CE style expansion is capturing a real transport mechanism of the discrete parity echo walk, rather than only giving a formal asymptotic hierarchy. Thus the inverse γ scaling should be understood as a description of the effective diffusion coefficient within the weak to moderate dephasing regime where the density profiles are sufficiently smooth and the variance growth is approximately linear.

The $\gamma = 0$ point is excluded from the inverse γ fit because the overdamped closure requires a non-zero dephasing rate. Very weak dephasing strengths ($\gamma < 0.03$) are also

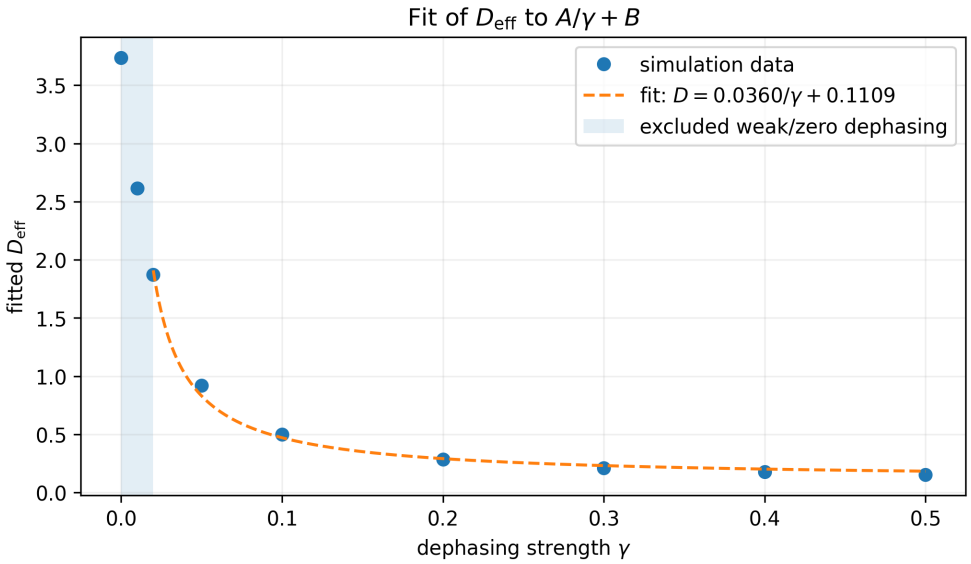

Figure 5.4: Inverse dephasing fit for the effective diffusion coefficient. The fitted diffusion coefficients are compared with a curve of the form $D_{\text{eff}}(\gamma) = A/\gamma + B$. The $\gamma = 0$ point is excluded from the fit because the overdamped coherence closure requires a nonzero dephasing rate.

excluded, despite their fitted D_{eff} values following the expected inverse γ trend. This is because the corresponding probability distributions retain a pronounced three peaked structure, characteristic of coherent ballistic spreading, rather than the smooth profile expected for classical diffusion. Thus, agreement with the predicted scaling of D_{eff} does not necessarily imply genuinely diffusive behaviour, since the fitted coefficient characterises variance growth without fully describing the shape of the distribution. The most useful diffusive regime instead occurs at intermediate dephasing strengths, where the variance grows approximately linearly and the probability density agrees more closely with the matched LBM profile. At stronger dephasing, the coherences responsible for transport are increasingly suppressed, reducing the effective diffusion coefficient and limiting spatial spreading. Appendix A.5 presents the variance growth for the full range of dephasing strengths considered.

Appendix A.5 shows the variance growth plots for each value of γ in the dephasing sweep. For $\gamma = 0$, the DTQW variance reflects the underlying coherent spreading of the unitary parity echo walk, as expected. As γ is increased, the variance growth

becomes closer to linear, showing that dephasing helps regularise the dynamics and produces a more diffusive second moment behaviour. Interestingly, this improvement is not monotonic. One might expect that increasing γ would continue to suppress coherent oscillations and therefore make the variance increasingly linear. Instead, at larger values of γ, the variance begins to deviate slightly from an ideal linear fit again, while the overall spreading is reduced. This suggests that strong dephasing does not simply produce a better diffusion limit. Rather, it damps the coherences that mediate transport too strongly, so the walk becomes less effective at spreading across the lattice. This supports the interpretation of the parity echo model as requiring a balance between coherent generation and dissipative relaxation. Weak to moderate dephasing is sufficient to suppress the unwanted ballistic structure and produce nearly linear variance growth, but excessive dephasing suppresses the transport mechanism itself. The most useful regime is therefore an intermediate one, where the variance growth is close to linear while the packet still spreads appreciably.

The dephasing strength therefore controls a trade off. Weak dephasing allows rapid but coherent spreading, while strong dephasing suppresses transport. This is demonstrated in both the probability distribution plots and the variance plots. The best classical diffusion approximation occurs at weak to intermediate γ, where the density profile is smooth and the variance growth is approximately linear. This is closely analogous to the BGK relaxation structure in LBM. In standard LBM, the diffusion coefficient scales as $D = c_s^2(\tau - 1/2)$, where τ is the relaxation time, or equivalently $D = c_s^2(1/\omega - 1/2)$ in terms of the relaxation rate $\omega = 1/\tau$. In the parity echo model, the dephasing strength γ acts like a relaxation rate for the coherence sector, so $1/\gamma$ plays the role of a coherence relaxation time. The observed scaling $D_{\text{eff}} \propto 1/\gamma$ therefore has the same structure as the LBM result. Longer lived non-equilibrium corrections give a larger diffusive response.

5.4.2 Coin Angle

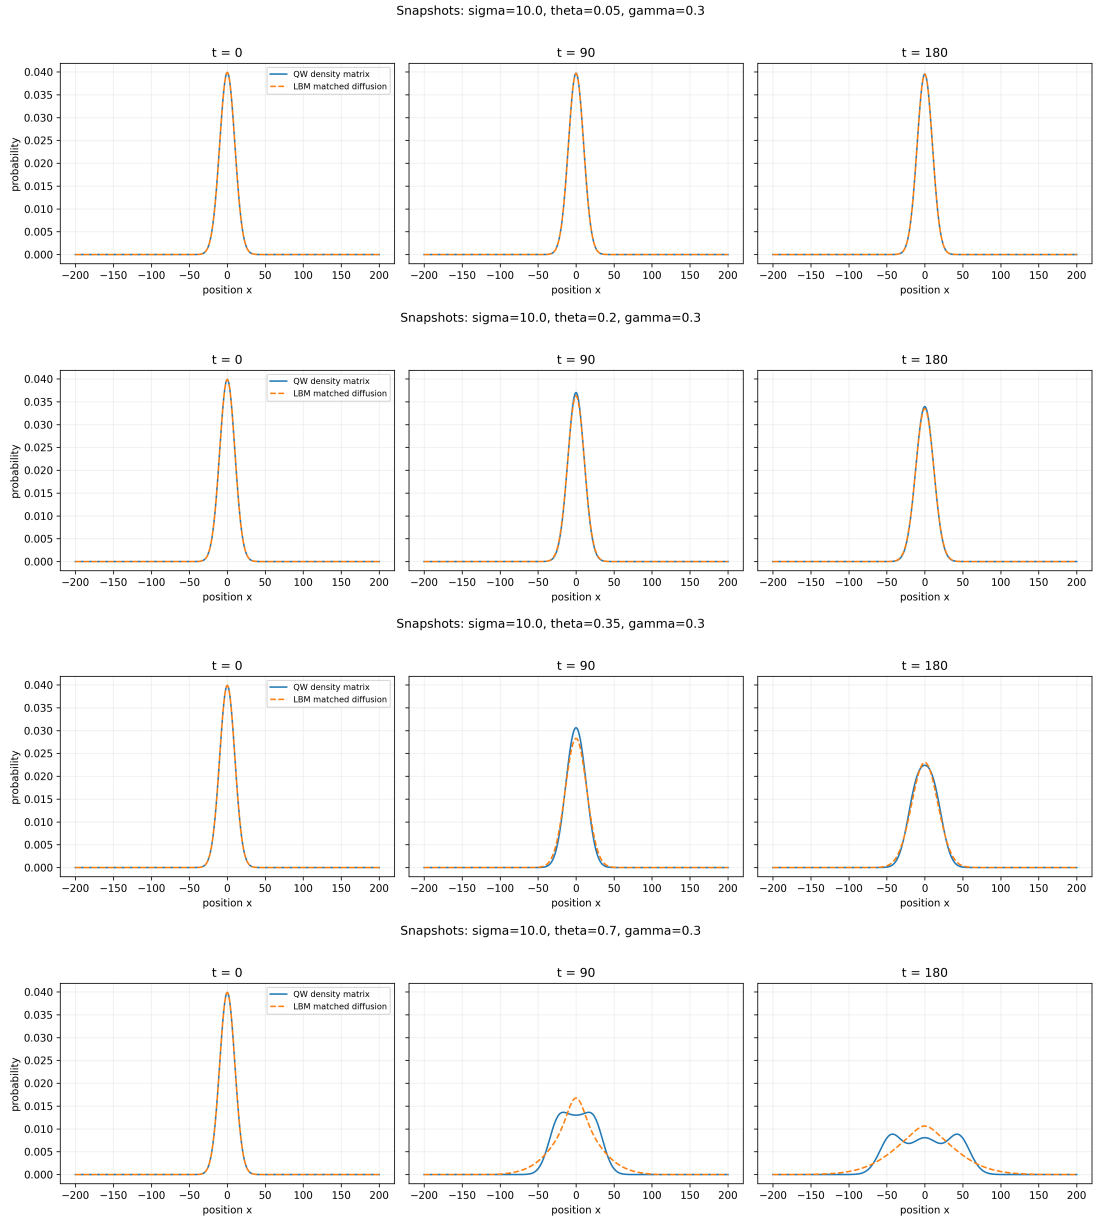

Figure 5.5: Probability density snapshots from the θ sweep for the unbiased parity echo quantum walk at fixed $\gamma = 0.3$. The quantum walk density is compared with the matched D1Q3 LBM diffusion solution at $t = 0$, $t = 90$, and $t = 180$. The rows show $\theta = 0.05$, $\theta = 0.2$, $\theta = 0.35$, and $\theta = 0.7$. The spatial initial state is a Gaussian of width $\sigma_0 = 10$, while the unbiased initial coin state is defined in Eq. (5.74).

The probability density snapshots for the parity echo model with ranging values of θ is shown in Figure 5.5, with $\gamma = 0.3$. Again, in each case, the quantum walk density is compared with a matched classical D1Q3 lattice Boltzmann solution, following the procedure described in Appendix A.4. These plots show that the coin angle controls the qualitative behaviour of the walk. For very small θ, there is only weak mixing between the rest state and the propagating states, so the packet spreads only slowly and the effective diffusion is small. For intermediate values of θ, the density becomes smoother and more diffusive, giving the best agreement with the fitted LBM profile. For larger values, such as $\theta = 0.7$, coherent quantum walk structure begins to reappear, and the probability density develops more visible multipeaked features. Thus, even at fixed dephasing strength, changing the coin angle can move the model between weak diffusion, a useful diffusive regime, and a more coherent, non Gaussian regime. Only representative values of θ are shown in this section to keep the main text focused. The full θ sweep, including all probability density snapshots and variance growth plots, is included in Appendix A.6. The appendix plots show the same trend across the full parameter range.

As shown in Fig. 5.6, the fitted diffusion coefficient increases strongly with the parity echo coin angle. This is consistent with the CE closure prediction $D_{\text{eff}} \propto \bar{\theta}^2/\bar{\gamma}$, since increasing θ strengthens the coherent mixing between the rest state and the propagating states. In the overdamped coherence picture, the first coherence correction is generated by the coherent term $k\bar{\theta}\mathcal{A}_2[\hat{\boldsymbol{\rho}}_0]$ and is then damped by the dephasing channel. Substituting this slaved coherence correction back into the population equation gives the approximate scaling $D_{\text{eff}} \propto \bar{\theta}^2/\bar{\gamma}$. Thus, at fixed dephasing strength, the effective diffusion coefficient is expected to scale approximately quadratically with the coin angle, provided the system remains in the perturbative regime assumed by the CE expansion. To test the predicted scaling, the fitted diffusion coefficients were compared with a quadratic fit of the form $D_{\text{eff}}(\theta) = A\theta^2 + B$. Fitting over the small to moderate angle range $\theta \leq 0.4$ gives $D_{\text{eff}}(\theta) \approx 4.7095\theta^2 - 0.0413$, with $R^2 \simeq 0.977$. This indicates that the quadratic form explains the perturbative angle data well, supporting the CE

prediction $D_{\text{eff}} \propto \theta^2$ in the regime where the assumptions of the expansion are expected to hold.

When the full range of angles is included, the fit becomes $D_{\text{eff}}(\theta) \approx 14.3082\theta^2 - 0.7545$, with a lower fit quality $R^2 \simeq 0.953$. Although this still captures the overall increasing trend, the larger θ values clearly deviate from the small angle scaling. This is expected because, again, the CE prediction comes from a truncated expansion in the coin angle. At larger θ, higher order terms in the unitary expansion become important, and the measured diffusion coefficient grows faster than the leading quadratic approximation predicts. In these cases, the probability density can again develop a more structured, multipeaked profile, showing that changing only the coin angle can move the model from a diffusive regime back towards a more coherent quantum walk regime. This is an important point. Even at fixed dephasing strength, the balance between coherent mixing and dephasing determines whether the observed dynamics is close to classical diffusion. The full θ sweep variance plots in Appendix A.6 show that the variance growth becomes less cleanly linear for the larger θ values, supporting the interpretation that the large angle regime lies outside the clean perturbative diffusion window. For this reason, the $\theta \leq 0.4$ fit is used as the perturbative small to moderate angle comparison, while the full range fit is shown only to illustrate the breakdown of the leading order scaling at larger coin angles. Together with the inverse γ trend, this provides numerical support for the analytic CE closure. The simulations reproduce the predicted parameter dependence $D_{\text{eff}} \sim \bar{\theta}^2/\bar{\gamma}$ in the regime where the expansion assumptions are valid. Outside this regime, the model does not simply become a better diffuser as θ is increased. Instead, large coherent mixing can reintroduce structured quantum walk features and reduce the quality of the classical diffusion approximation.

I now move on to introducing advection into the parity echo model. The first subsection presents a particular parameter choice and shows how the resulting walk can be used to simulate advection-diffusion. The following subsections then explore how the coin angle and decoherence strength affect the macroscopic behaviour of the model.

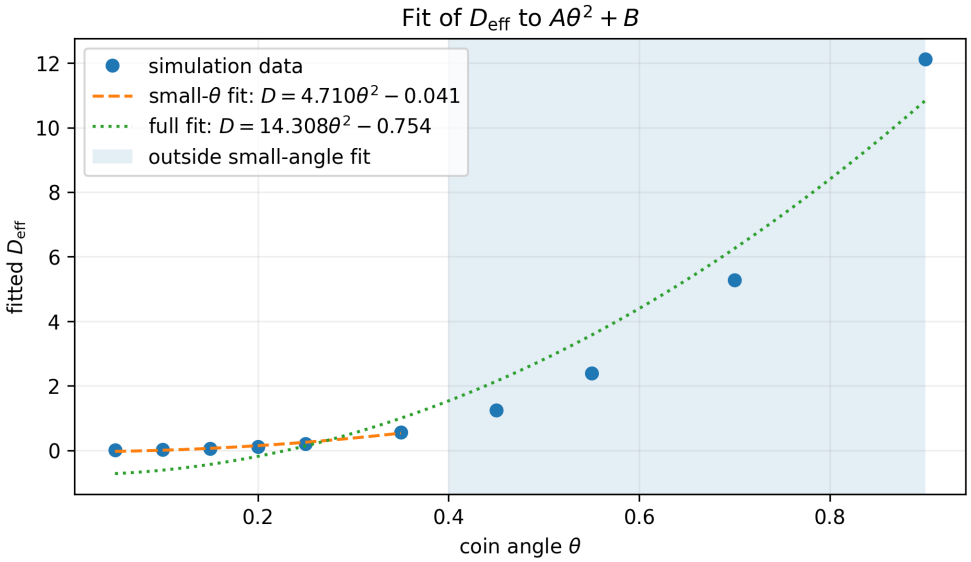

Figure 5.6: Quadratic fit for the dependence of the effective diffusion coefficient on the parity echo coin angle. The fitted diffusion coefficients are compared with curves of the form $D_{\text{eff}}(\theta) = A\theta^2 + B$. The small to moderate angle fit supports the CE prediction $D_{\text{eff}} \propto \theta^2$, while the larger-angle data show deviations from the leading perturbative scaling.

5.5 Advection-Diffusion Simulations

The previous simulations focused on the diffusive behaviour of the parity echo model with an unbiased initial coin state. Advection can also be introduced without changing the parity echo unitary itself, but instead by choosing an initial coin state with unequal left and right moving populations. In this section, I use such a biased initial state to test whether the same decoherence assisted parity echo mechanism can simulate an advection-diffusion process. In each case, the quantum walk density is compared with a matched classical D1Q3 lattice Boltzmann advection-diffusion solution, following the procedure described in Appendix A.4.

5.5.1 Advection-diffusion from a biased initial coin state

The evolution of the parity echo walk, given by Eqn. (4.9), is now used to simulate advection diffusion. The initial spatial state is a broad Gaussian wavepacket with width $\sigma_0 = 10$. To introduce a right moving drift, the initial coin state is biased towards the

right moving component. In the (L, S, R) basis, the initial coin state is

$$\chi_{\text{adv}} = \frac{1}{\mathcal{N}} \begin{pmatrix} 0.05 \\ 0.25 \\ 1.95 \end{pmatrix}, \tag{5.75}$$

where $\mathcal{N}$ is the normalisation constant. This initial state gives a much larger weight to the right moving state than to the left moving state, and therefore introduces a non-zero drift into the quantum walk probability density. Figure 5.7 shows snapshots

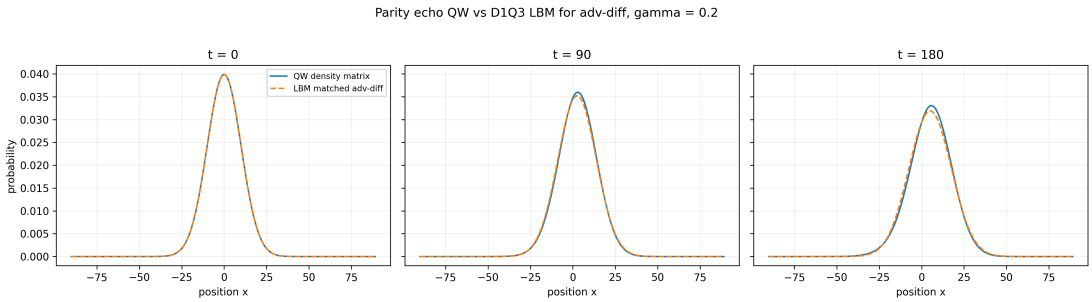

Figure 5.7: Comparison between the biased parity echo quantum walk and the matched D1Q3 lattice Boltzmann advection-diffusion model at selected time steps. The quantum walk is evolved using the direct density matrix dephasing channel.

of the probability density at selected times, comparing the biased parity echo quantum walk with the matched LBM advection-diffusion profile. The LBM drift velocity and diffusion coefficient are fitted from the measured quantum walk mean position and variance growth, respectively. The parity echo walk parameters are $\sigma_0 = 10$, $\gamma = 0.2$, $\theta = 0.25$, with total simulation time $T = 180$ on a lattice of length $N = 180$, with the coin initial state given by Eqn. (5.75). The agreement is very good across the simulated time window. The quantum and classical profiles move in the same direction, have nearly the same width, and remain visually close over the full evolution. This indicates that the biased parity echo model can reproduce a classical advection-diffusion profile to a relatively high degree of accuracy.

Figure 5.8 shows the variance growth of the biased parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model. Both curves grow approximately linearly in time, indicating diffusive spreading around the drifting cen-

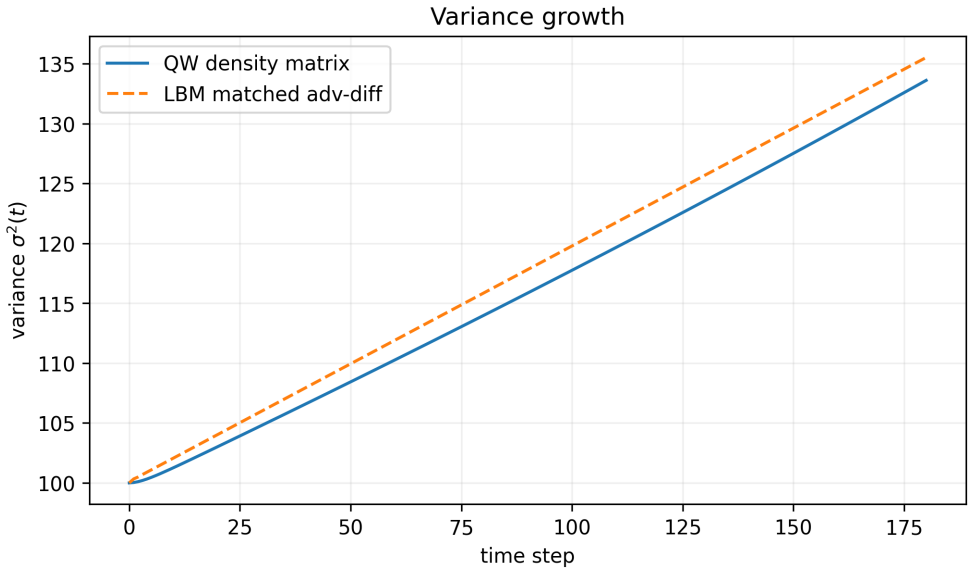

Figure 5.8: Variance growth of the biased parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model.

tre of mass. The quantum walk variance is very slightly lower than the LBM variance over this time window, meaning that the fitted classical model spreads slightly faster than the quantum walk. However, the two curves remain close and nearly parallel, showing that the parity echo walk reproduces the same overall diffusive scaling as the matched LBM model. Figure 5.9 shows the mean position and drift of the probability distributions. The mean position grows approximately linearly in time, confirming the presence of an effective advective component. The LBM and quantum walk curves are almost identical, with only a very small separation at late times. This indicates that the biased initial coin state introduces a controlled drift which is accurately captured by the fitted LBM advection velocity.

Taken together, these plots show that the biased parity echo quantum walk can closely reproduce a classical advection-diffusion process. The snapshots demonstrate agreement at the level of the full density profile, the mean position confirms the advective drift, and the variance confirms the diffusive spreading. The small discrepancy in the variance indicates that the classical LBM model spreads slightly faster than the quantum walk, but this difference is minor compared with the overall agreement. This

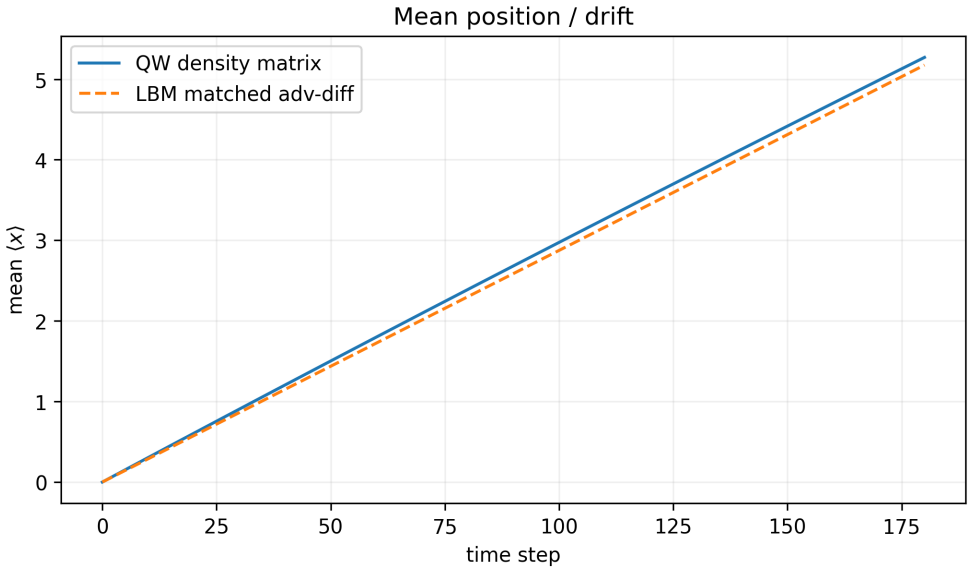

Figure 5.9: Mean position of the biased parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model.

supports the interpretation that initial state biasing introduces an effective advection term, while the decoherence assisted parity echo mechanism supplies the diffusive part of the dynamics.

5.5.2 Advection-diffusion parameter sweeps

The simulations above show that a biased parity echo walk can reproduce a matched classical advection-diffusion profile for one choice of parameters. To test whether this behaviour is robust over different parameter choices, additional simulations were performed over a range of dephasing strengths γ and coin angles θ. The full probability density snapshots and variance plots, and fitted summary curves are included in Appendix A.7 and A.8. These appendix plots are used to check how the advective and diffusive parts of the dynamics change as the model parameters are varied.

Decoherence Strength

Appendix A.7 shows the probability density snapshots and variance plots for the advection-diffusion simulations with different decoherence strengths γ. The snapshots

show that the very low dephasing cases retain more coherent, ballistic, and peaked quantum walk structure, so they do not match the fitted LBM profile as well. As γ is increased, the density becomes smoother and the agreement with the LBM advection-diffusion profile improves. Larger γ also reduces the amount of diffusive spreading, consistent with the decrease in D_{eff}, but the biased walk still retains a clear advective drift. The variance plots show the same trend. At very low γ, the variance growth is less diffusive, while at higher γ it becomes more linear and closer to the matched LBM variance.

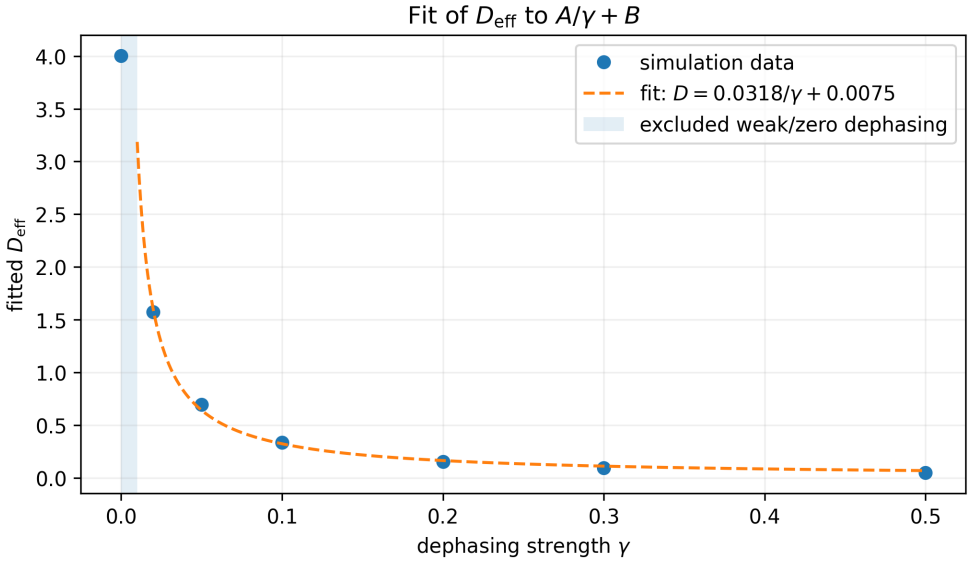

Figure 5.10: Inverse dephasing fit for the effective diffusion coefficient for the advection diffusion simulation. The fitted diffusion coefficients are compared with a curve of the form $D_{\text{eff}}(\gamma) = A/\gamma + B$. The $\gamma = 0$ point is excluded from the fit because the overdamped coherence closure requires a nonzero dephasing rate.

Figure 5.10 shows how the diffusion coefficient behaves when changing the decoherence strength γ. For the advection-diffusion γ sweep, the fitted effective diffusion coefficients were compared with an inverse dephasing model of the form $D_{\text{eff}} = A/\gamma + B$. Fitting over the range $0.01 \leq \gamma \leq 0.5$ gives $A = 0.03178307$ and $B = 0.00752610$, with $R^2 = 0.99773632$ and RMSE $= 2.53 \times 10^{-2}$. The high R^2 value shows that the inverse γ form captures the measured diffusion coefficients extremely well across this range. This agrees with the overdamped coherence picture again, as increasing γ damps the coher-

ences generated by the parity echo unitary more rapidly, reducing their feedback into the diagonal population sector and therefore lowering D_{eff}. The same inverse dephasing structure appears in the diffusion only simulations, suggesting that the diffusive part of the advection-diffusion dynamics is controlled by the same decoherence-assisted mechanism. The same sweep also shows that the fitted drift velocity changes with γ, shown

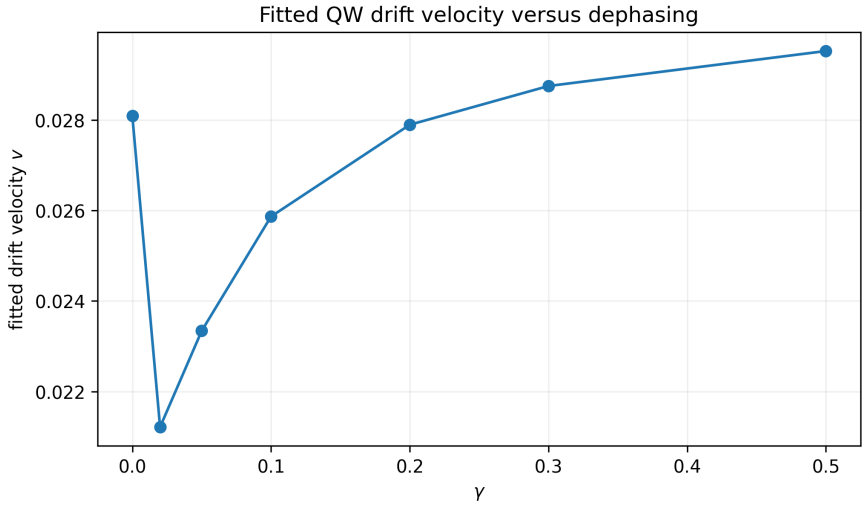

Figure 5.11: Fitted drift velocity for the biased parity echo advection-diffusion simulations as a function of the dephasing strength γ. The velocity is extracted from a late time linear fit to the mean position, $\langle x(t)\rangle \simeq \langle x(0)\rangle + vt$. The $\gamma = 0$ point corresponds to the fully coherent walk and is more affected by coherent oscillations, while for nonzero dephasing the fitted drift velocity increases and begins to level off at larger γ.

in Figure 5.11. The drift velocity is extracted from the first moment of the quantum walk probability distribution. At each time step, the mean position is computed as $\langle x(t)\rangle = \sum_x x\, p(x,t)$, and the drift velocity v is obtained by fitting the late-time mean position to a straight line, $\langle x(t)\rangle \simeq \langle x(0)\rangle + vt$. Apart from the $\gamma = 0$ point, which corresponds to the fully coherent walk and is therefore more sensitive to oscillatory structure in the mean position, the drift velocity increases with γ and then begins to level off for stronger dephasing. This suggests that dephasing helps regularise the biased transport generated by the initial coin imbalance, giving a cleaner effective advective motion. At larger γ, the fitted velocity changes more slowly, as larger decoherence values supress spreading.

Coin Angle

Appendix A.8 shows the corresponding snapshots and variance plots for the advection-diffusion θ sweep. Increasing θ generally increases the drift and spreading, because the coherent parity echo mixing becomes stronger. However, the largest θ cases do not match the LBM density profile as well. The packet moves and spreads too quickly, and more structured quantum walk features begin to appear. This is also reflected in the variance plots, which become progressively less close to the LBM variance as θ is increased. The smallest value, $\theta = 0.05$, is a slightly weak spreading case, while $\theta = 0.1$ gives an especially clean match to the LBM variance. For larger θ, the variance growth remains broadly diffusive, but it becomes less linear and less closely matched to the classical LBM comparison.

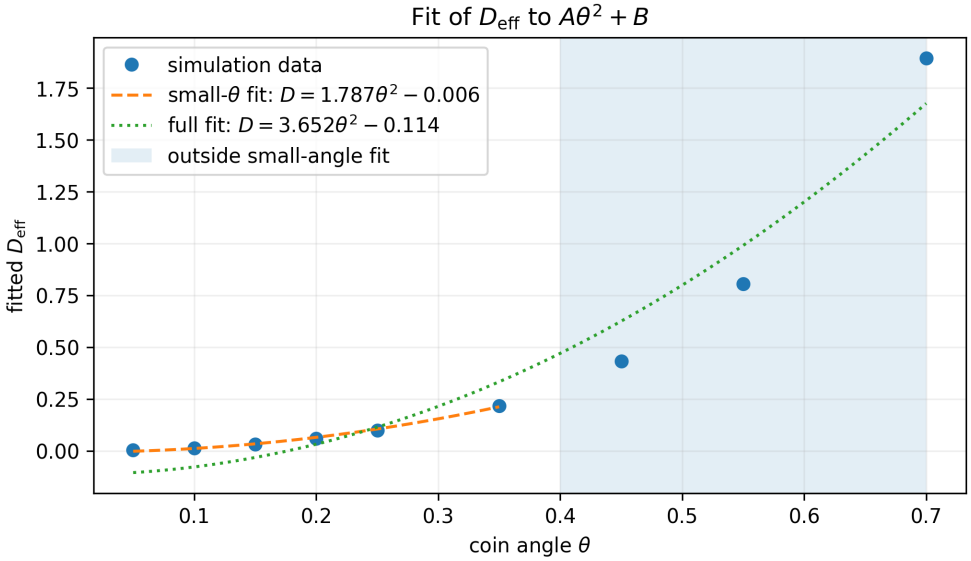

Figure 5.12: Quadratic fit for the dependence of the effective diffusion coefficient on the parity echo coin angle for the advection diffusion model. The fitted diffusion coefficients are compared with curves of the form $D_{\text{eff}}(\theta) = A\theta^2 + B$. The small to moderate angle fit supports the CE prediction $D_{\text{eff}} \propto \theta^2$, while the larger angle data show deviations from the leading perturbative scaling.

For the advection-diffusion θ sweep, the fitted effective diffusion coefficients were compared with a quadratic scaling model of the form $D_{\text{eff}} = A\theta^2 + B$. The parity echo model was simulation over a range of θ values, with a value of $\gamma = 0.3$. Fitting

only the small angle regime $\theta \leq 0.4$ gives $A = 1.78748403$ and $B = -0.00598039$, with $R^2 = 0.99578536$ and RMSE $= 4.71 \times 10^{-3}$. This high R^2 value shows that the small to moderate angle data agree very well with the CE prediction $D_{\text{eff}} \propto \theta^2$. When the full range of θ values is included, the fit becomes $D_{\text{eff}} = 3.65212670\theta^2 - 0.11365300$, with $R^2 = 0.94870104$ and RMSE $= 1.32 \times 10^{-1}$. Although the full range fit still captures the overall increasing trend, the lower R^2 and larger RMSE show that again, the larger θ data deviates from the leading perturbative scaling. This is expected, since the CE prediction is derived from a small angle expansion; at larger values of θ the simple quadratic trend is no longer expected to hold exactly.

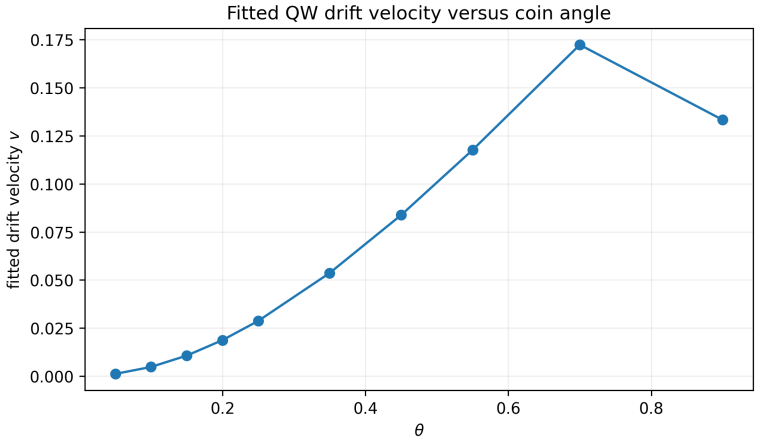

Figure 5.13: Fitted drift velocity for the biased parity echo advection-diffusion simulations as a function of the coin angle θ. The decrease at the largest value of θ should be interpreted with care, since the packet spreads strongly and finite-size or boundary effects can contaminate the late-time mean-position fit.

The fitted drift velocity and variance linearity diagnostics for the θ sweep are shown in Figs. 5.13 and 5.14. The drift velocity is again calculated in the same manner as the previous subsection. Over the small to intermediate angle range, the fitted velocity increases with θ. This is expected, since increasing θ strengthens the coherent mixing in the parity echo step, allowing the initial right moving coin bias to generate a stronger net drift.

The variance linearity plot shows that the variance growth is very close to linear for small and intermediate values of θ, with R^2 values close to one. This indicates that, in this regime, the biased parity echo walk has both an approximately linear mean

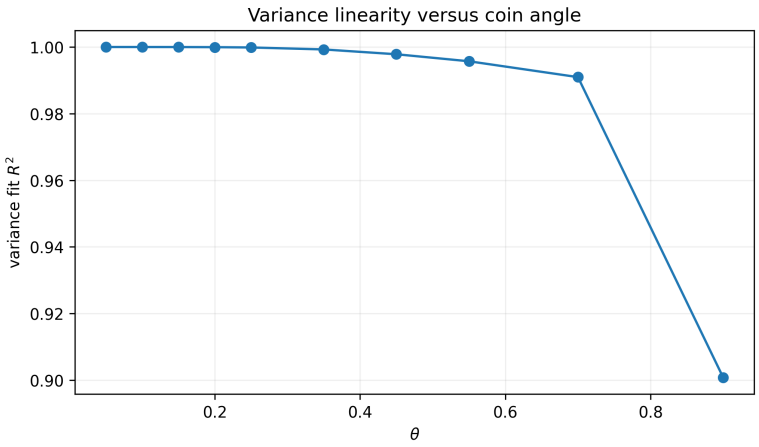

Figure 5.14: Variance linearity diagnostic for the biased parity echo advection-diffusion θ sweep.

displacement and an approximately linear variance growth, as expected for advection-diffusion. However, at larger θ the variance fit becomes less clean, and the fitted drift velocity also stops following the simple increasing trend. This should be interpreted carefully. For the largest coin angles the packet spreads much more strongly, and the late time dynamics can be affected by finite lattice size and boundary effects. In particular, if probability reaches the edge of the lattice and reappears or is distorted by the boundary treatment, the fitted mean position can be corrupted. Thus the large θ regime is not the cleanest regime for extracting an advection-diffusion limit, even though it produces large spreading and large fitted transport coefficients.

Overall, the advection-diffusion sweeps show the same basic mechanism as the diffusion-only case, but with an additional drift component introduced through the biased initial coin state. The dephasing strength controls the lifetime of the coherences and therefore the effective diffusion coefficient, while the coin angle controls the strength of the coherent mixing and hence the amount of spreading. The biased initial state then breaks the left/right symmetry of the density profile, shifting the centre of mass and producing an effective advection velocity. In the small to intermediate parameter regime, the biased parity echo walk therefore has the two main macroscopic signatures of advection-diffusion: approximately linear mean displacement and approximately linear variance growth. This is important because the analytical and numerical results

are consistent with one another. The Chapman-Enskog expansion predicts that diffusion arises from a balance between coherent generation and dephasing of off-diagonal coherences, while advection can be introduced by breaking the symmetry of the initial coin state. The simulations show precisely these features: the density profiles can be matched to D1Q3 LBM profiles, the fitted diffusion coefficient follows the expected parameter dependence, and the biased case produces a measurable drift. Thus the CE expansion is not just a formal calculation; it provides useful guidance for choosing the walk structure and parameters needed to reproduce macroscopic transport behaviour.

However, reproducing advection-diffusion at the level of density profiles is only one part of the problem. To interpret the model as a quantum algorithm, it is also necessary to ask how the initial density is encoded, how many quantum resources are required to implement the walk, and what information can be efficiently extracted at the end. The following section therefore moves from the continuum and numerical validation of the parity echo model to its computational interpretation.

5.6 Dependence of the Effective Diffusion Coefficient on Initial Conditions

I further investigate how the initial coin state influences the effective diffusion coefficient, D_{eff}, for both diffusion and advection-diffusion. Figure 5.15 shows that changing the initial coin state produces modest but noticeable variations in D_{eff} for the unbiased diffusion case. This differs from the classical LBM, where the macroscopic diffusion coefficient is determined by the relaxation time and lattice parameters, independently of the initial distribution. In the quantum model, the initial coin populations and coherences can influence the measured transport behaviour.

For the advection-diffusion case, Fig. **??** shows considerably smaller variations in D_{eff}. This is partly because the initial states considered are restricted to strongly biased configurations, limiting the range of coin populations explored. Nevertheless, both cases show that D_{eff} is primarily influenced by the coin angle θ and dephasing strength γ, with the initial state providing an additional dependence. Note that the

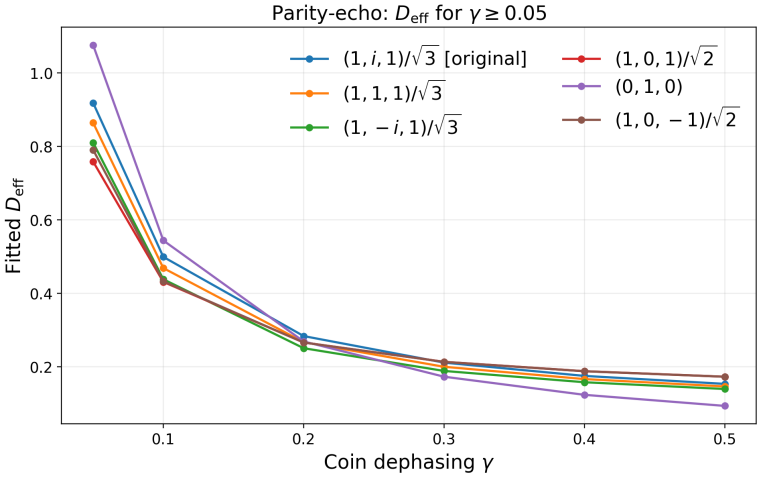

Figure 5.15: Effective diffusion coefficient D_{eff} as a function of the dephasing strength γ for different unbiased initial coin states in the parity echo model. The diffusion coefficient is extracted from the linear growth of the spatial variance, with $\theta = 0.25$ and $\sigma_0 = 10$.

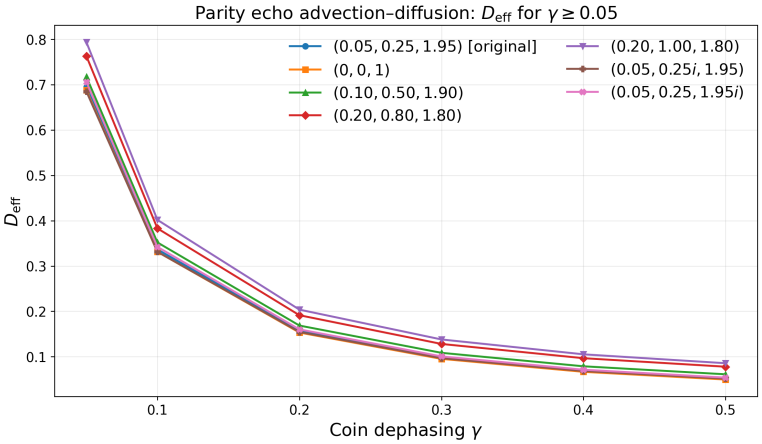

Figure 5.16: Effective diffusion coefficient D_{eff} as a function of the dephasing strength γ for different strongly biased initial coin states in the parity-echo model. The diffusion coefficient is extracted from the linear growth of the spatial variance, with $\theta = 0.25$ and $\sigma_0 = 10$.

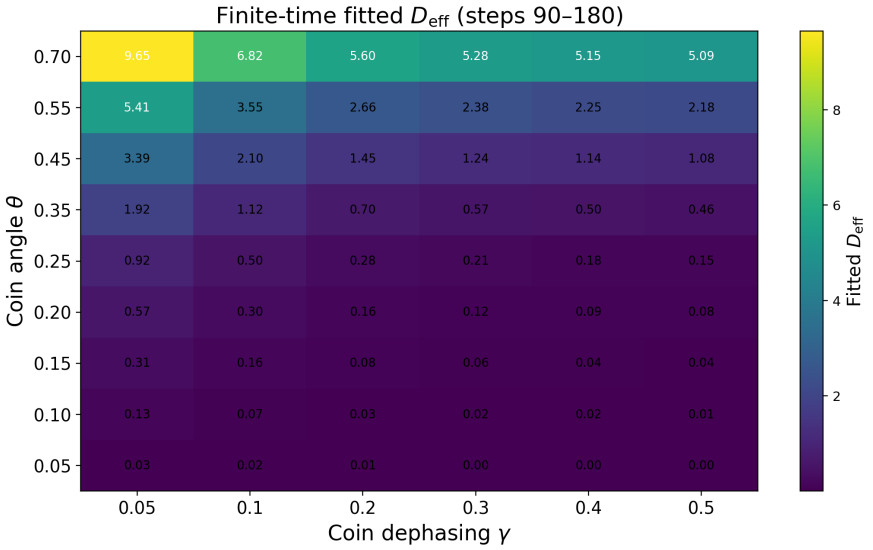

Figure 5.17: Parameter sweep of the effective diffusion coefficient D_{eff} in the parity echo model as a function of the coin angle θ and dephasing strength γ, for the initial coin state $\chi = (1, i, 1)^T/\sqrt{3}$. The effective diffusion coefficient is extracted from a linear fit to the variance growth over the chosen fitting interval.

$\gamma = 0$ points have been excluded from both plots, as the absence of dephasing leads to predominantly ballistic spreading, for which a diffusion coefficient extracted from a linear variance fit is not physically meaningful.

Figure 5.17 shows the dependence of the effective diffusion coefficient on the coin angle θ and dephasing strength γ. Within the parameter ranges investigated, chosen to explore approximately diffusive behaviour, the fitted coefficients span $2.7 \times 10^{-3} D_{\text{eff}} 9.65$. The coefficient generally increases with θ and decreases with γ, demonstrating the tunability of the parity echo model. However, the larger values obtained at weak dephasing may retain significant coherent contributions and therefore do not necessarily represent purely diffusive transport. Unlike classical LBM, the parity echo model exhibits some sensitivity to initial state preparation, but its effective diffusion coefficient remains predominantly tunable through the coin angle and dephasing strength.

5.7 Encoding and Computational Cost

It is important to comment on the computational cost of the parity echo model, and in particular on the role of state preparation and readout. A classical D1Q3 lattice Boltzmann simulation of advection-diffusion on a one dimensional lattice with N sites has memory cost $\mathcal{O}(N)$ and computational cost $\mathcal{O}(NT)$ for T time steps. This is because only three populations, f_-, f_0, f_+, are stored at each lattice site, and each timestep consists of local collision and streaming operations. Classical LBM is therefore already very efficient, with linear scaling in the number of lattice sites.

The parity echo walk has a different cost structure on a quantum computer. A lattice with N spatial sites can be encoded using $\log_2 N$ position qubits, together with a small coin register representing the three coin states $|L, S, R\rangle$. Since three coin states require $\lceil \log_2 3 \rceil = 2$ qubits, an an ancilla qubit for the decoherence, the total Hilbert space register requires $\lceil \log_2 N \rceil + 3$ qubits. In this sense, the quantum representation of the lattice is exponentially more compact than storing an N site classical probability distribution.

One parity echo timestep is also structurally local. The coin operations act only on the small coin register, while the conditional shift updates the position register depending on the coin state. At the circuit level, a parity echo step contains two coin operations and two conditional shifts. Compared with a standard lazy Grover step $U_G = SC_G$, the parity echo step is therefore only a constant factor larger. The echo construction does not change the asymptotic dependence on the number of lattice sites. It adds a small fixed overhead while improving the continuum ordering of the model.

However, this compact quantum encoding should not be confused with an automatic computational advantage over classical LBM. The total cost also includes preparing the initial quantum state, implementing the dephasing channel, and measuring the desired output. If an arbitrary classical density profile must be loaded into amplitudes, the state preparation cost can be large and may remove any advantage from the $\log N$ spatial encoding. Similarly, if the full probability density must be reconstructed at every time step, many measurements are required, and the readout cost can again become

comparable to storing or computing the classical distribution directly. The measurement cost must also be considered. If the goal is to reconstruct the full probability density, then many measurements are required, and this can reduce the practical benefit of the quantum representation. However, if only macroscopic quantities such as the mean position, variance, or low order moments are needed, then the readout cost may be much lower. This is particularly relevant for transport simulations, where the quantities of interest are often effective drift velocities, diffusion coefficients, or coarse grained density profiles rather than the full microscopic state.

For the advection-diffusion examples considered here, the state preparation assumption is more reasonable. Transport problems often use simple smooth initial profiles, such as Gaussian wavepackets, rather than arbitrary data. The simulations in this chapter use Gaussian spatial states, together with simple three component coin states. These are much easier to specify and, in principle, to prepare than a completely arbitrary density field. The biased advection-diffusion case also only requires changing the initial coin amplitudes, so the extra cost of introducing drift is negligible.

Therefore, the parity echo model should not be interpreted as giving an automatic quantum speedup for classical diffusion or advection-diffusion. As with all quantum algorithms, its potential advantage depends on the full algorithmic pipeline. Efficient preparation of simple initial states, efficient implementation of the local walk and dephasing channel, and measurement of a restricted set of observables. Within this setting, the model is still useful because it shows that a small, structured open quantum walk can reproduce diffusive and advection-diffusion transport using coherent dynamics together with controlled decoherence. The main cost caveat is that this only remains efficient if the initial state and the required output observables are themselves efficiently accessible.

5.8 Discussion

The parity echo model developed in this chapter addresses the limitations identified in the previous chapter. The Grover and twistless models demonstrated that decoherence, rest states, plastic scaling, and stroboscopic cancellation are useful ingredients

for constructing diffusive quantum walks, but do not automatically produce the desired CE ordering. In these models, coherent transport and dephasing often appeared at the same first non-trivial order, preventing dephasing from acting independently as an equilibrium selection mechanism. The parity echo construction was therefore designed to suppress the leading coherent contribution through cancellation within the unitary step itself.

The central improvement is the vanishing first order unitary contribution. Unlike a generic lazy quantum walk, where the conditional shift introduces leading directional transport, the two halves of the parity echo step are related through a parity swap and opposite sign rest state mixing. This partially reverses the directional dynamics, suppressing the leading ballistic contribution and delaying coherent transport to higher order. Physically, the walk retains its quantum unitary structure while exhibiting reduced spreading, providing a more suitable starting point for decoherence assisted diffusion. This is supported numerically by the smaller variance and standard deviation compared with the standard lazy Grover walk, even without dephasing. Although reduced spreading alone does not establish diffusion, it demonstrates the effectiveness of the cancellation mechanism.

The symmetric rest state coupling also improves upon the three state twistless model, where rest state mixing entered at the same order as transport and produced noticeable numerical drift. In contrast, the parity reflected coupling preserves the left/right balance more effectively in the unbiased case, avoiding an unwanted advective contribution. This is important when introducing a D1Q3 rest state to model diffusion without directional bias.

The resulting CE hierarchy provides the main analytical advantage of the parity echo construction. At leading order, dephasing selects the coin diagonal sector as the slow manifold, corresponding to the nullspace of the relaxation operator. This sector emerges naturally rather than being imposed as an equilibrium condition, giving coin dephasing a role analogous to local relaxation in LBM. At the next order, coherent dynamics generates off-diagonal density matrix corrections, which act as the non-equilibrium variables and are subsequently damped by dephasing. Higher order

terms then feed these corrections back into the population sector, producing second derivative contributions associated with diffusion. The resulting relaxation-correction-diffusion ordering is precisely what was missing from the earlier constructions.

This mechanism parallels the classical CE expansion, where non-equilibrium corrections are determined at lower order and substituted into higher order moment equations to obtain diffusion. Here, the corrections are quantum coherences rather than classical population deviations. The overdamped closure assumes that these coherences relax sufficiently rapidly to become slaved to the slow diagonal populations, allowing an effective diffusive response to be obtained. However, this is an additional asymptotic approximation rather than an exact consequence of the discrete walk. The natural continuum description remains a coupled density matrix hierarchy, with a scalar diffusion equation arising only under sufficient timescale separation. The closure is therefore expected to be most appropriate for smooth, broad initial states dominated by long-wavelength modes, and for decoherence strengths that damp coherences without excessively suppressing transport.

The numerical dephasing sweep supports this interpretation. Increasing γ suppresses oscillatory features and smooths the probability distributions, but stronger dephasing does not necessarily improve diffusion. Weak dephasing retains significant coherent behaviour, while excessive dephasing suppresses the coherences responsible for transport. The most suitable intermediate regime exhibits approximately linear variance growth and good agreement with matched D1Q3 LBM diffusion profiles. Furthermore, the overdamped closure predicts an approximate inverse dependence of the effective diffusion coefficient on dephasing strength, $D_{\text{eff}} \propto 1/\gamma$, supported numerically by the approximately linear relationship between D_{eff} and $1/\gamma$. This agreement provides quantitative evidence for the proposed coherence-mediated diffusion mechanism. The $\gamma = 0$ case lies outside this interpretation, since the overdamped closure requires non-zero dephasing and coherent spreading does not constitute the same diffusive mechanism.

The coin angle dependence provides further validation. In the perturbative regime, the closure predicts $D_{\text{eff}} \propto \theta^2$, consistent with the numerical results at small and mod-

erate angles. This reflects the distinct roles of θ, which controls coherent generation of non-equilibrium coherences, and γ, which determines their relaxation. At larger angles, higher order coherent effects become significant and the distributions deviate from diffusive behaviour, identifying the limitations of the small angle continuum approximation. Together, these results demonstrate that effective diffusion is governed by the balance between coherent generation and dissipative damping, rather than decoherence alone.

The advection-diffusion simulations demonstrate that the same parity echo unitary can also support directed transport through a biased initial coin state. Unlike the unbiased diffusion case, these simulations exhibit approximately linear mean displacement alongside approximately linear variance growth, reproducing the characteristic moments of advection-diffusion without modifying the underlying unitary. Comparisons with matched D1Q3 LBM simulations further assess both the transport moments and probability density profiles. By extracting the effective diffusion coefficient from the quantum walk variance and using it in the LBM comparison, the agreement demonstrates compatibility with classical diffusive transport over the simulated time intervals, rather than merely producing smooth distributions.

Overall, the parity echo model is the most successful construction developed in this thesis, with numerical behaviour supporting the hierarchy predicted analytically. Nevertheless, important limitations remain. The scalar diffusive description relies on an overdamped closure valid only within an appropriate long-wavelength regime, while the model remains linear and does not reproduce the nonlinear collision dynamics of full LBM fluid solvers. Its behaviour also depends on carefully engineered symmetry and cancellation, and it should therefore not be interpreted as a complete quantum lattice Boltzmann algorithm.

More broadly, the results establish several design principles for diffusive quantum walks: symmetry can suppress unwanted coherent orders, a rest state provides a useful kinetic velocity structure, parity preservation helps prevent unintended drift, and controlled decoherence selects and stabilises the slow sector. Crucially, these ingredients must be organised to produce the appropriate continuum hierarchy rather than sim-

ply combined within a quantum walk. Decoherence is therefore used constructively as a relaxation mechanism, producing transport that is neither purely coherent ballistic spreading nor simply classical random walk diffusion. Although the model does not yet provide a complete quantum fluid simulation algorithm, it demonstrates how carefully engineered open quantum dynamics can reproduce features of classical kinetic diffusion and advection–diffusion, providing a foundation for future quantum kinetic models.

Chapter 6

Discussion and Conclusion

6.1 Summary of Results

6.1.1 Continuum limits of lazy open quantum walks

Chapter 3 established the continuum limit foundations for the subsequent models. Continuous spacetime limits were derived for lazy discrete time quantum walks subject to coin or spatial decoherence, extending the usual unitary continuum limit picture to open three state walks. Although the coherent transport contribution originates from the same unitary dynamics, the dissipative terms depend on the noise channel: coin dephasing suppresses coherences between internal coin states while preserving spatial coherences, whereas spatial dephasing suppresses coherences between lattice sites and more directly inhibits interference and quantum transport. These differences were also apparent in the numerical simulations, which showed distinct routes towards more classical spreading.

An important result was the numerical validation of the continuum equations against the corresponding discrete walks in the appropriate continuum regime. These limits are valuable independently of the LBM comparison: they identify which continuous equations the discrete dynamics approximate and demonstrate that different forms of decoherence are not interchangeable. They also establish a methodological principle used throughout the thesis: continuum analysis can validate a chosen quantum walk and, more importantly, guide the design of a walk with a desired macroscopic response,

even when the continuum equations cannot be solved analytically in closed form.

6.1.2 Grover and twistless quantum walks

Chapter 4 developed a quantum CE style expansion to examine whether open quantum walks reproduce the ordering of equilibrium, transport and diffusion familiar from lattice Boltzmann methods. The initial two state and lazy three state Grover walks provided D1Q2 and D1Q3 inspired test cases, respectively. Introducing a rest state was beneficial in this comparison: the operator $J_z^2 = \mathrm{diag}(1, 0, 1)$ distinguishes moving from stationary populations, giving a useful second moment structure, while the rest population contributes to the density without directly carrying flux. Numerically, the lazy three-state Grover walk exhibited density profiles and variance growth closer to the matched LBM diffusion benchmark than the corresponding two state walk. Nevertheless, in both Grover models coherent transport, coin mixing and dephasing entered the first non-trivial CE equation together. Although higher order k^2 terms indicated diffusion-like spreading, the hierarchy did not isolate a diagonal local-equilibrium sector before transport or yield a clean scalar diffusion closure.

Plastic scaling and stroboscopic cancellation were then investigated as ways to reorganise this hierarchy. With a timestep of order q^2, an uncancelled order q unitary term becomes singular after division by the timestep, giving a q^{-1} contribution. Thus rescaling alone is insufficient: the walk must satisfy a regularity condition that removes these lower order terms. The twistless construction achieved this cancellation and produced a substantially cleaner expansion. In the two state case, the first non-vanishing unitary contribution took the form of a $k\sigma_y$ Hamiltonian term, followed by a $k^2\sigma_x$ correction. This showed how engineered cancellation can suppress the earliest ballistic contributions while retaining coherent dynamics and higher order spreading terms.

Extending the twistless walk to three states introduced a rest component and associated mixing terms, but their coupling also modified the leading transport and produced unwanted directional bias. The drift seen in the numerical simulations was consistent with the CE analysis, demonstrating that the expansion captured a feature of the discrete dynamics. Such a bias is undesirable for pure diffusion, although it

could potentially be exploited for advection-diffusion. More generally, both twistless models still placed coherent transport and weakly scaled dephasing at the same first non-trivial order. They improved the regularity of the hierarchy without establishing the desired separation between initial relaxation and subsequent transport. The Grover and twistless constructions therefore served as useful examples of both the benefits and the limitations of rest states, scaling and unitary cancellation.

6.1.3 The parity echo model

Chapter 5 introduced the parity echo construction to address these limitations. By relating the two parts of a walk step through a parity reflected sequence with symmetric rest state coupling, the model cancels the leading directional unitary contribution and reduces the rapid ballistic spreading of a standard lazy walk. It also avoids the unwanted left/right bias seen in the three-state twistless construction when studying unbiased diffusion. The resulting CE hierarchy came closest to the intended LBM-inspired sequence: the leading dephasing operator selects the coin diagonal slow sector; coherent evolution subsequently generates off-diagonal, non-equilibrium density matrix corrections; and higher orders couple these corrections back into the coin populations through k^2 dependent terms. Under an overdamped coherence assumption, the off-diagonal corrections can be approximately slaved to the slow sector, giving an effective diffusive response. In this quantum analogue of a kinetic closure, coherences play the role of rapidly relaxing non-equilibrium variables rather than classical population corrections.

The simulations supported the proposed mechanism. In the unitary limit, parity echo reduced the variance and spreading relative to the standard lazy Grover walk; this suppression alone is not diffusion, but it provides a better starting point for obtaining it through controlled decoherence. Increasing coin dephasing smoothed the probability distributions, with the most useful diffusive behaviour occurring at intermediate strengths: weak dephasing retained substantial interference, whereas strong dephasing suppressed the coherences needed to mediate transport. Over the appropriate regime, the variance was approximately linear in time and the density profiles agreed well with

D1Q3 LBM diffusion profiles matched using the measured diffusion coefficient. The fitted effective coefficient also showed the dependencies predicted by the overdamped analysis, approximately $D_{\text{eff}} \propto 1/\gamma$ and $D_{\text{eff}} \propto \theta^2$ for suitable dephasing strengths and small-to-moderate coin angles. The agreement between these analytical trends and numerical fits supports the interpretation of diffusion as a balance between coherence generation and damping; deviations at larger angles are consistent with the breakdown of the perturbative approximation.

Finally, a biased initial coin state generated a moving distribution without changing the parity echo unitary. Approximately linear growth of both the mean position and variance provided numerical signatures of advection-diffusion, whereas unbiased initial states isolated the diffusive regime. Comparisons with matched D1Q3 LBM profiles examined not only the variance used to set D_{eff}, but also the shape of the distribution over the simulated time windows; the initial packet width and coin-state bias were further relevant controls. Of the constructions considered, parity echo therefore gave the strongest correspondence between CE ordering and diffusive numerical behaviour. Its interpretation remains conditional, however: the overdamped closure requires an appropriate long wavelength and timescale separation regime, and a k^2 dependent equation for coin populations should not, by itself, be taken as proof of a closed scalar diffusion equation for the traced spatial density.

6.2 Overarching Discussion and Conclusions

The aim of this thesis was not to reproduce the full lattice Boltzmann algorithm step by step, but to determine whether decohered discrete time quantum walks could realise the underlying kinetic mechanisms of diffusion and advection-diffusion. Directly mapping quantum dynamics onto classical LBM is difficult, in part because the usual fluid collision equilibrium depends nonlinearly on macroscopic fields. Existing approaches may therefore linearise the collision step or use measurement and reinitialisation, with consequent restrictions on the dynamics or coherence. Quantum walks provide a natural high-dimensional, wave-like starting point for such kinetic models, but their unitary dynamics are typically ballistic rather than diffusive. By focusing first on linear trans-

port and its diffusive limit, this work isolated the ingredients required for a quantum kinetic analogue of relaxation, streaming and macroscopic transport without assuming that a complete nonlinear fluid solver had already been obtained.

The first principle is that the location of decoherence matters. In LBM, streaming transports populations while local collision relaxes them towards equilibrium. Keeping the position shift unitary and applying dephasing to the internal coin preserves this separation more closely than spatial dephasing. Coin dephasing selects a diagonal fixed sector and damps the off-diagonal coherences generated by the walk, allowing it to be interpreted as a local relaxation mechanism. This does not mean that any decohered walk automatically becomes an LBM model: the equilibrium sector must be selected at a suitable order of the multiscale expansion if the analogy is to provide a useful kinetic closure.

The second principle concerns velocity structure. A stationary coin state can make a lazy walk closer to the D1Q3 setting than a two state walk is to D1Q2, providing a population that contributes to density without directly carrying flux and improving the available second-moment structure. The Grover results demonstrated a numerical benefit from this addition, but the twistless results showed why a rest state is not sufficient on its own. Its coupling to moving states must respect the symmetries of the desired macroscopic equation; otherwise, a nominally diffusive model may acquire drift. The objective is therefore not simply to increase the number of coin states, but to control how they are mixed.

The third principle is that the ordering of the CE hierarchy is a design constraint. A density matrix equation containing dephasing and k^2 terms is not automatically a macroscopic diffusion equation. In the Grover hierarchy, coherent mixing and relaxation remained coupled at leading order. Twistless cancellation removed problematic lower-order plastic terms, but still left relaxation and transport at the same first non-trivial order. Parity echo improved this separation by suppressing the early unitary contributions, so that the diagonal sector could be selected before higher-order coherence-mediated transport appeared. In this way, the continuum expansion became a constructive tool for choosing the microscopic sequence, rather than solely a way to

describe it afterwards.

These observations distinguish three levels of success that should not be conflated: observing numerically smooth, approximately diffusive spreading; deriving a consistent density matrix kinetic hierarchy with relaxation and k^2 contributions; and proving that a specified macroscopic density obeys a closed scalar transport PDE. The present models realise these objectives to differing degrees. In particular, the parity echo population-level closure supplies a strong mechanism and numerical evidence for diffusion-like transport, while the precise projection onto an autonomous traced-density equation remains an important mathematical qualification. This distinction does not undermine the continuum or numerical results; it defines their scope and identifies what would be required for a stronger macroscopic claim.

A further conclusion is that the mathematical hierarchy and numerical results reinforce one another. The early coupling of transport and dephasing in the simpler walks helps explain their less controlled spreading; the drift in the three-state twistless walk is visible both analytically and numerically; and the improved parity echo ordering is accompanied by smoother profiles, more nearly linear variance growth and parameter trends consistent with coherence-mediated diffusion. The comparisons with LBM are useful here because visual smoothness alone is not a sufficient benchmark. Matching an effective coefficient from variance tests the moments, while comparison of the remaining profile shape gives additional evidence about the transport process, without constituting an independent prediction of the fitted coefficient.

The work also illustrates the limits of what the present models establish computationally. All of the constructions considered are linear open quantum systems, and none provides the nonlinear collision dynamics required for full Navier Stokes simulation. Nor has a computational advantage over classical LBM or CFD been demonstrated. In many fluid applications the desired output is a high resolution field at numerous timesteps, whereas extracting a full quantum spatial distribution requires repeated measurements; intermediate measurements can also disrupt the coherence used in the dynamics. Input preparation, readout, error tolerance and classical pre and postprocessing must therefore be treated as part of the algorithm rather than as incidental

implementation details.

For these reasons, a plausible longer term role for quantum transport methods may lie in targeted observables or hybrid workflows rather than direct replacement of classical CFD. Moments, correlations, final-time quantities and selected local observables may be more natural quantum outputs than complete flow fields. Classical HPC methods also continue to improve, including through tensor network, probabilistic, machine learning and AI-based approaches, so comparisons must be made against strong contemporary classical alternatives. Even without quantum speed-up, the development of these models can provide new mathematical perspectives on open transport and potentially useful quantum inspired algorithms.

The central contribution is therefore a set of linked design principles, not a finished quantum lattice Boltzmann solver. Controlled coin dephasing provides local relaxation; rest states can improve kinetic moment structure; symmetry controls drift and, in higher dimensions, isotropy; and carefully designed stroboscopic or echo cancellation determines whether equilibrium, transport and diffusion appear on compatible scales. None of these elements works reliably in isolation. The broader proposal is to compare and engineer continuum hierarchies, rather than demand microscopic identity between classical and quantum updates. Grover, twistless and parity echo walks should be viewed as representative constructions that reveal these principles, not as an exhaustive search over possible quantum walks.

6.3 Future Work

One direction is to develop more flexible dissipative channels. Projective coin dephasing was chosen because it is simple, completely positive and trace preserving, and has a clear interpretation as local relaxation. Biased, partial or subspace dephasing, or channels that relax towards non-uniform coin equilibria, could permit more general macroscopic dynamics. Any such choice would have to retain complete positivity and trace preservation while avoiding spurious drift or the loss of symmetries needed by the target transport equation. Establishing the resulting slow manifold and its connection to a closed spatial-density equation would also provide a more rigorous test of the

proposed kinetic mechanism.

A second direction is to extend the constructions to two or more spatial dimensions. A D2Q9 walk, containing rest, axial and diagonal states, would be closer to the velocity sets used in practical lattice Boltzmann calculations. Unlike one dimensional parity, a higher dimensional design must also control rotational isotropy and lattice artefacts. A central question is whether suitable coin couplings, dephasing and echo symmetries can produce an isotropic $k_x^2 + k_y^2$ response without introducing unwanted directional transport.

A third possibility is to introduce nonlinear macroscopic dynamics. The present linear models naturally address diffusion and linear advection-diffusion, but Burgers', Euler and Navier Stokes equations require additional structure. Possible routes include hydrodynamic or Madelung-type reformulations of quantum walk density and phase variables [Hatifi et al., 2019], or carefully controlled measurement, feedback and hybrid quantum-classical coupling. These ideas would need to be examined for whether they generate genuine nonlinear closure without destroying the quantum dynamics that motivates the approach.

A fourth direction concerns practical encoding, measurement and benchmarking. The preparation of relevant spatial distributions, efficient evolution, and extraction of useful answers must be assessed together. Full state tomography or repeated reconstruction of entire spatial profiles is unlikely to be the most attractive readout strategy. Global moments, hitting probabilities, correlations, selected local observables or final time quantities may be better candidates. Benchmarking should also quantify the cost of state preparation, noise, measurements and classical processing, not only compare the resulting density profiles.

Finally, these walks could be explored as specialised components of hybrid HPC-quantum algorithms. Potential tasks include estimating transport observables, sampling kinetic distributions, assisting with linear subproblems or accessing spectral information within a larger classical calculation. Determining whether any such component is genuinely useful will require comparison with efficient classical methods and a clear account of the full computational workflow. More generally, the thesis leaves open the

search for other coin, shift and decoherence constructions that improve on parity echo or produce a more rigorous scalar macroscopic closure.

Overall, the results support the use of controlled decoherence as a constructive ingredient rather than only as an obstacle to quantum transport. They show both why generic quantum walks do not automatically reproduce kinetic diffusion and how their microscopic dynamics may be engineered to approach it. The value of this framework lies in clarifying the connection between discrete open quantum dynamics and continuum transport, while identifying the further analytical and computational steps needed before practical quantum fluid simulation can be claimed.

Appendix A

Appendix

A.1 Gell Mann Matrices

Explicitly, the Gell Mann matrices are given by Wilson [2024]:

$$\lambda_1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \lambda_2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \lambda_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$\lambda_4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \lambda_5 = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \quad \lambda_6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\lambda_7 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \quad \lambda_8 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}$$

A.2 Numerical Simulations of the Continuum Equation

To compare the discrete lazy quantum walk with its continuum approximation, we solved the continuum density matrix master equation (3.35) numerically on a finite one dimensional spatial grid.

The spatial derivative operator D_x was constructed using second order centred

differences in the bulk,

$$(\partial_x f)_i \approx \frac{f_{i+1} - f_{i-1}}{2\Delta x}, \tag{A.1}$$

and first order one sided differences at the two boundaries,

$$(\partial_x f)_0 \approx \frac{f_1 - f_0}{\Delta x}, \qquad (\partial_x f)_{N_x-1} \approx \frac{f_{N_x-1} - f_{N_x-2}}{\Delta x}. \tag{A.2}$$

The initial state in the continuum simulation was chosen to match the discrete initial condition: a Gaussian wavepacket in position multiplied by the Fourier coin state,

$$|\chi_F\rangle = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ \omega \\ \omega^2 \end{pmatrix}, \qquad \omega = e^{2\pi i/3}, \tag{A.3}$$

with spatial profile

$$\psi_0(x) \propto \exp\left(-\frac{x^2}{2\sigma_0^2}\right). \tag{A.4}$$

The continuum state was normalised such that

$$\int \|\psi(x,0)\|^2 \, dx = 1. \tag{A.5}$$

For each value of the discrete dephasing probability γ_{disc}, the corresponding continuum dephasing rate was chosen as

$$\gamma_{\text{cont}} = -\frac{\ln(1-\gamma_{\text{disc}})}{\Delta t_{\text{step}}}, \tag{A.6}$$

so that the continuum Lindblad damping reproduces the same effective dephasing per discrete time step.

Equation (3.35) was integrated in time using the classical explicit fourth order Runge Kutta (RK4) method. Writing

$$\dot{\rho} = F(\rho), \tag{A.7}$$

the update from time t_n to $t_{n+1} = t_n + \Delta t_{\text{int}}$ is

$$k_1 = F(\rho_n), \tag{A.8}$$

$$k_2 = F\left(\rho_n + \frac{\Delta t_{\text{int}}}{2} k_1\right), \tag{A.9}$$

$$k_3 = F\left(\rho_n + \frac{\Delta t_{\text{int}}}{2} k_2\right), \tag{A.10}$$

$$k_4 = F(\rho_n + \Delta t_{\text{int}} k_3), \tag{A.11}$$

$$\rho_{n+1} = \rho_n + \frac{\Delta t_{\text{int}}}{6} (k_1 + 2k_2 + 2k_3 + k_4). \tag{A.12}$$

After each time step, Hermiticity was restored numerically by symmetrisation,

$$\rho \leftarrow \frac{1}{2}\left(\rho + \rho^\dagger\right). \tag{A.13}$$

The final probability density was obtained from the diagonal position blocks of the density matrix,

$$P(x_i, t) = \text{Tr}_{\text{coin}}[\rho_{x_i x_i}(t)], \tag{A.14}$$

and was then normalised so that

$$\int P(x, t)\, dx = 1. \tag{A.15}$$

A.3 Full lattice Boltzmann CE expansion to recover the Navier Stokes equation

This section demonstrates how the Navier Stokes equation for fluid flow may be derived from the lattice Boltzmann equation. We firstly consider the BGK lattice Boltzmann equation

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t,\, t + \Delta t) - f_i(\mathbf{x}, t) = -\frac{\Delta t}{\tau} [f_i(\mathbf{x}, t) - f_i^{\text{eq}}(\mathbf{x}, t)]. \tag{A.16}$$

Appendix A. Appendix

Table A.1: Numerical parameters used in the discrete quantum walk and continuum RK4 simulation.

Parameter	Meaning	Value
N_{steps}	Number of discrete time steps	100
P	Number of spatial lattice points	$2N_{\text{steps}} + 1 = 201$
x_i	Spatial grid	$x_i = -100, \ldots, 100$
Δx	Spatial grid spacing	1
θ	Coin / Hamiltonian parameter	$\pi/2$
σ_0	Initial Gaussian width	5
Initial coin state	Fourier coin state	$\frac{1}{\sqrt{3}}(1, \omega, \omega^2)^T$, $\omega = e^{2\pi i/3}$
γ_{disc}	Discrete dephasing values	[0, 0.01, 0.05, 0.10, 0.20, 0.40]
γ_{cont}	Continuum dephasing rate	$-\ln(1 - \gamma_{\text{disc}})/\Delta t_{\text{step}}$
Δt_{int}	RK4 internal timestep	0.01
N_x	Continuum grid size	201
Time integration	Numerical method	Fourth order Runge Kutta (RK4)
Hermiticity correction	Post-step symmetrisation	$\rho \leftarrow (\rho + \rho^\dagger)/2$
Probability extraction	Final density	$P(x_i) = \text{Tr}_{\text{coin}}[\rho_{x_i x_i}]$
Normalisation	Continuum probability normalisation	$\int P(x,t)\, dx = 1$

The hydrodynamic fields are defined by the lowest moments

$$\rho_{LB} = \sum_i f_i, \qquad \rho_{LB}\mathbf{u} = \sum_i \mathbf{c}_i f_i. \tag{A.17}$$

To derive the corresponding continuum equations, we perform a Taylor expansion of the left-hand side about $(\mathbf{x}, t)$,

$$(\partial_t + c_{i\alpha}\partial_\alpha)\, f_i + \frac{\Delta t}{2}\,(\partial_t + c_{i\alpha}\partial_\alpha)^2\, f_i = -\frac{1}{\tau}\,(f_i - f_i^{\text{eq}}) + \mathcal{O}(\Delta t^2). \tag{A.18}$$

We now introduce the Chapman-Enskog multiscale expansion. Time and space derivatives are separated into slow scales,

$$\partial_t = \varepsilon\,\partial_{t_1} + \varepsilon^2\,\partial_{t_2} + \cdots, \qquad \partial_\alpha = \varepsilon\,\partial_{\alpha 1}, \tag{A.19}$$

and the distribution functions are expanded as

$$f_i = f_i^{(0)} + \varepsilon f_i^{(1)} + \varepsilon^2 f_i^{(2)} + \cdots. \tag{A.20}$$

Appendix A. Appendix

At leading order, $\mathcal{O}(\varepsilon^0)$, the system is in local equilibrium,

$$f_i^{(0)} = f_i^{\mathrm{eq}}. \tag{A.21}$$

At this order, no macroscopic evolution equation appears; instead, the equilibrium moments determine the thermodynamic structure of the system. In particular, the second moment yields

$$\sum_i c_{i\alpha} c_{i\beta} f_i^{(0)} = \rho_{LB} c_s^2 \delta_{\alpha\beta} + \rho_{LB} u_\alpha u_\beta, \tag{A.22}$$

corresponding to the isothermal equation of state $p = \rho_{LB} c_s^2$.

At first order, $\mathcal{O}(\varepsilon)$, defining the streaming operator

$$D_{1i} \equiv \partial_{t_1} + c_{i\alpha} \partial_{\alpha 1}, \tag{A.23}$$

we obtain

$$D_{1i} f_i^{(0)} = -\frac{1}{\tau} f_i^{(1)}, \qquad f_i^{(1)} = -\tau D_{1i} f_i^{\mathrm{eq}}. \tag{A.24}$$

Taking moments of this equation yields the macroscopic conservation laws

$$\partial_{t_1} \rho_{LB} + \partial_{\alpha 1} (\rho_{LB} u_\alpha) = 0, \tag{A.25}$$

$$\partial_{t_1} (\rho_{LB} u_\beta) + \partial_{\alpha 1} \left(\rho_{LB} u_\alpha u_\beta + p\, \delta_{\alpha\beta} \right) = 0, \tag{A.26}$$

which correspond to the Euler equations describing inviscid transport.

At second order, $\mathcal{O}(\varepsilon^2)$, nonequilibrium corrections generate dissipative contributions. The nonequilibrium momentum flux tensor is defined as

$$\Pi_{\alpha\beta}^{(1)} = \sum_i c_{i\alpha} c_{i\beta} f_i^{(1)}, \tag{A.27}$$

which evaluates, using standard lattice isotropy conditions, to the Newtonian stress

tensor

$$\Pi^{(1)}_{\alpha\beta} = -\rho_{LB}\nu\left(\partial_{\alpha 1}u_\beta + \partial_{\beta 1}u_\alpha - \frac{2}{d}\delta_{\alpha\beta}\partial_{\gamma 1}u_\gamma\right), \tag{A.28}$$

with kinematic viscosity

$$\nu = c_s^2\left(\tau - \frac{\Delta t}{2}\right). \tag{A.29}$$

The second order equation then gives

$$\partial_{t_2}(\rho_{LB}u_\beta) + \partial_{\alpha 1}\Pi^{(1)}_{\alpha\beta} = 0. \tag{A.30}$$

Combining the $\mathcal{O}(\varepsilon)$ and $\mathcal{O}(\varepsilon^2)$ contributions and restoring physical derivatives yields the (isothermal) compressible Navier–Stokes equations,

$$\partial_t(\rho_{LB}u_\beta)+\partial_\alpha(\rho_{LB}u_\alpha u_\beta) = -\partial_\beta p+\partial_\alpha\left[\rho_{LB}\nu\left(\partial_\alpha u_\beta + \partial_\beta u_\alpha - \frac{2}{d}\delta_{\alpha\beta}\partial_\gamma u_\gamma\right)\right], \qquad p = \rho_{LB}c_s^2. \tag{A.31}$$

In the weakly compressible (low Mach) regime, this reduces to

$$\partial_t\mathbf{u} + (\mathbf{u}\cdot\nabla)\mathbf{u} = -\frac{1}{\rho_{LB}}\nabla p + \nu\nabla^2\mathbf{u}, \tag{A.32}$$

showing explicitly that the Chapman-Enskog expansion generates diffusive behaviour through relaxation toward equilibrium. This emergent diffusion mechanism is the key feature of lattice Boltzmann dynamics that we aim to reproduce within the quantum walk framework.

A.4 Classical Lattice Boltzmann Comparison Models

Throughout this chapter, the quantum walk simulations are compared with small classical D1Q2 and D1Q3 lattice Boltzmann models. These comparisons are not intended to claim that the quantum walk is exactly equivalent to LBM at the microscopic level. Instead, the LBM models provide simple classical reference systems with the same fitted macroscopic drift and/or diffusion behaviour. This allows the quantum-walk probability density to be compared with a familiar classical transport profile.

The two state quantum walks are compared with a D1Q2 LBM model, while the

lazy three state walks are compared with a D1Q3 LBM model. This choice mirrors the velocity structure of the corresponding quantum walks. The D1Q2 model has only left- and right-moving populations, whereas the D1Q3 model also contains a stationary rest population.

The D1Q2 model uses the two velocities

$$c_i \in -1, +1,$$

with weights

$$w_i = \left(\frac{1}{2}, \frac{1}{2}\right).$$

The corresponding lattice sound speed is

$$c_s^2 = 1.$$

The density is defined as

$$\rho_{LB}(x,t) = \sum_i f_i(x,t),$$

where $f_i(x,t)$ are the LBM populations associated with the lattice velocities.

The D1Q3 model uses the three velocities

$$c_i \in -1, 0, +1,$$

with weights

$$w_i = \left(\frac{1}{6}, \frac{4}{6}, \frac{1}{6}\right),$$

and lattice sound speed

$$c_s^2 = \frac{1}{3}.$$

Again, the density is defined as

$$\rho_{LB}(x,t) = \sum_i f_i(x,t).$$

In both cases, the LBM update consists of a BGK relaxation step followed by streaming. The relaxation time is chosen so that the LBM diffusion coefficient matches the effective diffusion coefficient measured from the quantum walk simulation.

A.4.1 Diffusion comparison

For the diffusion only comparisons, no advective drift is imposed in the LBM model. The equilibrium distribution is therefore taken to be the zero velocity equilibrium

$$f_i^{\mathrm{eq}}(x,t) = w_i \rho(x,t),$$

with the appropriate weights w_i for either D1Q2 or D1Q3.

The effective diffusion coefficient is fitted directly from the quantum walk variance using

$$\sigma^2(t) \simeq \sigma^2(0) + 2D_{\mathrm{eff}}t, \qquad D_{\mathrm{eff}} = \frac{1}{2}\frac{d\sigma^2}{dt}.$$

Once D_{eff} has been obtained from the quantum walk, the LBM relaxation time is fixed by

$$\tau = \frac{1}{2} + \frac{D_{\mathrm{eff}}}{c_s^2}.$$

Equivalently, the BGK relaxation rate is

$$\omega = \frac{1}{\tau}.$$

The resulting LBM density is then evolved from the same initial spatial density as the quantum walk and compared with the quantum walk probability density at selected times. This comparison should be interpreted carefully. Since D_{eff} is fitted from the quantum walk variance, the LBM model is calibrated to match the second moment of the quantum walk. The comparison therefore tests whether the full quantum walk density profile is close to a classical diffusion profile once the measured variance growth has been matched. A good variance fit alone does not prove that the full density is classical diffusive; this is why probability density snapshots and distance metrics are also used.

A.4.2 Advection-diffusion comparison

For simulations where the quantum walk has a non-zero drift, the comparison model is extended to a D1Q3 advection-diffusion LBM model. In this case, both the drift velocity and the diffusion coefficient are fitted from the quantum walk moments. The mean position is fitted using

$$\langle x(t)\rangle \simeq \langle x(0)\rangle + ut,$$

so that the fitted advection velocity is

$$u = \frac{d\langle x\rangle}{dt}.$$

The diffusion coefficient is again fitted from the variance using

$$\sigma^2(t) \simeq \sigma^2(0) + 2D_{\text{eff}}t.$$

Both fits are performed over the same late-time fitting window.

The LBM equilibrium is then taken to be the linearised D1Q3 advection equilibrium

$$f_i^{\text{eq}}(x,t) = w_i\rho(x,t)\left(1 + \frac{c_i u}{c_s^2}\right).$$

The fitted D_{eff} fixes the relaxation time through

$$\tau = \frac{1}{2} + \frac{D_{\text{eff}}}{c_s^2},$$

while the fitted u fixes the imposed drift in the LBM equilibrium. This gives a classical advection-diffusion reference solution whose mean drift and variance growth are matched to the quantum walk. As in the diffusion only case, this should be understood as a macroscopic comparison rather than a microscopic equivalence. The LBM model is fitted to the quantum walk moments, and the non-trivial test is whether the full density profile of the quantum walk agrees with the matched classical advection-diffusion solution. This is useful because it separates moment level agreement from agreement

of the full probability distribution.

A.5 Full Decoherence Strength Sweep for the Parity Echo Diffusion Simulations

This appendix contains the full set of results from the dephasing strength sweep discussed in Section 5.4.1. In the main text, only representative values of γ were shown in order to keep the discussion focused. Here, I include the full range of simulations for $0 \leq \gamma \leq 0.5$, including probability density snapshots and variance growth plots for each value of the dephasing strength.

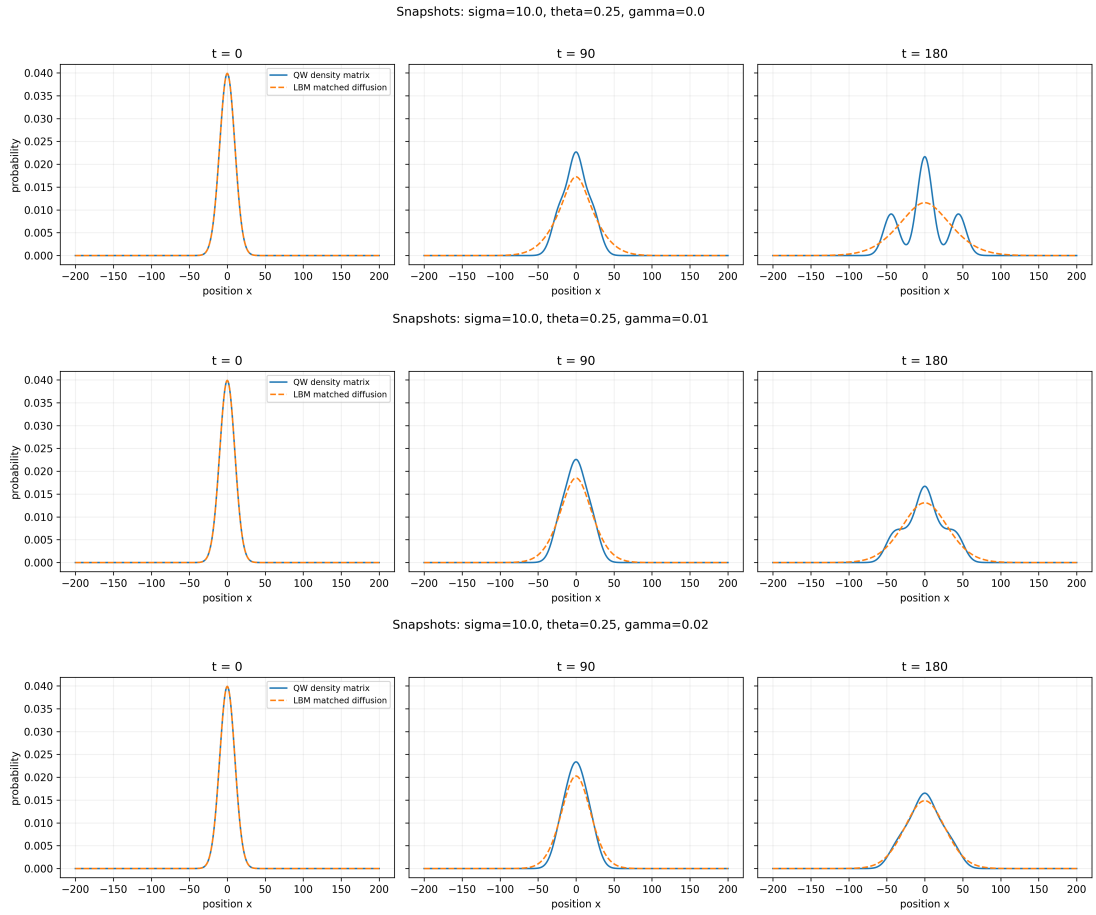

Appendix A. Appendix

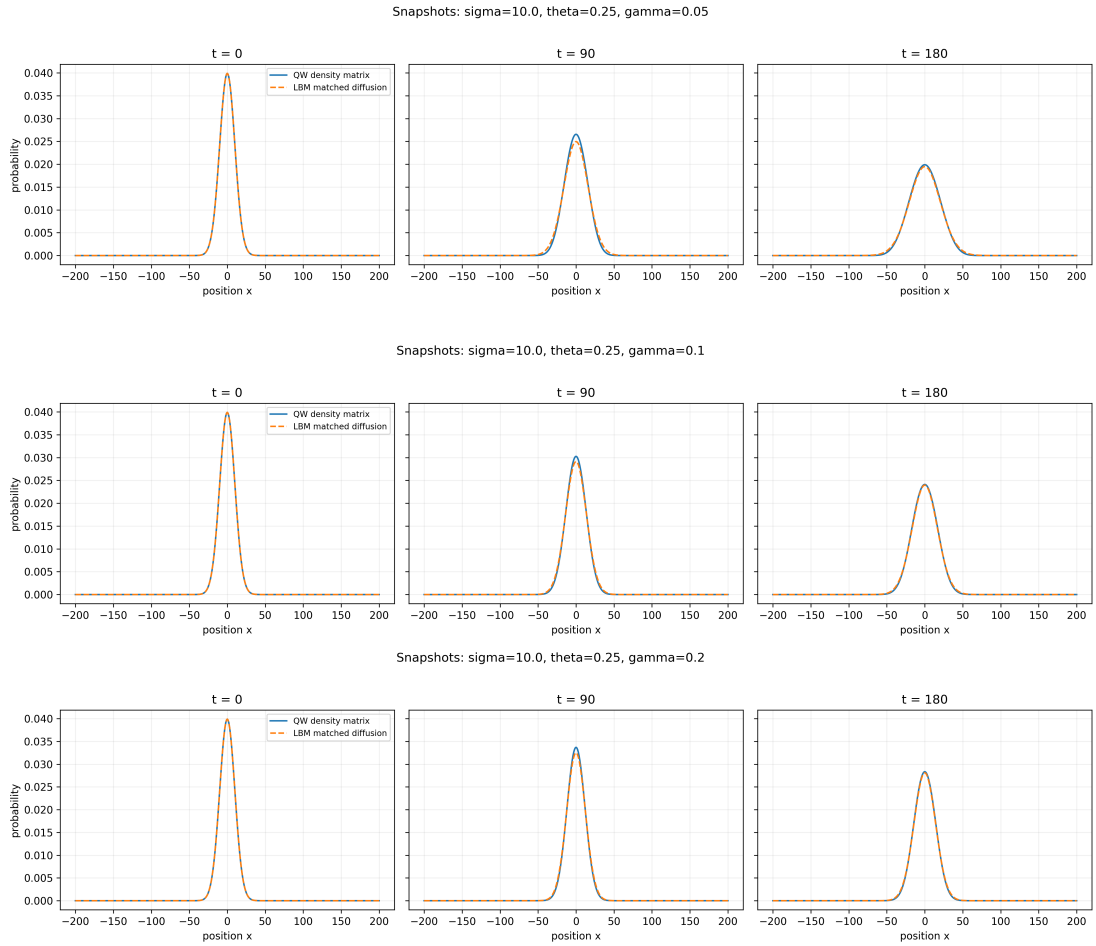

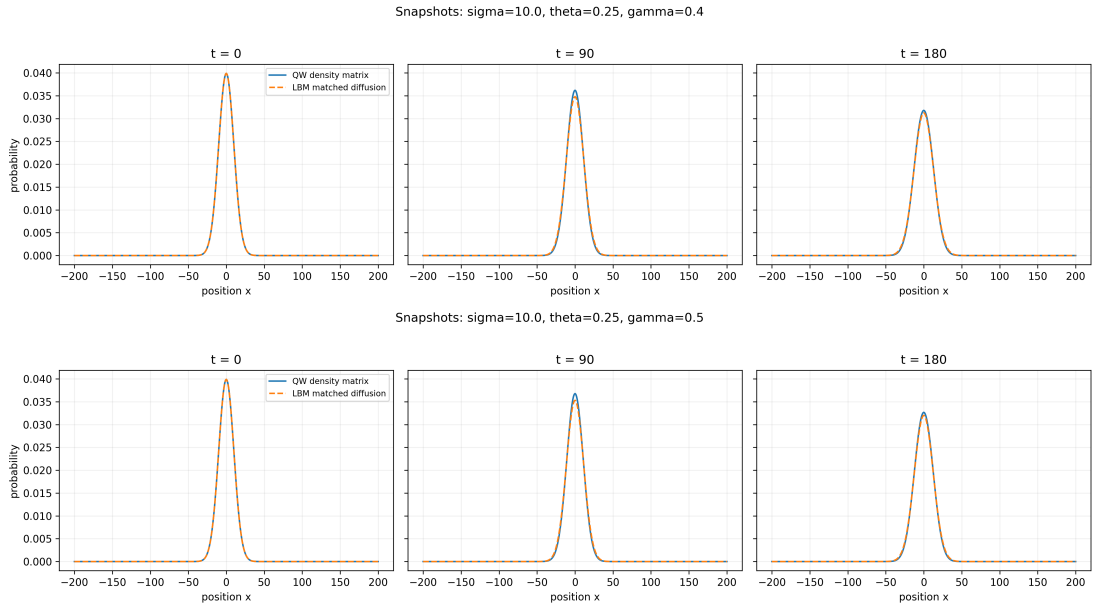

Figure A.1: Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths range from $\gamma = 0$ to $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve.

Figure A.1 shows the probability distribution functions for different values of γ. The variance plots, given by Figure A.2 are used to verify that $\sigma^2(t)$ is approximately linear over the fitted time window, while the snapshots show whether the full probability density is actually close to a smooth classical diffusion profile.

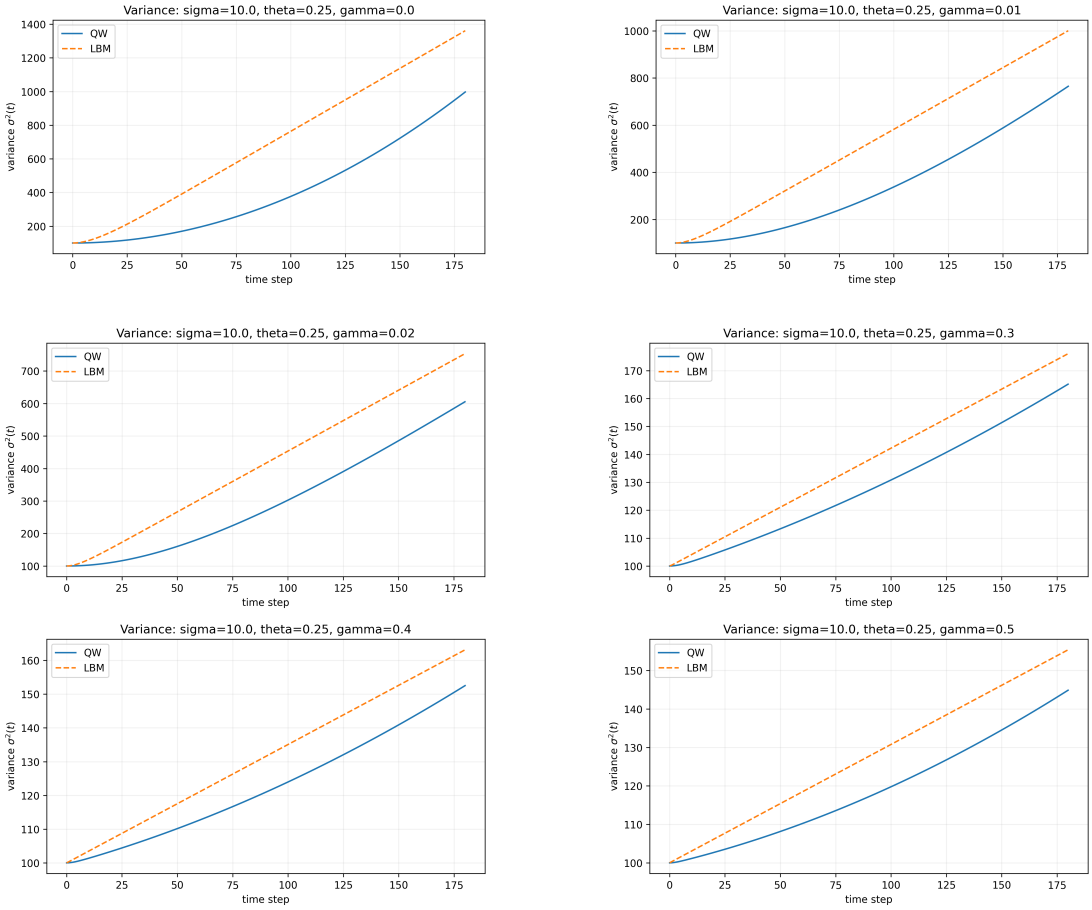

Figure A.2: Variance growth of the parity echo quantum walk for different dephasing strengths γ. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. The high linear fit quality across the sweep indicates approximately diffusive second moment scaling, although the full density snapshots are still needed to assess whether the probability profile is close to classical diffusion.

A.6 Full Coin Angle Sweep for the Parity Echo Diffusion Simulations

This appendix contains the full set of results from the coin angle sweep discussed in Section 5.4.2. In the main text, only representative values of θ were shown, together with the fitted scaling of the effective diffusion coefficient. Here, I include the full range of simulations, including probability density snapshots and variance growth plots for each value of the coin angle.

These additional plots show how the coherent part of the parity echo step controls the spreading rate of the walk. For small to moderate values of θ, the Chapman-Enskog closure predicts an approximate scaling $D_{\text{eff}} \propto \bar{\theta}^2$. The full sweep is therefore used to check both the fitted diffusion coefficient and the shape of the probability density as θ is varied. The snapshots are particularly useful because larger values of θ may move the model away from the small angle regime assumed in the analytic expansion, even when the variance growth still appears approximately linear. Figure A.4 shows the probability distributions of the models for varying values of θ while Figure A.4 shows the variance of the models for varying values of θ.

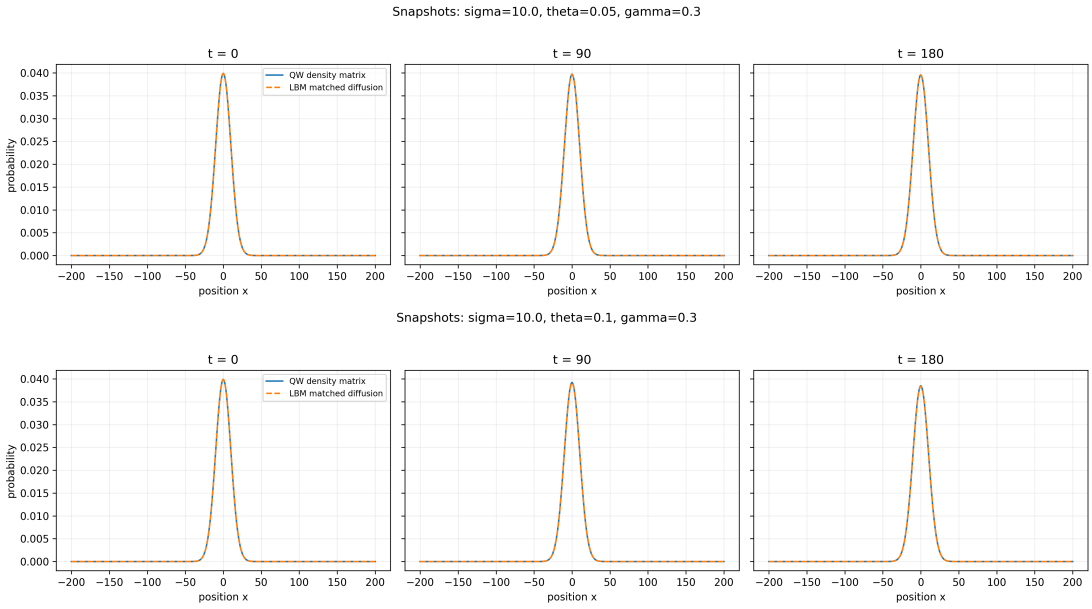

Appendix A. Appendix

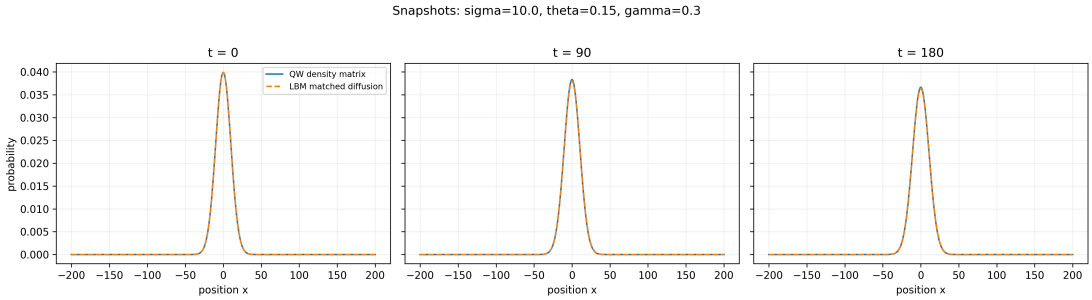

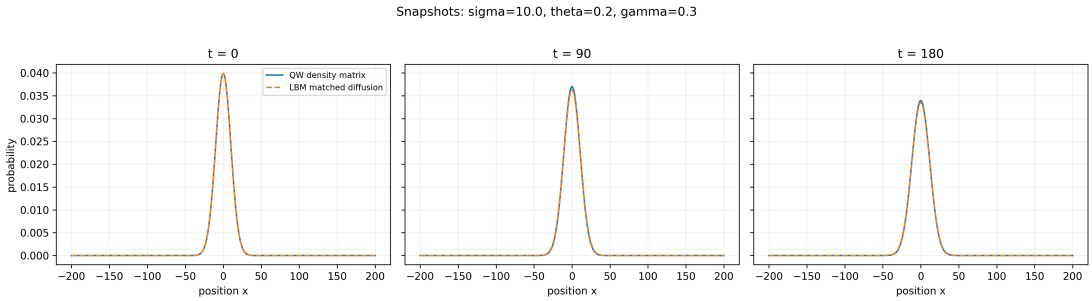

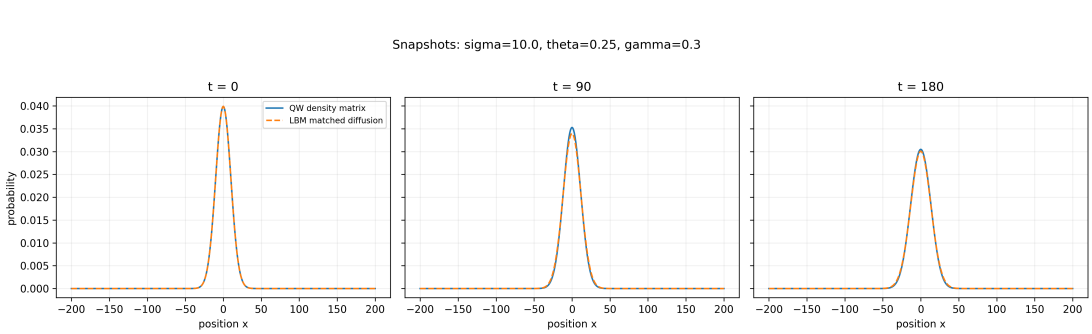

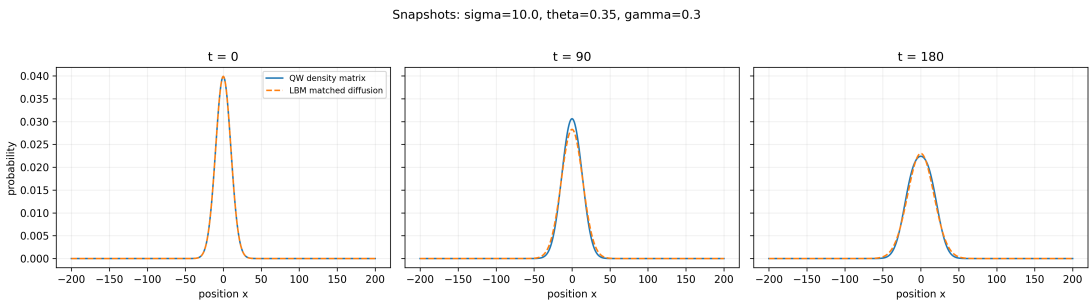

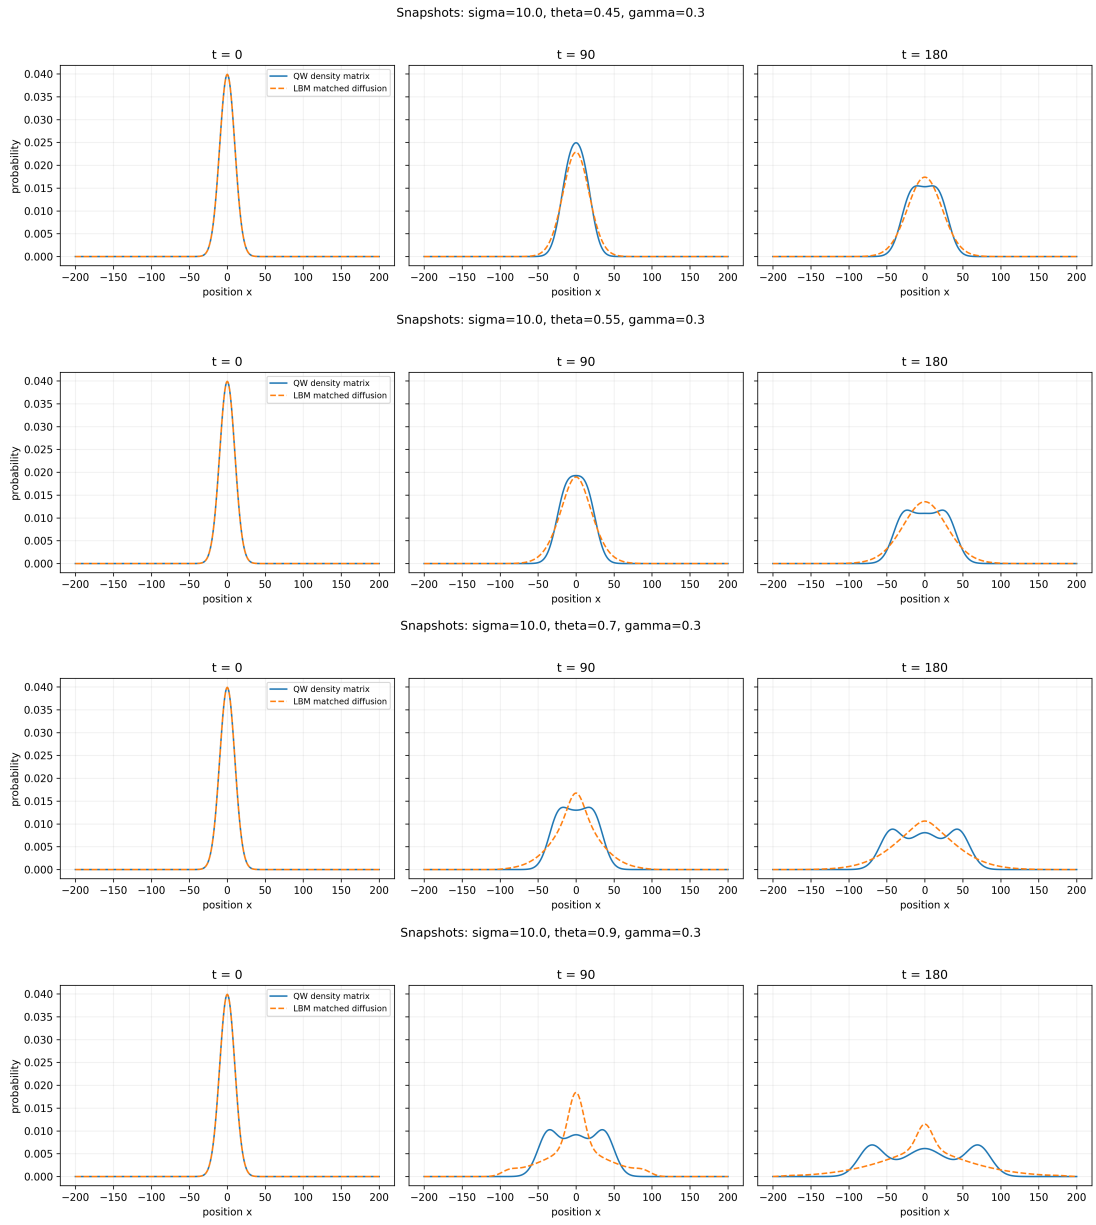

Figure A.3: Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM diffusion model for increasing coin angle θ. From top to bottom, the coin angles range from $\theta = 0.05$ to $\theta = 0.9$. The decoherence rate is fixed at $\gamma = 0.3$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve.

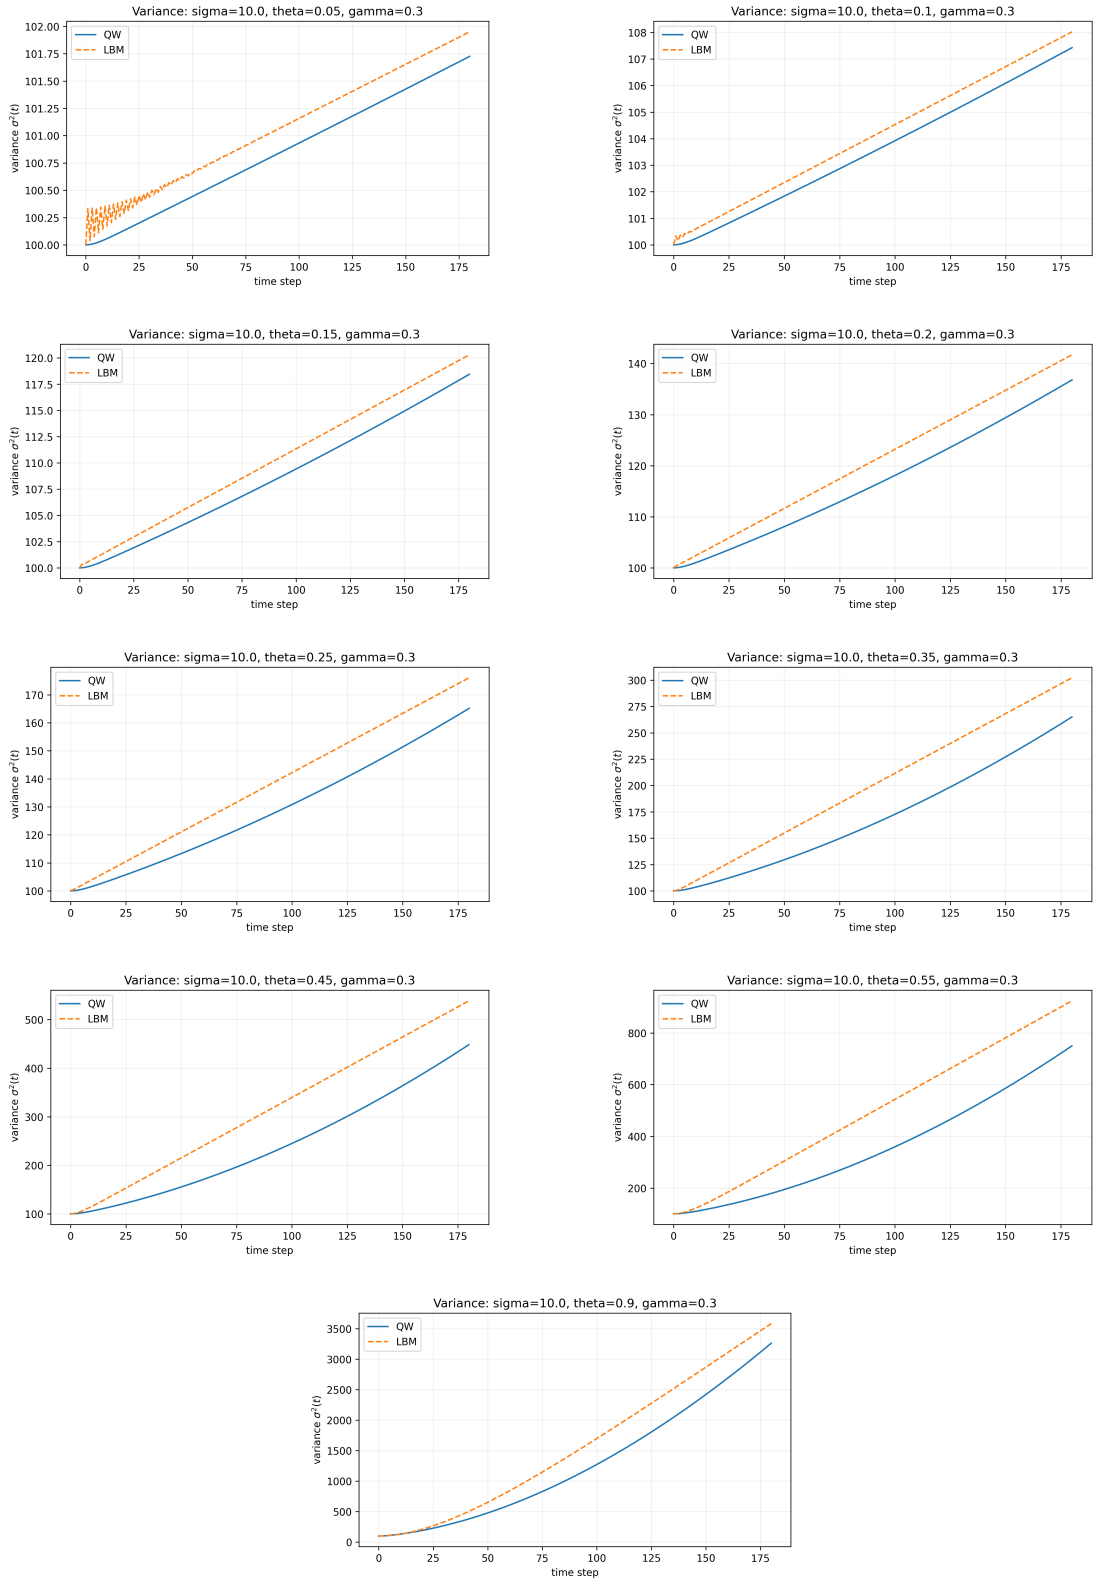

Figure A.4: Variance growth of the parity echo quantum walk for different coin angles θ, with $\gamma = 0.3$ and $\sigma_0 = 10$. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. Lower values of θ appear to give a higher linear fit, and extremely high angles seem to lead to sub-diffusive scaling.

A.7 Full Decoherence Strength Sweep for the Parity Echo Advection-Diffusion Simulations

This appendix contains the full set of results from the dephasing strength sweep discussed in Section 5.5.2. In the main text, only representative values of γ were shown in order to keep the discussion focused. Here, I include the full range of simulations for $0 \leq \gamma \leq 0.5$, including probability density snapshots and variance growth plots for each value of the dephasing strength. Figure A.5 shows the probability snapshots while Figure A.6 shows the corresponding variance plots.

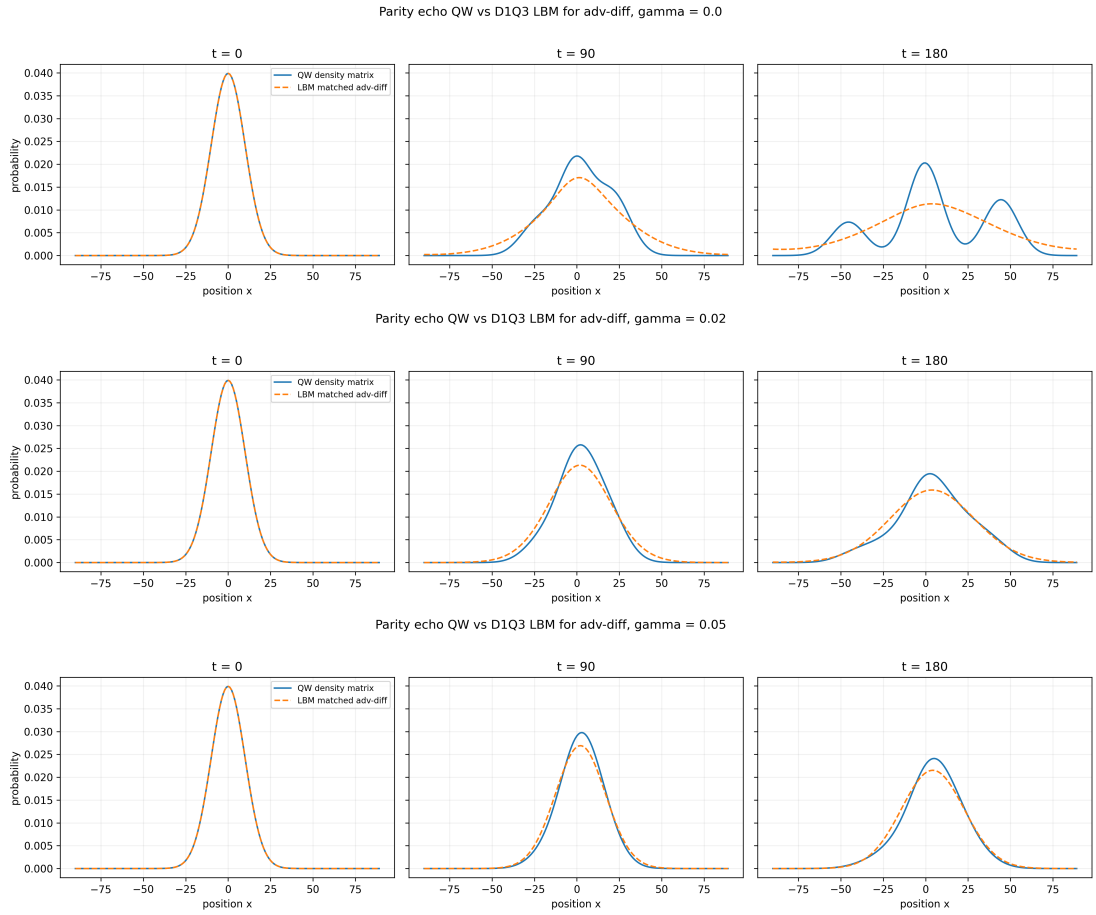

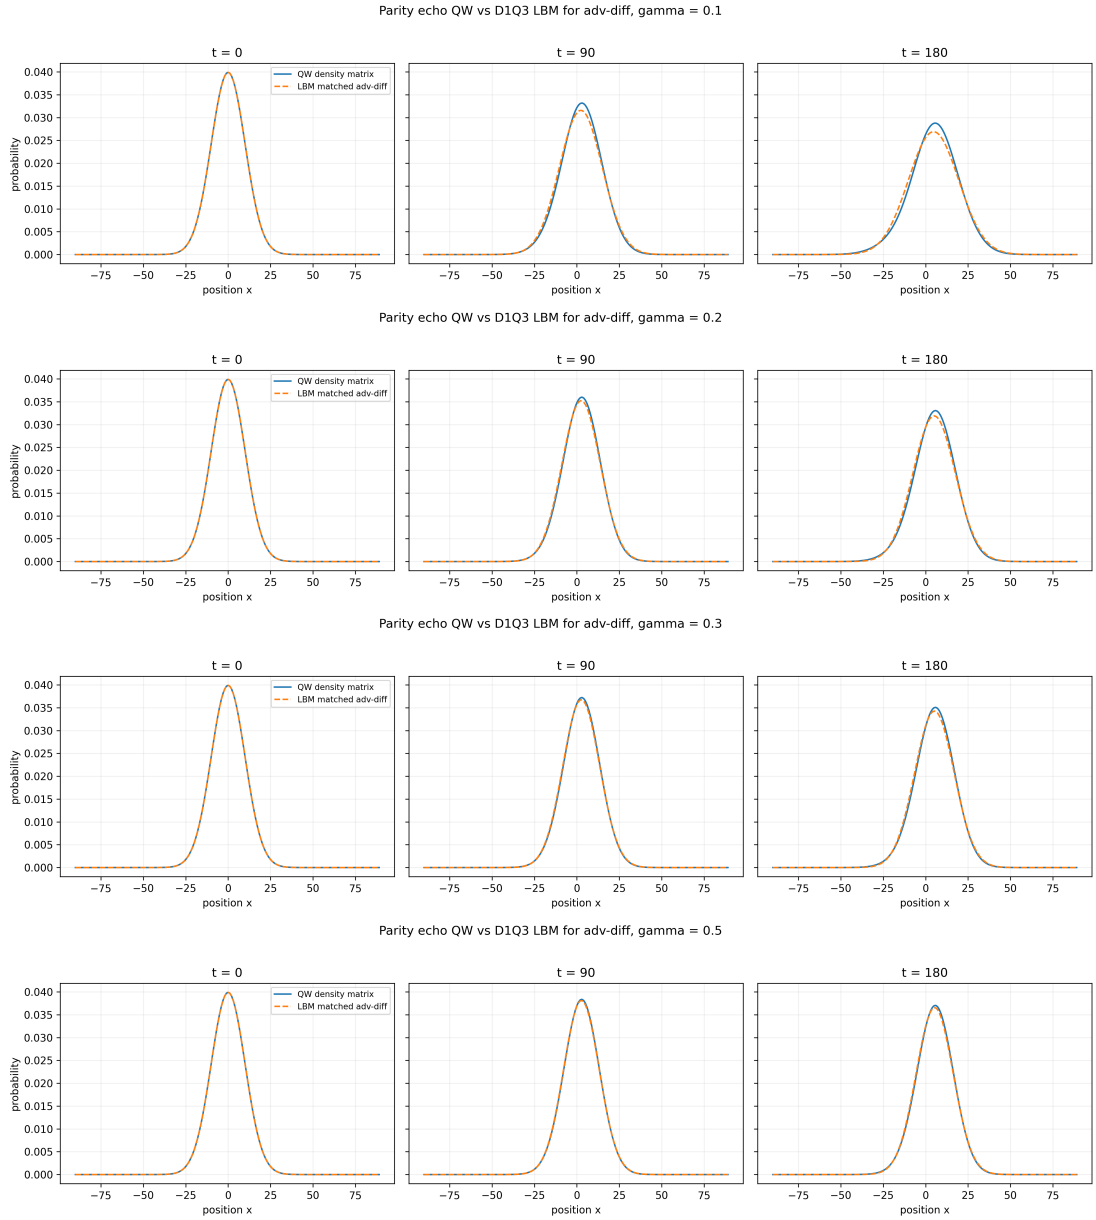

Figure A.5: Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model for increasing dephasing strength. From top to bottom, the dephasing strengths range from $\gamma = 0$ to $\gamma = 0.5$. The coin angle is fixed at $\theta = 0.25$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve.

Appendix A. Appendix

These additional plots show how the transition from coherent quantum spreading to smoother advection-diffusive behaviour depends on γ. They also provide a check that the trends discussed in the main text are not specific to the selected representative cases. In particular, the variance plots are used to verify that $\sigma^2(t)$ is approximately linear over the fitted time window, while the snapshots show whether the full probability density is actually close to a smooth classical diffusion profile.

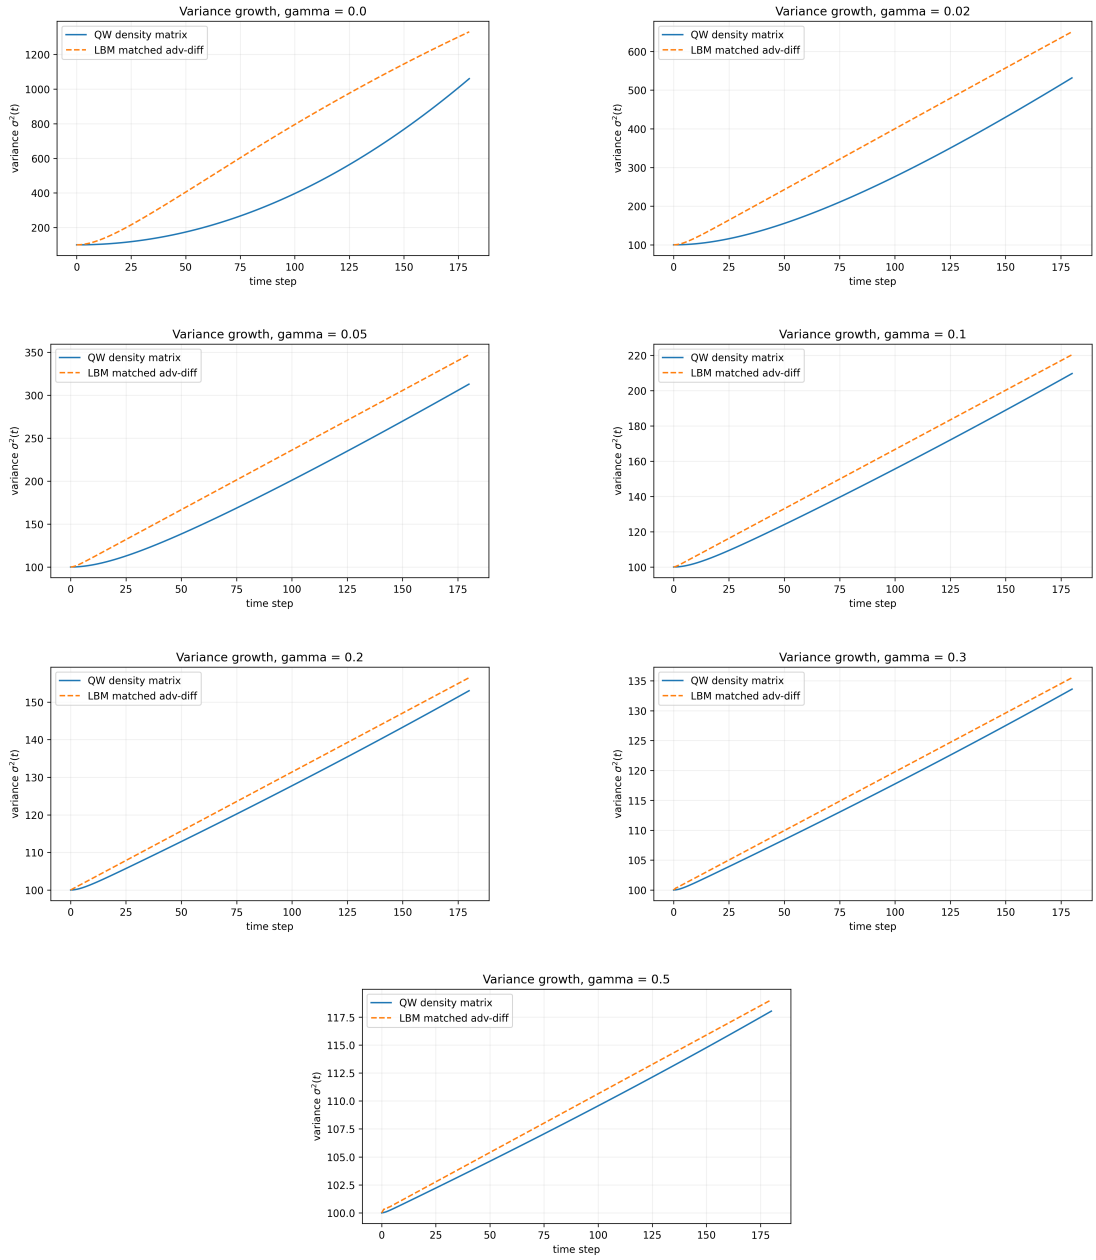

Figure A.6: Variance growth of the parity echo quantum walk for different dephasing strengths γ for the advection-diffusion model. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. The high linear fit quality across the sweep indicates approximately diffusive second moment scaling, although the full density snapshots are still needed to assess whether the probability profile is close to classical diffusion.

A.8 Full Coin Angle Sweep for the Parity Echo Advection-Diffusion Simulations

This appendix contains the full set of results from the coin angle sweep discussed in Section 5.5.2. In the main text, only representative values of θ were shown, together with the fitted scaling of the effective diffusion coefficient. Here, I include the full range of simulations, including probability density snapshots and variance growth plots for each value of the coin angle. Figure A.7 shows the probability snapshots while Figure A.8 shows the corresponding variance plots. These additional plots show how the coherent part of the parity echo step controls the spreading rate of the walk and advection. For small to moderate values of θ, the Chapman-Enskog closure predicts an approximate scaling $D_{\text{eff}} \propto \bar{\theta}^2$. The full sweep is therefore used to check both the fitted diffusion coefficient and the shape of the probability density as θ is varied. The snapshots are particularly useful because larger values of θ may move the model away from the small angle regime assumed in the analytic expansion, even when the variance growth still appears approximately linear.

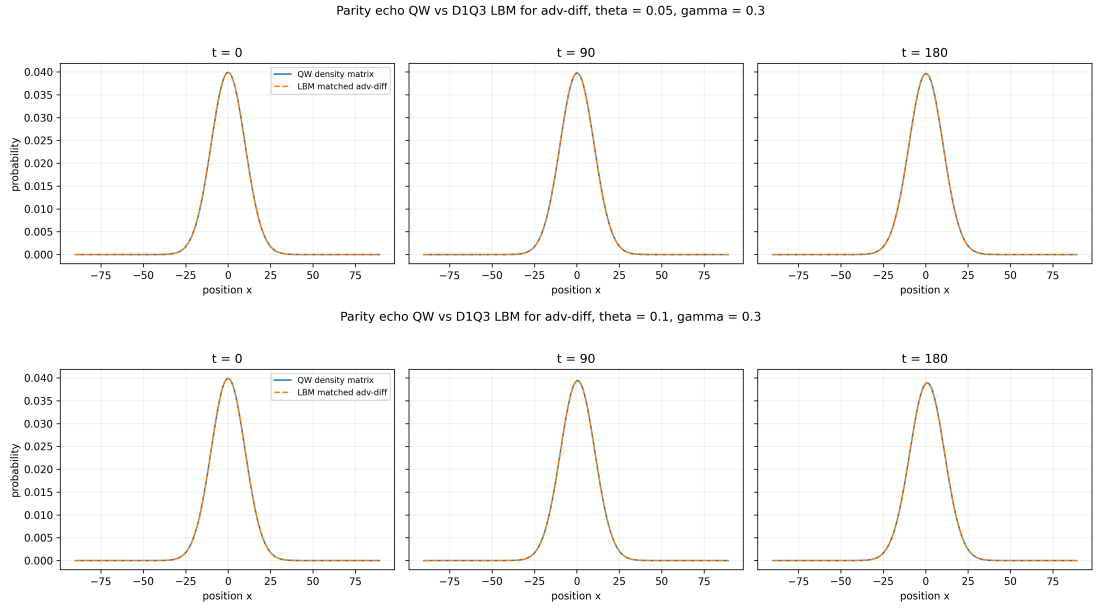

Appendix A. Appendix

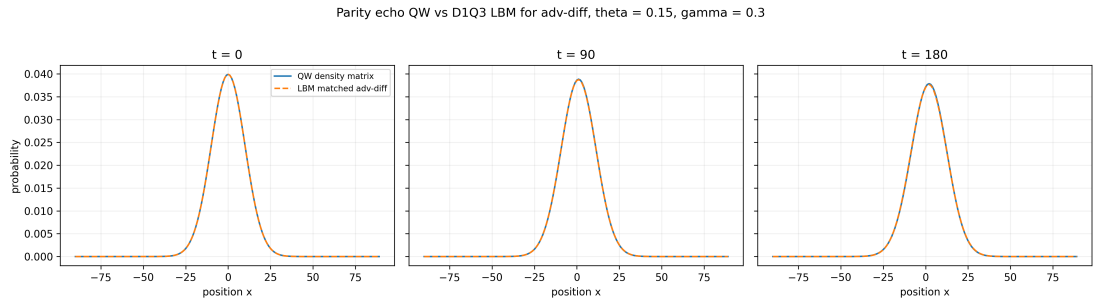

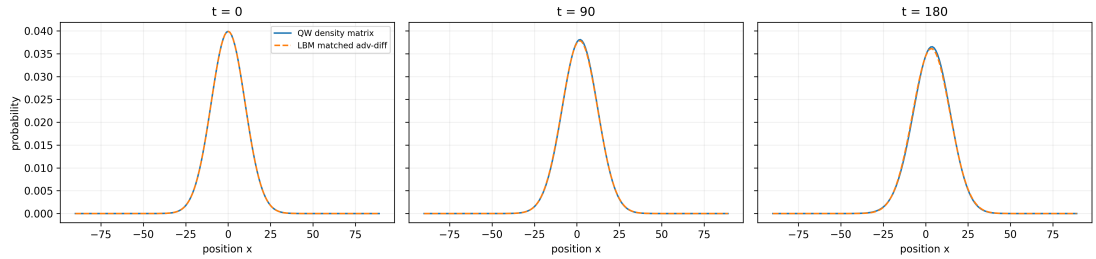

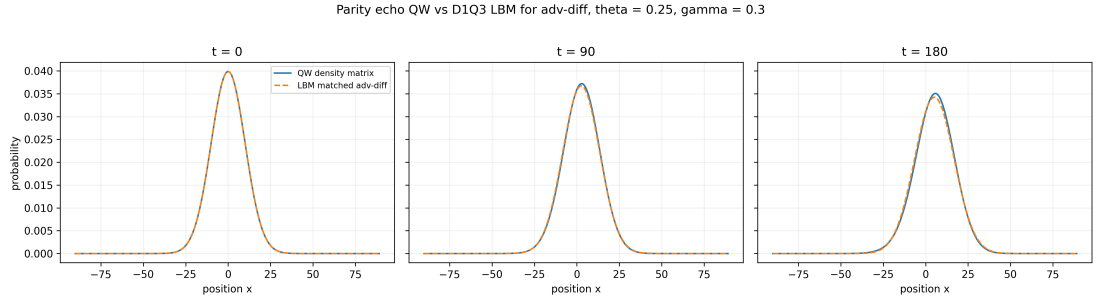

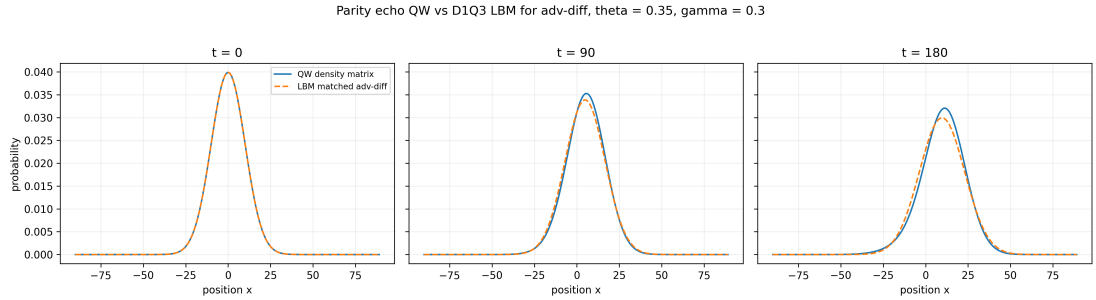

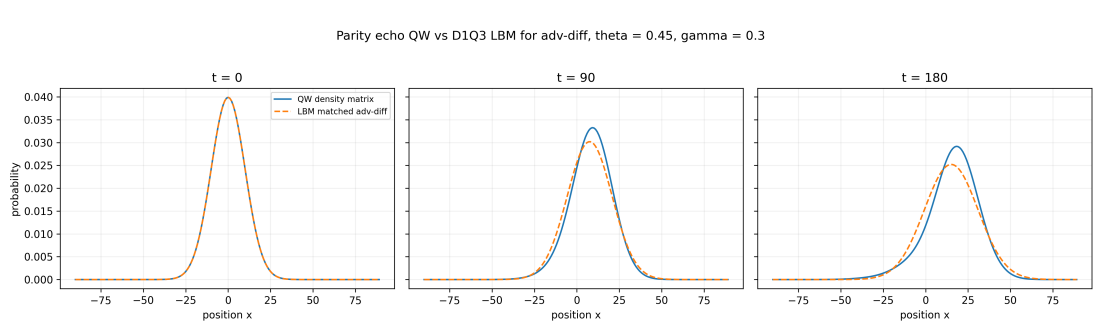

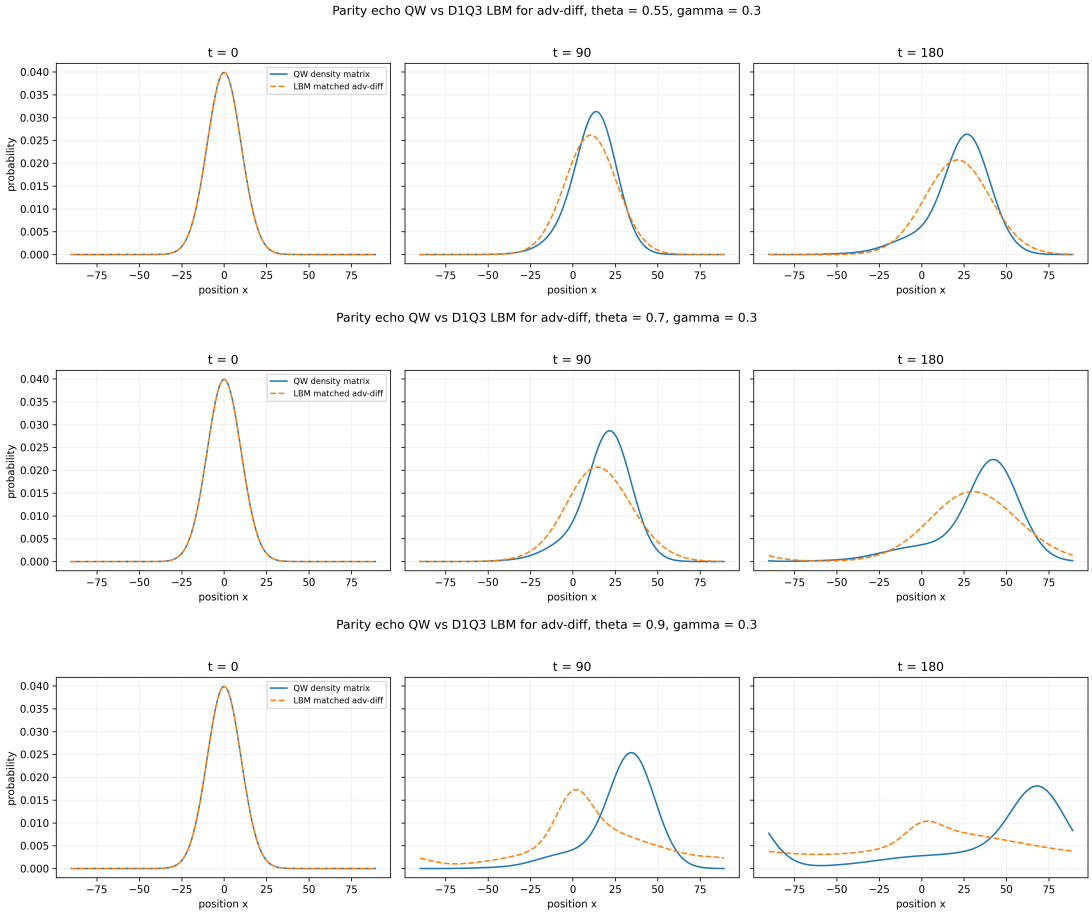

Figure A.7: Probability density snapshots for the parity echo quantum walk compared with the matched D1Q3 LBM advection-diffusion model for increasing coin angle θ. From top to bottom, the coin angles range from $\theta = 0.05$ to $\theta = 0.9$. The decoherence rate is fixed at $\gamma = 0.3$, and the initial spatial state is a Gaussian with width $\sigma_0 = 10$. In each case, the LBM diffusion coefficient is fitted from the variance growth of the quantum walk and then used as the comparison curve.

Appendix A. Appendix

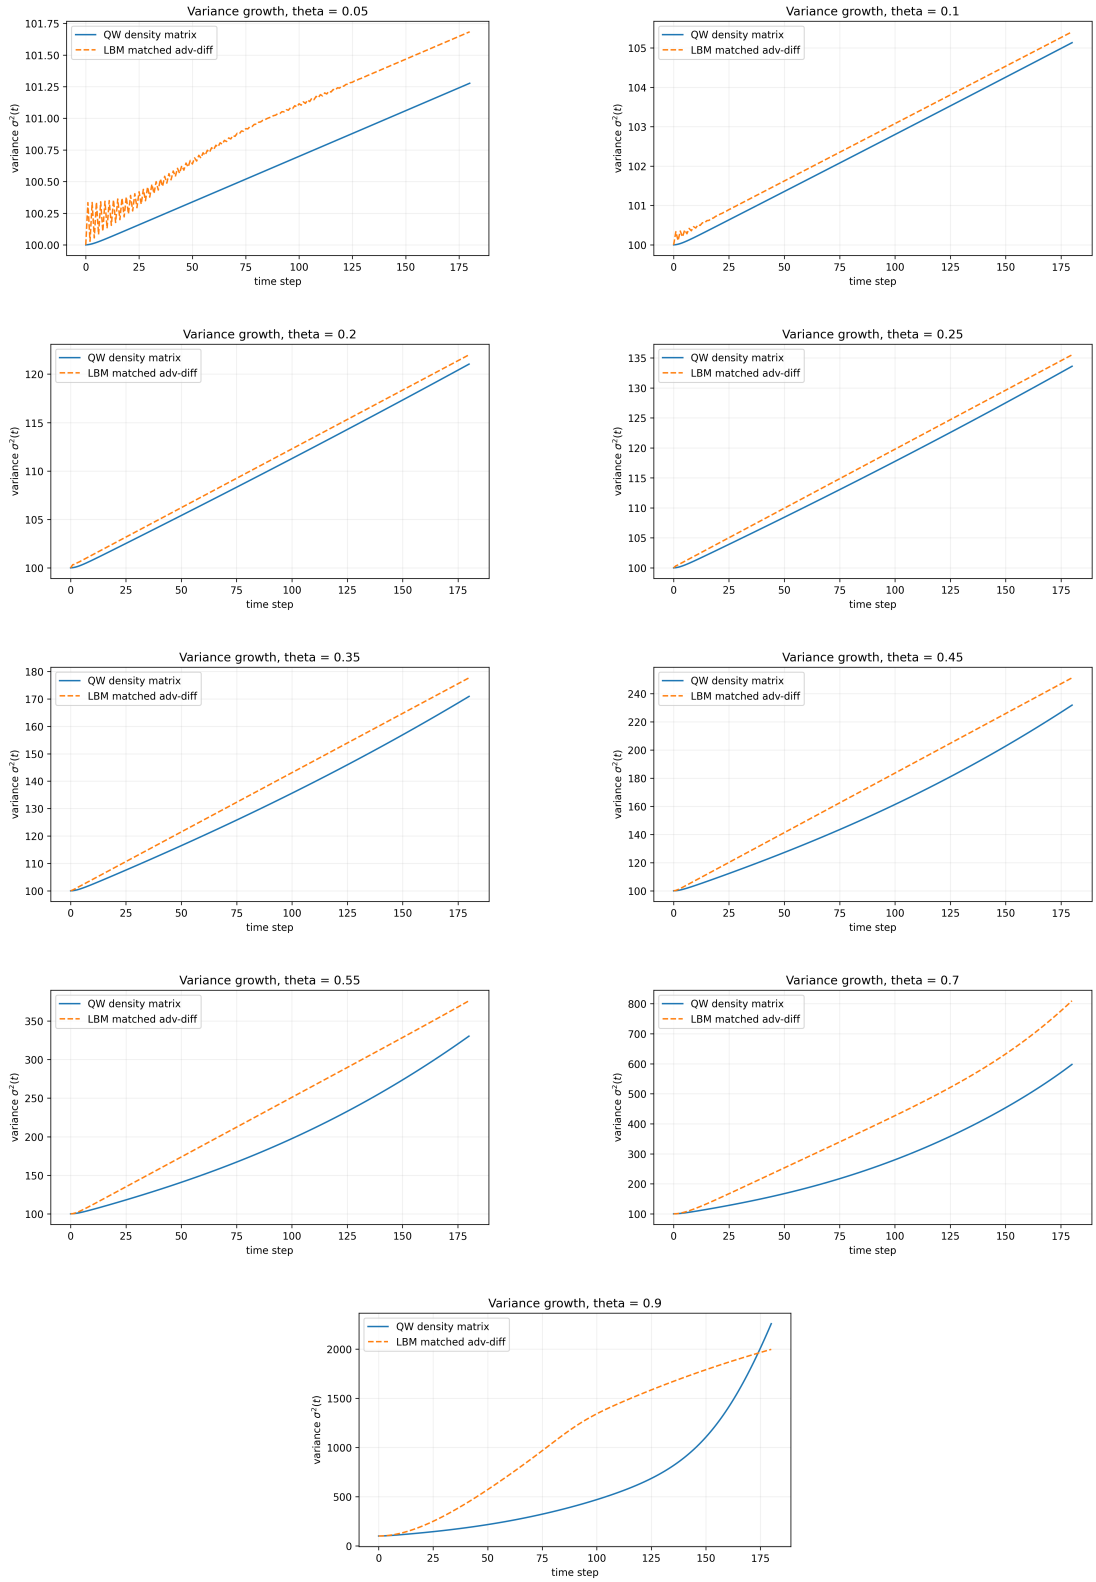

Figure A.8: Variance growth of the parity echo quantum walk for different coin angles θ, with $\gamma = 0.3$ and $\sigma_0 = 10$. Each panel shows the quantum walk variance together with the fitted linear growth used to extract D_{eff}. Lower values of θ appear to give a higher linear fit, and extremely high angles seem to lead to sub-diffusive scaling.

Bibliography

Richard P Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21, 1982. doi: 10.1007/BF02650179.

A Da Ronch, D Vallespin, M Ghoreyshi, and KJ Badcock. Evaluation of dynamic derivatives using computational fluid dynamics. *AIAA journal*, 50(2):470–484, 2012.

Mori Mani and Andrew J. Dorgan. A perspective on the state of aerospace computational fluid dynamics technology. *Annual Review of Fluid Mechanics*, 55 (Volume 55, 2023):431–457, 2023. ISSN 1545-4479. doi: https://doi.org/10.1146/annurev-fluid-120720-124800. URL `https://www.annualreviews.org/content/journals/10.1146/annurev-fluid-120720-124800`.

Giulio Malinverno, Javier Blasco Alberto, Jon Lecumberri SanMartin, and Lucas Lamata Manuel. Quantum computing CFD simulations: review and state of the art. *American Society of Thermal and Fluids Engineering*, 2025.

Amr Elsharkawy, Xiao-Ting Michelle To, Philipp Seitz, Yanbin Chen, Yannick Stade, Manuel Geiger, Qunsheng Huang, Xiaorang Guo, Muhammad Arslan Ansari, Christian B. Mendl, Dieter Kranzlmüller, and Martin Schulz. Integration of quantum accelerators with high performance computing—a review of quantum programming tools. *ACM Transactions on Quantum Computing*, 6(3), July 2025. doi: 10.1145/3743149. URL `https://doi.org/10.1145/3743149`.

Richard P. Feynman and Albert R. Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill, New York, 1965. Problem 2-6, pp. 34–38.

A. Ambainis. Quantum walks and their algorithmic applications, 2004. URL https://arxiv.org/abs/quant-ph/0403120. arXiv preprint.

Sauro Succi, François Fillion-Gourdeau, and Silvia Palpacelli. Quantum lattice Boltzmann is a quantum walk. *EPJ Quantum Technology*, 2(1):12, December 2015. doi: 10.1140/epjqt/s40507-015-0025-1. URL https://ui.adsabs.harvard.edu/abs/2015EPJQT...2...12S.

David A. Meyer. From quantum cellular automata to quantum lattice gases. *Journal of Statistical Physics*, 85(5-6):551–574, December 1996. ISSN 1572-9613. doi: 10.1007/bf02199356. URL http://dx.doi.org/10.1007/BF02199356.

David Wawrzyniak, Josef Winter, Steffen Schmidt, Thomas Indinger, Christian F. Janßen, Uwe Schramm, and Nikolaus A. Adams. A quantum algorithm for the lattice-boltzmann method advection-diffusion equation. *Computer Physics Communications*, 306:109373, 2025. ISSN 0010-4655. doi: https://doi.org/10.1016/j.cpc.2024.109373. URL https://www.sciencedirect.com/science/article/pii/S0010465524002960.

Antonio David Bastida Zamora, Ljubomir Budinski, Valtteri Lahtinen, and Pierre Sagaut. Quantum lattice boltzmann method for several time steps: A local carleman linearization algorithm. *Phys. Rev. E*, 113:035307, Mar 2026. doi: 10.1103/kbx3-c35d. URL https://link.aps.org/doi/10.1103/kbx3-c35d.

David Barral, F. Javier Cardama, Guillermo Díaz-Camacho, Daniel Faílde, Iago F. Llovo, Mariamo Mussa-Juane, Jorge Vázquez-Pérez, Juan Villasuso, César Piñeiro, Natalia Costas, Juan C. Pichel, Tomás F. Pena, and Andrés Gómez. Review of distributed quantum computing: From single QPU to high performance quantum computing. *Computer Science Review*, 57:100747, 2025. ISSN 1574-0137. doi: https://doi.org/10.1016/j.cosrev.2025.100747. URL https://www.sciencedirect.com/science/article/pii/S1574013725000231.

Tingting Liu, Dong Lu, Hao Zhang, Mingyue Zheng, Huaiyu Yang, Yechun Xu, Cheng Luo, Weiliang Zhu, Kunqian Yu, and Hualiang Jiang. Applying high-performance

computing in drug discovery and molecular simulation. *National science review*, 3 (1):49–63, 2016.

Antonella Galizia, Simone Cammarasana, Andrea Clematis, and Giuseppe Patane'. Evaluating accuracy and efficiency of HPC solvers for sparse linear systems with applications to pdes, 2022. URL `https://arxiv.org/abs/2201.05413`.

Laizet Sylvain. Lloyd Seth. Magri Luca Tennie, Felix. Quantum computing for nonlinear differential equations and turbulence. *Nature Reviews Physics*, 7, 2025. doi: 10.1038/s42254-024-00799-w.

Masoud Ghasemian, Z Najafian Ashrafi, and Ahmad Sedaghat. A review on computational fluid dynamic simulation techniques for darrieus vertical axis wind turbines. *Energy Conversion and Management*, 149:87–100, 2017.

Peter S Huyakorn. *Computational methods in subsurface flow*. academic press, 2012.

Timm Krüger, Halim Kusumaatmaja, Alexandr Kuzmin, Orest Shardt, Goncalo Silva, and Erlend Magnus Viggen. *The Lattice Boltzmann Method: Principles and Practice*. Graduate Texts in Physics. Springer International Publishing, 2016. ISBN 978-3-319-44647-9. doi: 10.1007/978-3-319-44649-3. URL `https://link.springer.com/book/10.1007/978-3-319-44649-3`.

MR Usikalu. Simulation of turbulent flow in pipes using the navier-stokes solver with finite volume method. *Association Journal of Interdisciplinary Technics in Engineering Mechanics*, 3(3):8–13, 2025.

Buyang Li, Shu Ma, and Weiwei Sun. Optimal analysis of finite element methods for the stochastic stokes equations. *Mathematics of Computation*, 94(352):551–583, 2025.

G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8): 114–117, 1965.

Giorgio Tosti Balducci, Boyang Chen, Matthias Möller, Marc Gerritsma, and Roeland De Breuker. Review and perspectives in quantum computing for partial differential

equations in structural mechanics. *Frontiers in Mechanical Engineering*, 8:914241, 2022.

Navamita Ray, Tirtha Banerjee, Balasubramanya Nadiga, and Satish Karra. Towards solving the navier-stokes equation on quantum computers. *arXiv preprint arXiv:1904.09033*, 2019.

Wolf Schwarz. Random walk and diffusion models. *Springer. https://doi. org/10*, 1007: 978–3, 2022.

Zhao-Li Guo, Bao-Chang Shi, and Neng-Chao Wang. Fully lagrangian and lattice boltzmann methods for the advection-diffusion equation. *Journal of scientific computing*, 14(3):291–300, 1999.

Feng Xia, Jiaying Liu, Hansong Nie, Yonghao Fu, Liangtian Wan, and Xiangjie Kong. Random walks: A review of algorithms and applications. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(2):95–107, 2019.

Mark Kac. Random walk and the theory of brownian motion. *The American Mathematical Monthly*, 54(7P1):369–391, 1947.

Edward A Codling, Michael J Plank, and Simon Benhamou. Random walk models in biology. *Journal of the Royal society interface*, 5(25):813, 2008.

Steven P Blurton, Søren Kyllingsbæk, Carsten S Nielsen, and Claus Bundesen. A poisson random walk model of response times. *Psychological Review*, 127(3):362, 2020.

George H Weiss. Random walks and their applications: Widely used as mathematical models, random walks play an important role in several areas of physics, chemistry, and biology. *American Scientist*, 71(1):65–71, 1983.

Annamaria Ficara, Giacomo Fiumara, Pasquale De Meo, and Salvatore Catanese. Classical and quantum random walks to identify leaders in criminal networks. In *International Conference on Complex Networks and Their Applications*, pages 190–201. Springer, 2022.

Gregory F Lawler and Vlada Limic. *Random walk: a modern introduction*, volume 123. Cambridge University Press, 2010.

Edward A Codling, Rachel N Bearon, and Graeme J Thorn. Diffusion about the mean drift location in a biased random walk. *Ecology*, 91(10):3106–3113, 2010.

Ryszard Wojnar. Random walk, diffusion and wave equation. *Acta physica polonica B*, 44(5), 2013.

Michael F Shlesinger. Asymptotic solutions of continuous-time random walks. *Journal of Statistical Physics*, 10(5):421–434, 1974.

Raymond T Pierrehumbert. Lattice models of advection-diffusion. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 10(1):61–74, 2000.

Abdelkader Mojtabi and Michel O Deville. One-dimensional linear advection–diffusion equation: Analytical and finite element solutions. *Computers & Fluids*, 107:189–195, 2015.

Loukas Vlahos, Heinz Isliker, Yannis Kominis, and Kyriakos Hizanidis. Normal and anomalous diffusion: A tutorial, 2008. URL `https://arxiv.org/abs/0805.0419`.

Todd A. Brun, Hilary A. Carteret, and Andris Ambainis. Quantum to classical transition for random walks. *Phys. Rev. Lett.*, 91:130602, Sep 2003a. doi: 10.1103/PhysRevLett.91.130602. URL `https://link.aps.org/doi/10.1103/PhysRevLett.91.130602`.

U. Frisch, B. Hasslacher, and Y. Pomeau. Lattice-gas automata for the navier-stokes equation. *Phys. Rev. Lett.*, 56:1505–1508, Apr 1986. doi: 10.1103/PhysRevLett.56.1505. URL `https://link.aps.org/doi/10.1103/PhysRevLett.56.1505`.

Guy R. McNamara and Gianluigi Zanetti. Use of the boltzmann equation to simulate lattice-gas automata. *Phys. Rev. Lett.*, 61:2332–2335, Nov 1988. doi: 10.1103/PhysRevLett.61.2332. URL `https://link.aps.org/doi/10.1103/PhysRevLett.61.2332`.

Sauro Succi, Roberto Benzi, and Francisco Higuera. The lattice boltzmann equation: A new tool for computational fluid-dynamics. *Physica D: Nonlinear Phenomena*, 47(1):219–230, 1991. ISSN 0167-2789. doi: https://doi.org/10.1016/0167-2789(91)90292-H. URL `https://www.sciencedirect.com/science/article/pii/016727899190292H`.

P. L. Bhatnagar, E. P. Gross, and M. Krook. A model for collision processes in gases. i. small amplitude processes in charged and neutral one-component systems. *Phys. Rev.*, 94:511–525, May 1954. doi: 10.1103/PhysRev.94.511. URL `https://link.aps.org/doi/10.1103/PhysRev.94.511`.

Cyrus K Aidun and Jonathan R Clausen. Lattice-boltzmann method for complex flows. *Annual review of fluid mechanics*, 42(1):439–472, 2010.

Xiaoyi He and Li-Shi Luo. Theory of the lattice boltzmann method: From the boltzmann equation to the lattice boltzmann equation. *Physical review E*, 56(6):6811, 1997.

F. J. Alexander, S. Chen, and J. D. Sterling. Lattice boltzmann thermohydrodynamics. *Phys. Rev. E*, 47:R2249(R)–R2252(R), Apr 1993. doi: 10.1103/PhysRevE.47.R2249. URL `https://link.aps.org/doi/10.1103/PhysRevE.47.R2249`.

F. J. Alexander, H. Chen, S. Chen, and G. D. Doolen. Lattice boltzmann model for compressible fluids. *Phys. Rev. A*, 46:1967–1970, Aug 1992. doi: 10.1103/PhysRevA.46.1967. URL `https://link.aps.org/doi/10.1103/PhysRevA.46.1967`.

Y. H. Qian, D. D'Humières, and P. Lallemand. Lattice bgk models for navier-stokes equation. *Europhysics Letters*, 17(6):479, feb 1992. doi: 10.1209/0295-5075/17/6/001. URL `https://doi.org/10.1209/0295-5075/17/6/001`.

Heintz A. Gebäck T. A lattice boltzmann method for the advection-diffusion equation with neumann boundary conditions. *Communications in Computational Physics*, 15(2):487–505, Mar. 2018. doi: 10.4208/cicp.161112.230713a. URL `https://www.global-sci.com/cicp/article/view/6165`.

Jun Li. Appendix: Chapman-enskog expansion in the lattice boltzmann method, 2015. URL https://arxiv.org/abs/1512.02599.

Sydney Chapman. *The Mathematical Theory of Non-Uniform Gases (3rd ed.)*. Cambridge University Press, 1970.

Theodore McKeever and Ahsan Nazir. An introduction to the foundations and interpretations of quantum mechanics, 2026. URL https://arxiv.org/abs/2603.09818.

Max Born. Zur Quantenmechanik der Stoßvorgänge. *Zeitschrift für Physik*, 37(12): 863–867, 1926. doi: 10.1007/BF01397477. URL https://doi.org.

Dipankar Bhattacharyya and Jyotirmoy Guha. Bra ket algebra of dirac. In *Quantum Optics and Quantum Computation*, 2053-2563, pages 1–1 to 1–45. IOP Publishing, 2022. ISBN 978-0-7503-2715-2. doi: 10.1088/978-0-7503-2715-2ch1. URL https://doi.org/10.1088/978-0-7503-2715-2ch1.

Guanru Feng, Dawei Lu, Jun Li, Tao Xin, and Bei Zeng. Quantum computing: principles and applications, 2023. URL https://arxiv.org/abs/2310.09386.

Paul Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 1980. doi: 10.1007/BF01011339.

Yu. I. Manin. Vychislimoe i nevychislimoe [computable and noncomputable] (in russian). *Sov.Radio.*, pages 13–15, 1980.

Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, Bangalore, India, 1984.

Daniel R. Simon. On the power of quantum computation. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 116–123. IEEE, 1994. doi: 10.1109/SFCS.1994.365700.

Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. ISSN 1095-7111. doi: 10.1137/s0097539795293172. URL `http://dx.doi.org/10.1137/S0097539795293172`.

Lov K. Grover. A fast quantum mechanical algorithm for database search, 1996. URL `https://arxiv.org/abs/quant-ph/9605043`.

Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996. doi: 10.1126/science.273.5278.1073. URL `https://www.science.org/doi/abs/10.1126/science.273.5278.1073`.

Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Ramiro Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574 (7779):505–510, 2019a. doi: 10.1038/s41586-019-1666-5.

Edwin Pednault, John A Gunnels, Giacomo Nannicini, L Hadopek, and Jay M Gambetta. Leveraging secondary storage to simulate deep quantum circuits using classical supercomputers. *arXiv preprint arXiv:1910.09534*, 2019.

Han-Shi Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian-Wen Qin, Dian Wu, Xing Ding, Yi Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. doi: 10.1126/science.abe8770.

Changhun Oh, Minzhao Liu, Yu Choi, Warit Asavanant, Mamoru Endo, Kosuke Junase, Jun-ichi Yoshiikawa, Akira Furusawa, Liang Jiang, and Bill Jiang. Classical algorithm for simulating experimental gaussian boson sampling. *Nature Physics*, 20 (9):1461–1467, 2024. doi: 10.1038/s41567-024-02520-y.

Joseph Tindall, Norbert Schuch, Hao P Frank, and F Dieterle. Efficient classical simulation of cluster state quantum circuits with tensor networks. *arXiv preprint arXiv:2306.03709*, 2024.

Hiu Yung Wong. *Quantum Gate Introduction: NOT and CNOT Gates*, pages 133–141. Springer International Publishing, Cham, 2024. ISBN 978-3-031-36985-8. doi: 10.1007/978-3-031-36985-8_15. URL https://doi.org/10.1007/978-3-031-36985-8_15.

Roselyn Nmaju, Fiona Speirits, and Sarah Croke. Low-depth measurement-based deterministic quantum state preparation, 2025. URL https://arxiv.org/abs/2510.08267.

Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 576, 2019b. doi: 10.1038/s41586-019-1666-5.

Anthony Ransford. A 98-qubit trapped-ion quantum computer with all-to-all connectivity. *Nature*, 655, 2026. doi: 10.1038/s41586-026-10676-4.

A. G. Radnaev et al. A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout, 2025. URL https://arxiv.org/abs/2408.08288.

Viv Kendon. Decoherence in quantum walks – a review. *Mathematical Structures in Computer Science*, 17(6):1169–1220, December 2007. ISSN 1469-8072. doi: 10.1017/s0960129507006354. URL http://dx.doi.org/10.1017/S0960129507006354.

Viv Kendon. How to compute using quantum walks. *Electronic Proceedings in Theoretical Computer Science*, 315:1–17, 04 2020. doi: 10.4204/EPTCS.315.1. URL https://arxiv.org/abs/2004.01329.

Giuseppe Di Molfetta. Quantum walks, limits and transport equations, 2021. URL https://arxiv.org/abs/2112.11828. arXiv preprint.

Andrew M Childs, Edward Farhi, and Sam Gutmann. An example of the difference between quantum and classical random walks. *Quantum Information Processing*, 1 (1):35–43, 2002.

Neil Shenvi, Julia Kempe, and K. Birgitta Whaley. Quantum random walk search

algorithm. *Phys. Rev. A*, 67:052307, May 2003a. doi: 10.1103/PhysRevA.67.052307. URL https://link.aps.org/doi/10.1103/PhysRevA.67.052307.

Andris Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007. doi: 10.1137/S0097539705447311. URL https://doi.org/10.1137/S0097539705447311.

Andrew M. Childs. Universal computation by quantum walk. *Physical Review Letters*, 102(18), May 2009a. ISSN 1079-7114. doi: 10.1103/physrevlett.102.180501. URL http://dx.doi.org/10.1103/PhysRevLett.102.180501.

Neil B. Lovett, Sally Cooper, Matthew Everitt, Matthew Trevers, and Viv Kendon. Universal quantum computation using the discrete-time quantum walk. *Physical Review A*, 81(4), April 2010. ISSN 1094-1622. doi: 10.1103/physreva.81.042330. URL http://dx.doi.org/10.1103/PhysRevA.81.042330.

Yakir Aharonov, Luiz Davidovich, and Nicim Zagury. Quantum random walks. *Physical Review A*, 48(2):1687–1690, 1993. doi: 10.1103/PhysRevA.48.1687. URL https://link.aps.org/doi/10.1103/PhysRevA.48.1687.

Edward Farhi and Sam Gutmann. Quantum computation and decision trees. *Physical Review A*, 58(2):915–928, August 1998. ISSN 1094-1622. doi: 10.1103/physreva.58.915. URL http://dx.doi.org/10.1103/PhysRevA.58.915.

A. Ambainis, E. Bach, A. Nayak, A. Vishwanath, and J. Watrous. One-dimensional quantum walks. In *Proceedings of the Thirty Third Annual ACM Symposium on Theory of Computing*, STOC '01, page 37–49, New York, NY, USA, 2001. Association for Computing Machinery. ISBN 1581133499. doi: 10.1145/380752.380757. URL https://doi.org/10.1145/380752.380757.

Julia Kempe. Quantum random walks hit exponentially faster, 2002. URL https://arxiv.org/abs/quant-ph/0205083. arXiv preprint.

Neil Shenvi, Julia Kempe, and K. Birgitta Whaley. Quantum random-walk search

algorithm. *Physical Review A*, 67(5), May 2003b. ISSN 1094-1622. doi: 10.1103/physreva.67.052307. URL http://dx.doi.org/10.1103/PhysRevA.67.052307.

Xiaogang Qiang, Shixin Ma, and Haijing Song. Quantum walk computing: Theory, implementation, and application. *Intelligent Computing*, 3, January 2024. ISSN 2771-5892. doi: 10.34133/icomputing.0097. URL http://dx.doi.org/10.34133/icomputing.0097.

Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, STOC03, page 59–68. ACM, 2003. doi: 10.1145/780542.780552. URL http://dx.doi.org/10.1145/780542.780552.

Vivien M Kendon. A random walk approach to quantum algorithms. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1849):3407–3422, October 2006. ISSN 1471-2962. doi: 10.1098/rsta.2006.1901. URL http://dx.doi.org/10.1098/rsta.2006.1901.

Ben Tregenna, Will Flanagan, Rik Maile, and Viv Kendon. Controlling discrete quantum walks: coins and initial states. *New Journal of Physics*, 5(1):83, jul 2003. doi: 10.1088/1367-2630/5/1/383. URL https://doi.org/10.1088/1367-2630/5/1/383.

Manami Yamagishi, Naomichi Hatano, Ken-Ichiro Imura, and Hideaki Obuse. Proposal of multidimensional quantum walks to explore Dirac and Schrödinger systems. *Phys. Rev. A*, 107:042206, Apr 2023. doi: 10.1103/PhysRevA.107.042206. URL https://link.aps.org/doi/10.1103/PhysRevA.107.042206.

Giuseppe Di Molfetta and François Debbasch. Discrete-time quantum walks: Continuous limit and symmetries. *Journal of Mathematical Physics*, 53(12):123302, 2012. ISSN 1089-7658. doi: 10.1063/1.4769171. URL https://doi.org/10.1063/1.4769171.

Pablo Arrighi, Giuseppe Di Molfetta, Iván Márquez-Martín, and Armando Pérez. Dirac equation as a quantum walk over the honeycomb and triangular lattices. *Phys. Rev.*

A, 97:062111, Jun 2018. doi: 10.1103/PhysRevA.97.062111. URL https://link.aps.org/doi/10.1103/PhysRevA.97.062111.

Frederick W. Strauch. Connecting the discrete- and continuous-time quantum walks. *Physical Review A*, 74(3), 2006. ISSN 1094-1622. doi: 10.1103/physreva.74.030301. URL http://dx.doi.org/10.1103/PhysRevA.74.030301.

Pablo Di Molfetta, Giuseppe. Arrighi. A quantum walk with both a continuous-time limit and a continuous-spacetime limit. *Quantum Information Processing*, 19, 2019. doi: 10.1007/s11128-019-2549-2. URL https://doi.org/10.1007/s11128-019-2549-2.

Giuseppe Di Molfetta, Marc Brachet, and Fabrice Debbasch. Quantum walks in artificial electric and gravitational fields. *Physica A: Statistical Mechanics and its Applications*, 397:157–168, March 2014. ISSN 0378-4371. doi: 10.1016/j.physa.2013.11.036. URL http://dx.doi.org/10.1016/j.physa.2013.11.036.

Nicolas Jolly and Giuseppe Di Molfetta. Twisted quantum walks, generalised dirac equation and fermion doubling. *The European Physical Journal D*, 77(5), May 2023. ISSN 1434-6079. doi: 10.1140/epjd/s10053-023-00659-9. URL http://dx.doi.org/10.1140/epjd/s10053-023-00659-9.

Mohamed Hatifi, Giuseppe Di Molfetta, Fabrice Debbasch, and Marc Brachet. Quantum walk hydrodynamics. *Scientific Reports*, 9:2989, 2019. doi: 10.1038/s41598-019-40059-x. URL https://www.nature.com/articles/s41598-019-40059-x.

Giuseppe Di Molfetta and Pablo Arrighi. A quantum walk with both a continuous-time limit and a continuous-spacetime limit. *Quantum Information Processing*, 19 (2):47, February 2019. ISSN 1573-1332. doi: 10.1007/s11128-019-2549-2. URL https://doi.org/10.1007/s11128-019-2549-2.

Giuseppe. di Molfetta and Fabrice. Debbasch. Discrete-time quantum walks: Continuous limit and symmetries. *Journal of Mathematical Physics*, 53(12):123302, 11

2012. ISSN 0022-2488. doi: 10.1063/1.4764876. URL https://doi.org/10.1063/1.4764876.

Giuseppe Di Molfetta. Continuous limits and plasticity in one dimension. In *Quantum Walks, Limits, and Transport Equations*, 2053-2563, pages 4–1 to 4–11. IOP Publishing, 2024. ISBN 978-0-7503-5744-9. doi: 10.1088/978-0-7503-5744-9ch4. URL https://doi.org/10.1088/978-0-7503-5744-9ch4.

Peter J. Love and Bruce M. Boghosian. From Dirac to diffusion: Decoherence in quantum lattice gases. *Quantum Information Processing*, 4:335–354, 2005. URL https://api.semanticscholar.org/CorpusID:1461684.

W. Dür, R. Raussendorf, V. M. Kendon, and H.-J. Briegel. Quantum random walks in optical lattices. *Phys. Rev. A*, 66:052319, 2002. URL https://arxiv.org/quant-ph/abs/0207137.

Viv Kendon and Ben Tregenna. Decoherence can be useful in quantum walks. *Phys. Rev. A*, 67:042315, Apr 2003. doi: 10.1103/PhysRevA.67.042315. URL https://link.aps.org/doi/10.1103/PhysRevA.67.042315.

Almut Beige, Daniel Braun, Ben Tregenna, and Peter L. Knight. Quantum computing using dissipation to remain in a decoherence-free subspace. *Phys. Rev. Lett.*, 85:1762–1765, Aug 2000. doi: 10.1103/PhysRevLett.85.1762. URL https://link.aps.org/doi/10.1103/PhysRevLett.85.1762.

Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. Quantum computation and quantum state engineering driven by dissipation. *Nature Physics*, 5(9):633–636, 2009. doi: 10.1038/nphys1342. URL https://doi.org/10.1038/nphys1342.

S. Attal, F. Petruccione, C. Sabot, and I. Sinayskiy. Open quantum random walks. *Journal of Statistical Physics*, 147(4):832–852, May 2012a. ISSN 1572-9613. doi: 10.1007/s10955-012-0491-0. URL http://dx.doi.org/10.1007/s10955-012-0491-0.

S. Attal, F. Petruccione, and I. Sinayskiy. Open quantum walks on graphs. *Physics Letters A*, 376(18):1545–1548, 2012b. ISSN 0375-9601. doi: https://doi.org/10.1016/

j.physleta.2012.03.040. URL https://www.sciencedirect.com/science/article/pii/S0375960112003453.

Ilya Sinayskiy and Francesco Petruccione. Properties of open quantum walks on $\mathbb{Z}$. *Physica Scripta*, 2012(T151):014077, nov 2012. doi: 10.1088/0031-8949/2012/T151/014077. URL https://doi.org/10.1088/0031-8949/2012/T151/014077.

Masoud Mohseni, Patrick Rebentrost, Seth Lloyd, and Alán Aspuru-Guzik. Environment-assisted quantum walks in photosynthetic energy transfer. *The Journal of Chemical Physics*, 129(17):174106, 2008. doi: 10.1063/1.3002335. URL https://doi.org/10.1063/1.3002335.

Ilya Sinayskiy and Francesco Petruccione. Efficiency of open quantum walk implementation of dissipative quantum computing algorithms. *Quantum Information Processing*, 13(4):805–813, 2014. ISSN 1573-1332. doi: 10.1007/s11128-013-0692-z. URL https://doi.org/10.1007/s11128-013-0692-z.

K. Blum. *Density Matrix Theory And Applications*. Plenum Press, 2012. URL https://books.google.co.uk/books?id=Lu5C0AEACAAJ.

Mostafa Annabestani, Seyed Javad Akhtarshenas, and Mohamad Reza Abolhassani. Decoherence in a one-dimensional quantum walk. *Phys. Rev. A*, 81:032321, Mar 2010. doi: 10.1103/PhysRevA.81.032321. URL https://link.aps.org/doi/10.1103/PhysRevA.81.032321.

Apoorva D. Patel. The quantum density matrix and its many uses: From quantum structure to quantum chaos and noisy simulators. *Journal of the Indian Institute of Science*, 103(2):401–417, April 2023. ISSN 0019-4964. doi: 10.1007/s41745-023-00406-4. URL http://dx.doi.org/10.1007/s41745-023-00406-4.

R Au-Yeung, B Camino, O Rathore, and V Kendon. Quantum algorithms for scientific computing. *Reports on Progress in Physics*, 87(11):116001, October 2024a. ISSN 1361-6633. doi: 10.1088/1361-6633/ad85f0. URL http://dx.doi.org/10.1088/1361-6633/ad85f0.

Omer Rathore, Alastair Basden, Nicholas Chancellor, and Halim Kusumaatmaja. Encoding strategies for quantum enhanced fluid simulations: opportunities and challenges, 2026. URL https://arxiv.org/abs/2604.24694.

Sauro Succi. Numerical solution of the schrödinger equation using discrete kinetic theory. *Phys. Rev. E*, 53:1969–1975, Feb 1996. doi: 10.1103/PhysRevE.53.1969. URL https://link.aps.org/doi/10.1103/PhysRevE.53.1969.

Jeffrey Yepez. An efficient quantum algorithm for the one-dimensional burgers equation, 2002. URL https://arxiv.org/abs/quant-ph/0210092.

René Steijl and George N. Barakos. Parallel evaluation of quantum algorithms for computational fluid dynamics. *Computers Fluids*, 173:22–28, 2018. ISSN 0045-7930. doi: https://doi.org/10.1016/j.compfluid.2018.03.080. URL https://www.sciencedirect.com/science/article/pii/S0045793018301841.

Blaga N. Todorova and René Steijl. Quantum algorithm for the collisionless boltzmann equation. *Journal of Computational Physics*, 409:109347, 2020. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2020.109347. URL https://www.sciencedirect.com/science/article/pii/S0021999120301212.

Ljubomir Budinski. Quantum algorithm for the advection–diffusion equation simulated with the lattice boltzmann method. *Quantum Information Processing*, 20, 02 2021. doi: 10.1007/s11128-021-02996-3.

Budinski Ljubomir. Quantum algorithm for the navier–stokes equations by using the streamfunction-vorticity formulation and the lattice boltzmann method. *International Journal of Quantum Information*, 20(02), February 2022. ISSN 1793-6918. doi: 10.1142/s0219749921500398. URL http://dx.doi.org/10.1142/S0219749921500398.

Wael Itani, Katepalli R. Sreenivasan, and Sauro Succi. Quantum carleman lattice boltzmann simulation of fluids, 2023. URL https://arxiv.org/abs/2301.05762.

Li Xu, Ming Li, Lei Zhang, Hai Sun, and Jun Yao. Improved quantum lattice boltzmann method for advection-diffusion equations with a linear collision model. *Phys. Rev. E*, 111:045305, Apr 2025. doi: 10.1103/PhysRevE.111.045305. URL https://link.aps.org/doi/10.1103/PhysRevE.111.045305.

Monica Lăcătuş and Matthias Möller. Surrogate quantum circuit design for the lattice boltzmann collision operator. *International Journal for Numerical Methods in Engineering*, 127(4):e70286, 2026. doi: https://doi.org/10.1002/nme.70286. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.70286.

Apurva Tiwari, Jason Iaconis, Jezer Jojo, Sayonee Ray, Martin Roetteler, Chris Hill, and Jay Pathak. Algorithmic advances towards a realizable quantum lattice boltzmann method, 2025. URL https://arxiv.org/abs/2504.10870.

Merel A. Schalkers and Matthias Möller. Momentum exchange method for quantum boltzmann methods, 2025. URL https://arxiv.org/abs/2404.17618.

R. Au-Yeung, A.J. Williams, V.M. Kendon, and S.J. Lind. Quantum algorithm for smoothed particle hydrodynamics. *Computer Physics Communications*, 294:108909, January 2024b. ISSN 0010-4655. doi: 10.1016/j.cpc.2023.108909. URL http://dx.doi.org/10.1016/j.cpc.2023.108909.

Cecilia C. López and Juan Pablo Paz. Phase space approach to the study of decoherence in quantum walks. *Phys. Rev. A*, 68:052305, Nov 2003. doi: 10.1103/PhysRevA.68.052305. URL https://link.aps.org/doi/10.1103/PhysRevA.68.052305.

Andrew M. Childs. On the relationship between continuous- and discrete-time quantum walk. *Communications in Mathematical Physics*, 294(2):581–603, October 2009b. ISSN 1432-0916. doi: 10.1007/s00220-009-0930-1. URL http://dx.doi.org/10.1007/s00220-009-0930-1.

Thomas G Wong. Grover search with lackadaisical quantum walks. *Journal of Physics A: Mathematical and Theoretical*, 48(43):435304, oct 2015. doi: 10.1088/1751-8113/48/43/435304. URL https://doi.org/10.1088/1751-8113/48/43/435304.

Thomas G. Wong. Faster search by lackadaisical quantum walk. *Quantum Information Processing*, 17(3), February 2018. ISSN 1573-1332. doi: 10.1007/s11128-018-1840-y. URL http://dx.doi.org/10.1007/s11128-018-1840-y.

Pulak Ranjan Giri and Vladimir Korepin. Lackadaisical quantum walk for spatial search. *Modern Physics Letters A*, 35(08):2050043, December 2019. ISSN 1793-6632. doi: 10.1142/s0217732320500431. URL http://dx.doi.org/10.1142/S0217732320500431.

Sauro Succi. *The Lattice Boltzmann Equation for Complex States of Flowing Matter*. Oxford University Press, Oxford, 2018. ISBN 9780199592357. doi: 10.1093/oso/9780199592357.001.0001. URL https://doi.org/10.1093/oso/9780199592357.001.0001.

Michael Manighalam and Mark Kon. General methods and properties to evaluate continuum limits of the 1D discrete time quantum walk. *Quantum Information Processing*, 19(10):379, October 2020. ISSN 1573-1332. doi: 10.1007/s11128-020-02880-6. URL https://doi.org/10.1007/s11128-020-02880-6.

Dan Li, Michael Mc Gettrick, Wei-Wei Zhang, and Ke-Jia Zhang. One-dimensional lazy quantum walks and occupancy rate*. *Chinese Physics B*, 24(5):050305, mar 2015. doi: 10.1088/1674-1056/24/5/050305. URL https://doi.org/10.1088/1674-1056/24/5/050305.

Garreth Kemp, Ilya Sinayskiy, and Francesco Petruccione. Lazy open quantum walks. *Phys. Rev. A*, 102:012220, Jul 2020. doi: 10.1103/PhysRevA.102.012220. URL https://link.aps.org/doi/10.1103/PhysRevA.102.012220.

Luísa Toledo Tude and Marcos César de Oliveira. Decoherence in the three-state quantum walk. *Physica A: Statistical Mechanics and its Applications*, 605:128012, November 2022. ISSN 0378-4371. doi: 10.1016/j.physa.2022.128012. URL http://dx.doi.org/10.1016/j.physa.2022.128012.

Ilya Sinayskiy and Francesco Petruccione. Microscopic derivation of open quantum

walks. *Phys. Rev. A*, 92:032105, Sep 2015. doi: 10.1103/PhysRevA.92.032105. URL https://link.aps.org/doi/10.1103/PhysRevA.92.032105.

Rhonda. Au-Yeung, Viv. Kendon, and Steven. J. Lind. Quantum smoothed particle hydrodynamics algorithm inspired by quantum walks. *Physics of Fluids*, 37(5), May 2025. ISSN 1089-7666. doi: 10.1063/5.0268240. URL http://dx.doi.org/10.1063/5.0268240.

Wael Itani and Sauro Succi. Analysis of Carlemann linearization of lattice Boltzmann, 2021. URL https://arxiv.org/abs/2111.11327. arXiv preprint.

Amit Saha, Sudhindu Mandal, Debasri Saha, and Amlan Chakrabarti. One-dimensional lazy quantum walk in ternary system. *IEEE Transactions on Quantum Engineering*, PP:1–1, 04 2021. doi: 10.1109/TQE.2021.3074707. URL https://doi.org/10.1109/TQE.2021.3074707.

Robert A. Wilson. A discrete model for Gell-Mann matrices, 2024. URL https://arxiv.org/abs/2401.13000. arXiv preprint.

Heinz-Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 2002. ISBN 9780199213900. URL https://academic.oup.com/book/27757.

Iwo Bialynicki-Birula. Photon wave function. volume 36 of *Progress in Optics*, pages 245–294. Elsevier, 1996. doi: https://doi.org/10.1016/S0079-6638(08)70316-0. URL https://www.sciencedirect.com/science/article/pii/S0079663808703160.

T. A. Brun, H. A. Carteret, and A. Ambainis. Quantum random walks with decoherent coins. *Phys. Rev. A*, 67:032304, Mar 2003b. doi: 10.1103/PhysRevA.67.032304. URL https://link.aps.org/doi/10.1103/PhysRevA.67.032304.

Andrea Alberti, Wolfgang Alt, Reinhard Werner, and Dieter Meschede. Decoherence models for discrete-time quantum walks and their application to neutral atom experiments. *New Journal of Physics*, 16(12):123052, December 2014. ISSN 1367-2630. doi:

10.1088/1367-2630/16/12/123052. URL http://dx.doi.org/10.1088/1367-2630/16/12/123052.

A. Romanelli, R. Siri, G. Abal, A. Auyuanet, and R. Donangelo. Decoherence in the quantum walk on the line. *Physica A: Statistical Mechanics and its Applications*, 347:137–152, 2005. ISSN 0378-4371. doi: https://doi.org/10.1016/j.physa.2004.08.070. URL https://www.sciencedirect.com/science/article/pii/S0378437104011422.

UK Department for Science, Innovation and Technology. UK Compute Roadmap. https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap, 2025. Outlines investment of up to £2 billion through 2030 for advanced and quantum computing infrastructure.

UK Department for Science, Innovation and Technology. National Quantum Strategy. https://www.gov.uk/government/publications/national-quantum-strategy/national-quantum-strategy-accessible-webpage, 2023. Sets out the UK's 10-year plan for quantum technologies, including targets for large-scale, fault-tolerant quantum computers by the early 2030s.

Steph Foulds and Viv Kendon. Lazy quantum walks with native multiqubit gates, 2025. URL https://arxiv.org/abs/2511.21608. arXiv preprint.

Robert Rubinstein and Li-Shi Luo. Theory of the lattice boltzmann equation: Symmetry properties of discrete velocity sets. *Phys. Rev. E*, 77:036709, Mar 2008. doi: 10.1103/PhysRevE.77.036709. URL https://link.aps.org/doi/10.1103/PhysRevE.77.036709.

Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. Sympy: symbolic computing

in python. *PeerJ Computer Science*, 3:e103, January 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103. URL https://doi.org/10.7717/peerj-cs.103.