Essays on Derivative Pricing and Risk Management

Dániel Léber

Doctoral School of Economics and Business Informatics

Supervisors:

Péter Csóka Ph.D.

Institute of Finance

Mihály Ormos Ph.D.

Faculty of Economics (Eötvös Loránd University)

© Dániel Léber

Corvinus University of Budapest

Doctoral School of Economics and Business Informatics

Essays on Derivative Pricing and Risk Management

Doctoral Dissertation

Dániel Léber

Budapest, 2026

Contents

1 Introduction **13**

1.1 Research Questions and Hypotheses 14

2 Efficient High Precision Pseudo-Hermitian Hamiltonian Simulation of Path-Dependent Options **17**

2.1 Introduction 17

2.2 Pseudo-Hermitian Representation of the Black-Scholes-Merton Dynamics 23

2.3 Arrow-Debreu Prices in Hilbert Space 26

2.4 Numerical Simulations of the State-Price Densities and Derived Quantities 30

2.4.1 Comparison of the Estimated Arrow-Debreu State Prices 35

2.4.2 Accuracy Against Analytical Solution of Option Price Valuation and Runtime-Error Diagnosis 39

2.4.3 The Double Knock-Out Runtime-Convergence Comparison 42

2.4.4 Risk Sensitivity Diagnosis and Digital Option Valuation 46

2.5 Conclusion 49

3 Quantum Mechanical Approach to Pricing Multi-Asset Path-Dependent Options **52**

3.1 Introduction 52

3.2 Temporal evolution of multi-asset European options as a quantum state . . 55

3.3 Solution of Multi-Asset European Options and Free Particles in Space . . 59

3.4 Solution of Multi-Asset Double Knock-Out Options and Particles in a Box 64

3.5 Conclusion 73

4 The Sentiment Augmented GARCH-LSTM Hybrid Model for Value-at-Risk Forecasting **75**

4.1 Introduction 75

4.2 Literature Review 79

4.3 Data Sources, Computational Resources and Software 82

4.4 Methodology 83

4.4.1 Realised Volatility 83

4.4.2 Sentiment Indices . . . 84
4.4.3 Measures of Volatility Prediction Performance . . . 84
4.4.4 Granger Causality . . . 85
4.4.5 Generalized AutoRegressive Conditional Heteroskedasticity Model (GARCH) . . . 87
4.4.6 Sentiment Augmented GARCH Model . . . 88
4.4.7 Multilayer Perceptron . . . 89
4.4.8 Long Short-Term Memory (LSTM) . . . 91
4.4.9 Value-at-Risk Estimation . . . 93
4.4.10 Value-at-Risk Backtesting . . . 93
4.5 Predictive Power of Sentiment Indices . . . 96
4.5.1 Linear Granger Causality . . . 96
4.5.2 Nonlinear Granger Causality . . . 99
4.6 Performance Evaluation and Backtesting of Predictive Models . . . 100
4.6.1 Hyperparameters of Neural Networks . . . 101
4.6.2 Results on Randomly Selected Stocks . . . 102
4.6.3 Aggregated Results of Stocks in the S&P 500 Index . . . 110
4.7 Conclusion and Further Suggestions . . . 113

5 Conclusion and Future Research **116**

References **120**

Publications Underlying the Dissertation **131**

List of Figures

2.1 Estimation of the plain-vanilla Arrow-Debreu state prices by different algorithms . . . 38
2.2 Double knock-out Arrow-Debreu state prices . . . 40
2.3 Runtime versus absolute error against the analytical Black-Scholes-Merton price . . . 42
2.4 Double Knock-Out price convergence as a function of the sensitivity parameter . . . 43
2.5 Double Knock-Out price convergence as a function of the sensitivity parameter without Monte Carlo . . . 44
2.6 Runtime as a function of sensitivity for double knock-out pricing . . . 45
2.7 Runtime as a function of sensitivity for double knock-out pricing without Monte Carlo and Crank-Nicolson algorithms . . . 46

3.1 The surface plot of the two-asset double knock-out worst-of rainbow option price as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.05. . . . 70
3.2 The heatmap of the two-asset double knock-out worst-of rainbow option price as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.05. . . . 70
3.3 The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.6. . . . 71
3.4 The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be -0.6. . . . 71
3.5 The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. S_1 has limits prices between \$100 and \$140, while S_2 has limit prices between \$90 and \$130. The correlation between the two underlying assets is assumed to be 0.6. . . . 72

4.1 The schematic architecture of an arbitrary neural network with 4 inputs, one hidden layer with 3 neurons and an output layer with 1 neuron. 90
4.2 The schematic architecture of one LSTM cell. 92
4.3 Forecasts of the 2.5% Value-at-Risk for Alphabet Inc. stock in the test dataset as predicted by multiple models. 104
4.4 Forecasts of the 2.5% Value-at-Risk for International Business Machines Corporation stock in the test dataset as predicted by multiple models. . . . 105
4.5 Forecasts of the 2.5% Value-at-Risk for Tesla Inc. stock in the test dataset as predicted by multiple models. 106
4.6 The volatility prediction capability of the given models compared to that of the GARCH model 114

List of Tables

2.1 Model and numerical sensitivity parameters 37

2.2 Discrepancy measures of the Arrow-Debreu state prices for the plain vanilla option, the reference is the analytical Black-Scholes-Merton state price density 38

2.3 Discrepancy measures of the Arrow–Debreu state prices in the case of the double knock-out option, the reference is the Hamiltonian spectral estimation 40

2.4 Plain-vanilla European call price comparison against the analytical Black-Scholes-Merton value 41

2.5 Plain-vanilla delta and digital-call estimations from the state-price densities, the relative errors are estimated against the analytical Black–Scholes–Merton value 48

2.6 Double knock-out survival probabilities, delta and digital-call estimations, the relative errors are estimated against the Hamiltonian spectral estimation 49

4.1 The proportion of companies in the S&P 500 whose variance and standardised sentiment indices are together stationary at the 5% significance level 97

4.2 Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 1% significance level across different lags. 98

4.3 Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 5% significance level across different lags. 98

4.4 Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 10% significance level across different lags. 99

4.5 Multilayer Perceptron's Hyperparameters of Non-linear Granger Causality Test 99

4.6 Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 1% significance level across different lags. 101
4.7 Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 5% significance level across different lags. 101
4.8 Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 10% significance level across different lags. 102
4.9 The hyperparameters of neural networks for volatility forecasting 103
4.10 Metrics measuring the forecasting ability of volatility predicted by different models 107
4.11 Results of different statistical tests at the 5% significance level 108
4.12 Metrics measuring the expected loss of Value-at-Risk predictions predicted by different models 109
4.13 The proportion of stocks for which the realised volatility is better predicted by the given model than by the other models based on the specified metric. 111
4.14 The proportion of stocks for which the loss is smaller than the given model compared to the other model based on the specified metric. 112
4.15 The proportion of stocks for which the null hypothesis of the given test statistics was rejected more or less frequently compared to the other model. 113

Acknowledgements

This dissertation grew out of many years of study, discussion, and writing. During that time, I was fortunate to receive the support, guidance, and encouragement of many people. Although I cannot mention everyone individually, I would like to express my sincere gratitude to those whose help and influence have been especially important to me.

First and foremost, I am deeply grateful to Péter Csóka. His steady mentorship and guidance, both professional and personal, have been invaluable. I am equally grateful to Mihály Ormos for his professional guidance and mentorship. His generosity with his time and his willingness to engage in long scholarly conversations greatly enriched my research. I also wish to thank Antal Jakovác, who, although not formally my supervisor, offered many valuable insights, inspired new ideas, and helped me overcome several difficulties along the way.

I would also like to express my gratitude to the giants on whose work I was able to build, including Schrödinger, Arrow, Linetsky, Baaquie, Campbell, Cont, Hamilton, Hull, Lánczos, Rubinstein, Wilmott, Tetlock, and many others. Their pioneering ideas laid the foundations for this dissertation. Without their work, none of the results presented here would have been possible.

Many other teachers and friends have also contributed to this work. I am grateful to László Udvardi for always being willing to answer my questions on quantum mechanics. I am especially grateful to Bálint Juhász, my external supervisor during my master's thesis, whose excellent guidance first introduced me to the fascinating subjects that later became the foundation of this dissertation. I also gratefully acknowledge the important and indispensable assistance of Zsolt Bihary, Tamás Sándor Biró, János Asbóth, András Fülöp, Zsolt Darvas, and Zsolt Illési. My co-author, Balázs Egyed, made a highly valuable contribution to our joint paper already during his BSc studies, for which I am sincerely grateful.

I would also like to thank my supervisors and leaders in my professional environment, Péter Tóth, Dániel Homolya, and Gyula Lencsés, for their flexibility and support, which made it possible for me to balance my professional responsibilities with the demands of doctoral research.

My family has been a constant source of strength and support throughout this process. I am deeply grateful to my mother, my sister Anna, my grandmother, and my grandfather for their love and unwavering support. I am especially grateful to my girlfriend, Nóra, for

her patience, encouragement, and for standing by me through this long and demanding period.

Finally, I would like to express my gratitude to all those who have helped me in ways both great and small. Their kindness, patience, and encouragement have carried me through this challenging journey.

Chapter 1

Introduction

For both risk managers and portfolio managers, the continuous emergence of increasingly complex financial instruments poses a persistent challenge. While financial innovation expands the set of tools available for hedging, risk transfer, and portfolio construction, it also makes valuation and risk measurement more demanding. This is particularly true for exotic derivatives whose payoffs depend on the path of one or more underlying assets. Lookback, Asian, barrier, and rainbow options are well-known examples of such contracts, and their valuation typically requires methods that go beyond the standard closed-form results available for plain-vanilla claims. From a practical perspective, the problem is not merely to obtain derivative prices, but to obtain prices and risk measures with sufficient precision and speed for trading, hedging, portfolio rebalancing, and risk monitoring. For complex path-dependent and multi-asset products, the relevant dynamics are often too rich for simple analytical treatment, while standard numerical procedures may become costly when high accuracy is required. This creates a natural demand for methodologies that can deliver accurate valuation and practically useful risk measures in a computationally efficient way (Hull and Basu (2016), Hull (2023), Wilmott (2013)).

The challenge is compounded by the empirical properties of financial markets themselves. Asset returns are known to exhibit stylised facts such as volatility clustering, heavy tails, leverage effects, and intermittency, all of which require complex modelling and forecasting. At the same time, market dynamics are influenced not only by past prices but also by information flows and investor sentiment reflected in news, corporate language, internet message boards, and social media. This is particularly important because financial institutions must forecast risk measures with a high degree of accuracy, while also meeting increasingly stringent regulatory requirements. In this context, volatility forecasts play a central role in Value-at-Risk forecasting used for internal risk control, capital allocation, and regulatory backtesting (Cont (2001), Basel Committee on Banking Supervision (2013, 2019), Tetlock (2007), Kranefuss and Johnson (2021), Deveikyte et al. (2022)).

1.1 Research Questions and Hypotheses

These challenges motivate the present dissertation, which examines how analytical, numerical, and data-driven methods can contribute to high-precision valuation and risk measurement of path-dependent derivatives, as well as to sentiment-based volatility forecasting. More specifically, the dissertation seeks answers to the following questions:

- How can pseudo-Hermitian Hamiltonian simulations be turned into practical high-precision algorithms for pricing and risk measurement of path-dependent options?
- How can this methodology be generalised to the case of multiple correlated underlying assets in order to price multi-asset path-dependent derivatives?
- How can volatility and Value-at-Risk forecasting be improved by combining GARCH-type models, recurrent neural networks, and sentiment information extracted from news and social media?

The first two research questions give rise to one methodological hypothesis each, while the third research question is examined through two empirical hypotheses concerning the predictive content of sentiment information and its contribution to forecasting performance. Accordingly, the dissertation formulates the following hypotheses:

H1. It is hypothesised that, within the examined single-asset Black–Scholes–Merton settings, the pseudo-Hermitian Hamiltonian spectral method will attain numerical accuracy at least comparable to that of the Crank–Nicolson, binomial-tree, and Monte Carlo benchmark methods for option valuation, with particularly strong performance for quantities that depend sensitively on the reconstructed Arrow–Debreu state-price density. On sufficiently fine grids, the shift-invert Lánczos approximation is expected to retain essentially the accuracy of the full-spectral solution while exhibiting a more favourable empirical runtime profile.

H2. It is hypothesised that the pseudo-Hermitian similarity transformation considered for the single-asset problem in H1 can be generalised into a Hermitian framework for pricing simple and path-dependent options on multiple correlated underlying assets, thereby enabling the derivation of semi-analytical pricing solutions and efficient approximations for path-dependent multi-asset derivatives.

H3. It is hypothesised that positive and negative sentiment indices derived from news and Twitter/X contain statistically significant predictive information for the conditional variance of daily returns among the S&P 500 stocks included in the empirical sample.

H4. It is hypothesised that augmenting the GARCH and GARCH–LSTM-based models with sentiment information will reduce out-of-sample realised-volatility forecast errors for stocks included in the empirical sample relative to the corresponding models without sentiment. The resulting improvements in volatility forecasts are, in turn, expected to improve the performance of Value-at-Risk forecasts, as assessed using standard Value-at-Risk backtests.

The three substantive chapters examine these hypotheses from closely connected but distinct angles. Chapter 2 examines H1 by developing a practical pseudo-Hermitian Hamiltonian simulation framework for single-asset path-dependent derivatives and evaluating it against established benchmark methods. This study builds on the results reported in my earlier thesis Léber (2021), which examined the pricing of path-dependent options through the general solution of the Schrödinger equation and analysed the convergence of the proposed method. The present work extends that research through further methodological developments and comparisons with several classical pricing methods. Chapter 3 examines H2 by generalising the pseudo-Hermitian framework to several correlated underlying assets and applying it to European terminal claims and multi-asset barrier options. Chapter 4 examines H3 and H4 by testing the linear and nonlinear predictive content of news- and social-media-based sentiment indices and by evaluating whether sentiment augmentation improves volatility and Value-at-Risk forecasts produced by GARCH and GARCH–LSTM-based models.

Beyond these specific contributions, the dissertation's added value lies in bringing together two strands of research that are often studied separately. The first two chapters show how a framework inspired by quantum mechanics can be developed from a practical high-precision implementation for single-asset derivatives into a generalised methodology applicable to multiple correlated assets. The third chapter shows how traditional econometric models and modern neural network architectures can be combined with sentiment information to forecast volatility, thereby capturing nonlinear patterns that are difficult to model with conventional methods.

The dissertation is organised into four subsequent chapters. Chapter 2, Efficient High Precision Pseudo-Hermitian Hamiltonian Simulation of Path-Dependent Options, develops the numerical pseudo-Hermitian Hamiltonian framework and benchmarks it against standard pricing methods. Plain-vanilla options are used to assess numerical accuracy against the analytical Black-Scholes-Merton benchmark, while a double knock-out option serves as a representative example of a path-dependent contract. Chapter 3, Quantum Mechanical Approach to Pricing Multi-Asset Path-Dependent Options, continues this methodological line by generalising the framework to multiple correlated underlying assets and deriving solutions for general multi-asset European options and a representative

path-dependent rainbow option. Chapter 4, The Sentiment Augmented GARCH-LSTM Hybrid Model for Value-at-Risk Forecasting, investigates the linear and nonlinear predictive power of sentiment indices and proposes a sentiment-augmented econometric and hybrid forecasting model for volatility and Value-at-Risk forecasting. Finally, Chapter 5 concludes the results and outlines possible directions for future research.

Chapter 2

Efficient High Precision Pseudo-Hermitian Hamiltonian Simulation of Path-Dependent Options

This chapter develops a high-precision pseudo-Hermitian Hamiltonian simulation framework for pricing single-asset path-dependent options and computing associated risk measures. By representing Arrow–Debreu state prices in a metric Hilbert space and implementing a Hamiltonian spectral simulation and its shift-invert Lánczos version, we show how quantum-mechanical-inspired simulations can deliver accurate option prices, deltas, digital option prices, and knock-out probabilities with a relatively fast runtime. This work lays the theoretical and numerical foundation for the next chapter, in which we generalise the framework to multiple correlated assets.

2.1 Introduction

Recent work in derivative pricing encompasses a broad spectrum of theoretical innovations and numerical methods. Although closed-form solutions exist for certain standard derivatives, such as plain-vanilla European-style options, the pricing and sensitivity analysis of path-dependent instruments, such as limit-price options, often rely on computationally expensive numerical procedures. Classical approaches based on Monte Carlo simulation, lattice trees, and finite-difference methods are well established and widely used in the finance literature. However, in settings where both high precision and fast computational turnaround are required, these standard methods are not always the most suitable, since their accuracy may improve only at the cost of substantial additional runtime or increased numerical complexity. This issue becomes particularly important when derivatives are used not only for indicative valuation but also for hedging, rebalancing, and risk control, because even small local pricing errors may propagate into materially

different risk sensitivities or knock-out event probabilities (Wilmott (2013), Hull and Basu (2016), Hull (2023)). In response to these challenges, derivative pricing theory is evolved into a rich methodological landscape in which tools from spectral theory, operator methods, and quantum-mechanically inspired models complement classical valuation approaches. This extensive literature includes pseudo-Hermitian reformulations, supersymmetric frameworks, Feynman path-integral methods, and a variety of alternative numerical schemes, all aimed at developing tractable and accurate solution frameworks for pricing problems that go beyond the standard solutions (Jana and Roy (2012, 2011), Linetsky (1997), Araneda and Villena (2021), Baaquie (2007)).

One of the earliest explicit uses of the Black-Scholes-Merton Schrödinger reformulation is due to Haven (2002), who introduces an arbitrage-free potential and interprets the derivative value as a quantum state. Building on endogenous-arbitrage formulations, Contreras et al. (2010) propose a so-called Black Scholes Schrödinger framework, in which arbitrage effects enter through an additional potential term and induce a Hamiltonian description of pricing dynamics. A related but conceptually distinct line of research retains the Schrödinger correspondence while focusing directly on the operator reformulation of the original Black-Scholes-Merton generator and on the analytical and computational implications of its Hamiltonian representation. Vukovic (2015) explores the mathematical correspondence between the Schrödinger and Black-Scholes-Merton partial differential equations through their Hamiltonian representations, showing that the quantum-mechanical Hamiltonian is Hermitian, whereas the Hamiltonian governing Black-Scholes-Merton dynamics is non-Hermitian. Jana and Roy (2012) refine this point by showing that the Black-Scholes-Merton Hamiltonian is nevertheless pseudo-Hermitian, in the sense that one can construct a η metric operator under which the pricing generator becomes similar to a genuinely Hermitian operator, thereby providing a consistent inner product setting for eigenfunction-based representations with eigenvalues that are typically real in the financial context. The authors further emphasise that pseudo-Hermiticity modifies the appropriate completeness and inner-product relations, which in turn yield an explicit eigenfunction-based representation of the option values within the η-weighted framework. In subsequent work, Jana and Roy (2011) exploit supersymmetric methods to construct isospectral partner Hamiltonians for the Black-Scholes-Merton operator. The authors illustrate the procedure using a path-dependent double knock-out barrier option, represented by an infinite square-well potential, and derive partner operators whose eigenfunctions satisfy the same boundary conditions and whose spectra coincide, up to the usual supersymmetry-based exception of a possible missing ground state. Related extensions based on Darboux transformations and supersymmetric constructions are discussed by Yeşiltaş (2023).

Complementing the Schrödinger-based reformulations, Feynman's path-integral formalism represents another important quantum-inspired approach in quantitative finance. In this framework, derivative values are expressed through propagators obtained by aggregating contributions from ensembles of admissible trajectories. Linetsky (1997) presents path-integral valuation as both a unifying conceptual framework and a practical computational methodology for path-dependent derivatives, showing that Gaussian cases may admit closed-form semiclassical expressions, whereas more general settings can be handled numerically via quadrature or Monte Carlo evaluation along the classical path. The methodology is illustrated on canonical path-dependent contracts such as weighted geometric Asian options, floating barrier options, and barrier options with ladder-like barriers. In a closely related Hamiltonian formulation, Baaquie et al. (2004) incorporate payoff constraints as potential terms and evaluate the resulting path integrals numerically, while Baaquie (2007) further develops this perspective across a broader range of financial applications. Using a semiclassical path-integral approximation tailored to local-volatility dynamics, Araneda and Villena (2021) also show that such methods can deliver accurate and computationally efficient pricing in the constant elasticity of variance (CEV) setting.

Both the Schrödinger-mapping and path-integral viewpoints foreground the role of linear operator structure in derivative valuation. In Markovian models, prices evolve under a pricing semigroup generated by an infinitesimal operator, and Linetsky (2007) formulates the associated boundary-value problems in terms of the generator's spectral analysis, emphasising that when the state-price density is unavailable in closed form, valuation and hedging on large grids of states and maturities become computationally demanding. Within this operator framework, derivative values are obtained by propagating payoffs through the pricing semigroup, so that valuation reduces to characterising the eigenvalues and eigenfunctions of the generator in close analogy with diagonalising a Hamiltonian. In the case of Black-Scholes-Merton diffusion models, the Liouville transformation converts the associated Sturm-Liouville eigenproblem into a one-dimensional Schrödinger normal form, in which an effective potential encodes the model dynamics together with the relevant boundary conditions. Building on this spectral perspective, Davydov and Linetsky (2003) introduce eigensecurities, defined as eigenfunctions of the present-value operator, which diagonalise the operator and yield a spectral representation of the state-price density. For regular Sturm-Liouville problems, they show that high accuracy typically requires only a modest number of terms, that convergence improves at longer maturities because higher modes are exponentially suppressed, and that the resulting eigenpairs can be reused across strikes, maturities, and hedge calculations.

Linetsky (2004b,a) extends spectral methods to path-dependent derivatives such as continuously sampled Asian and lookback options, where the relevant payoff distributions

are generally not available in closed form and standard analytical approximations often lack reliable error control. In these settings, pricing is reduced to the spectral analysis of the underlying diffusion generator, leading to semi-analytical valuation formulas expressed through eigenfunction expansions rather than purely classical simulation-based procedures. For arithmetic Asian options, the approach yields tractable spectral representations when the average is not lognormally distributed under geometric Brownian motion. The author further remarks that the same spectral structure admits alternative interpretations, including a mapping to a Schrödinger equation with a Morse-type potential and a connection to Brownian motion on the hyperbolic plane in the presence of a magnetic field. For lookback-style contracts, the method exploits the connection between running extrema and first-hitting-time problems. A central practical advantage of this methodology is that risk sensitivities can be obtained by analytical differentiation of the expansion without additional loss of accuracy. The author also emphasises that convergence improves at longer maturities because higher modes are increasingly suppressed, making spectral methods particularly attractive for long-dated path-dependent contracts. These contributions also help clarify the distinction between semi-analytical spectral pricing and the numerical approach pursued in the present study. In the spectral literature surveyed by Linetsky, the emphasis is explicitly on those symmetric Markov models for which the underlying Sturm–Liouville problem can be solved in closed form, typically after reduction to a Schrödinger equation with a known special-function potential. By contrast, outside that analytically tractable class, the state-price densities are generally not available in explicit form, and valuation must instead rely on numerical procedures (Linetsky (2007)).

Recent studies indicate that spectral methodology extends beyond closed-form eigenfunction expansions to include high-order numerical discretisations and semi-analytical transform techniques. For the time-fractional Black-Scholes-Merton equation with a Caputo time derivative, An et al. (2021) develop a high-order space-time spectral method based on a variational formulation, combining Jacobi-polynomial temporal expansions with Fourier-like spatial basis functions, and establish stability and convergence results supported by numerical experiments. In a related direction, after recasting the pricing partial differential equation in heat-equation form, Carr and Itkin (2021) derive semi-analytical pricing formulas for barrier and American options under a time-dependent Ornstein-Uhlenbeck process using a generalised integral transform, obtaining stable one-dimensional integral representations with a Fourier-series structure involving Jacobi theta functions. Their method is validated against Crank-Nicolson benchmarks and is particularly attractive when prices or Greeks are required across many strikes and maturities, especially in challenging configurations where finite-difference schemes may lose accuracy near the barrier. More recently, Mollahasani (2023) proposes a hybrid scheme for the time-fractional Black-Scholes-Merton equation that uses a second-order finite-difference

discretisation in time and a spatial spectral approximation based on fractional-order shifted Hahn functions, with numerical tests indicating improved accuracy and efficiency relative to competing methods.

Recent years have seen a growing body of work on quantum computing for finance that develops quantum hardware oriented algorithms for derivative pricing by expressing derivative valuation in operator form suited to implementation on quantum computers, while highlighting a central implementation challenge in that quantum processors naturally realise unitary evolutions generated by Hermitian operators whereas Black-Scholes-Merton type generators are generally non Hermitian and therefore correspond to intrinsically non unitary dynamics (Gonzalez-Conde et al. (2023)). Herman et al. (2023) provide a broad review of quantum computing for financial applications with particular emphasis on stochastic modelling, optimisation, and machine learning, outlining standard classical workflows used in practice and clarifying both the prospective advantages and the current limitations of quantum computational methods. In a representative option pricing application, Stamatopoulos et al. (2020) develop an end-to-end quantum gate-based pricing circuit built around quantum amplitude estimation for vanilla, multi-asset, and path-dependent payoffs, including barrier options. By contrast, Gonzalez-Conde et al. (2023) present also an end-to-end gate-based quantum algorithm that maps the Black-Scholes-Merton equation to Schrödinger-type evolution and implements Hamiltonian simulation on a quantum computer. The resulting non-Hermitian propagator is handled through unitary dilation in an enlarged Hilbert space with a single ancilla qubit, and the option value estimation can be quadratically accelerated by quantum amplitude estimation. More recently, Kumar and Wilmott (2025) extend quantum imaginary-time methods to general non-Hermitian Hamiltonians, indicating that future quantum approaches to derivative pricing may proceed through direct simulation of non-unitary pricing dynamics.

Motivated by the pseudo-Hermitian formulation of the Black-Scholes-Merton operator developed by Jana and Roy (2012, 2011), we cast the associated similarity transform as a practical numerical methodology for pricing vanilla and path-dependent derivatives, along with risk sensitivities. The central idea is to work with the pseudo-Hermitian transform of the Black-Scholes-Merton Hamiltonian, estimate the corresponding Arrow-Debreu prices numerically via a spectral decomposition, inspired by Davydov and Linetsky (2003)'s eigensecurities, and then use state prices to compute derivative values and express sensitivities from a risk-management perspective.

In particular, the pseudo-Hermitian similarity transform used here yields a self-adjoint operator on a Hilbert space, which allows the method to leverage robust numerical Hermitian eigensolvers and effective Krylov subspace algorithms. From a computational

perspective, this reformulation is attractive because it delivers two complementary advantages that simplify and accelerate both propagation and post-processing. First, the discretised Hamiltonian is sparse and Hermitian, with an orthogonal eigenbasis, which supports stable spectral propagation and enables the systematic reuse of computed eigenpairs to obtain option values, sensitivities, and barrier survival probabilities within a single computational framework. Second, Hermiticity makes diagonalisation substantially cheaper, since the dominant eigenpairs can be extracted efficiently via Krylov-subspace projections, in contrast to non-Hermitian settings, where more general and typically more computationally demanding procedures, such as the Arnoldi algorithm, are required. Indeed, since higher modes are exponentially suppressed as time to maturity increases, their contribution becomes negligible beyond a moderate index (Davydov and Linetsky (2003), Linetsky (2007)). We therefore propose a shift–invert Lánczos Krylov strategy that targets only the eigenpairs required to satisfy a prescribed spectral-tail tolerance, avoiding full diagonalisation while retaining the stability benefits of an orthogonal eigenbasis. In contrast to eigenfunction expansions based on Liouville-transformed Sturm–Liouville problems, which are generally confined to a narrow class of analytically tractable models and often require model-specific derivations and numerical procedures, the proposed pseudo-Hermitian approach provides a more general, broadly applicable framework. In typical implementations, its main computational burden is limited to the numerical evaluation of a relatively small number of Hermitian eigenpairs (Linetsky (2007, 2004b)). To quantify accuracy, convergence, and runtime in practice, we benchmark our results against standard numerical pricing algorithms, including Cox–Ross–Rubinstein binomial trees, Crank–Nicolson finite-difference solvers, and Monte Carlo simulations. We also compare not only option values but also derived sensitivities, including state-price densities, deltas, and barrier-event probabilities, which provide stricter validation than price-only comparisons. This broader numerical Hamiltonian framework is practically relevant because many financially important valuation problems lie outside the class of analytically solvable spectra. In the general Markovian setting, pricing expectations are typically unavailable in explicit analytical form, whereas extended specifications, such as Markovian HJM models with accrued-variance states or GARCH stochastic-volatility option models, require numerical simulations (Linetsky (2007), Lin and Ritchken (2006)).

To illustrate the effectiveness of the proposed algorithms, we first revisit the pseudo-Hermitian reformulation of the Black-Scholes-Merton framework and show how the associated similarity transform yields a Hermitian Hamiltonian representation. Building on this representation, we derive the Arrow-Debreu state prices by solving the associated time-independent Schrödinger equation. We then introduce the classical benchmark methodologies used in the numerical comparisons to compute Arrow-Debreu prices and derived values. In accordance with the established methodology, the numerical presenta-

tion is organised around two complementary case studies. We begin with a plain-vanilla European call, using the Black-Scholes-Merton closed form to assess approximation error and runtime across the competing numerical approaches. We then turn to a genuinely path-dependent example, a double knock-out barrier call option, and apply the same methodology to pricing and sensitivity analysis. In both settings, we compare the proposed algorithms with the classical numerical baselines in terms of the inferred Arrow-Debreu state-price densities, option values, convergence behaviour, computational cost, and the resulting state-price-derived risk measures, including sensitivities and knock-out probabilities.

2.2 Pseudo-Hermitian Representation of the Black-Scholes-Merton Dynamics

Representing an option-value function as a quantum state requires recasting the Euclidean-space formulation conventionally used in financial mathematics within the Hilbert-space framework of quantum mechanics (Vukovic (2015), Baaquie (2007)). Let $g(S,t)$ denote the value of the option when the price of the underlying asset (the stock price) is S at valuation time t. In the Hilbert space, all relevant information about an option can be represented by an infinite-dimensional complex vector, the so-called ket vector illustrated in Equation (2.1).

$$|g(S,t)\rangle = \begin{bmatrix} \vdots \\ g(S-2\Delta S,t) \\ g(S-\Delta S,t) \\ g(S,t) \\ g(S+\Delta S,t) \\ \vdots \end{bmatrix} \tag{2.1}$$

where, $|g(S,t)\rangle$ represents the quantum state of the option (Baaquie (2007)).

In financial mathematics, the Black-Scholes-Merton partial differential Equation 2.2 serves as a tool to determine the value of a European option over an infinitesimally small time interval, effectively capturing the dynamics of the option's value. Analogously to the fundamental laws governing physical dynamics, the equation governing the time evolution of an option within the Black-Scholes-Merton framework is inherently deterministic. In quantum mechanics, the Schrödinger differential Equation 2.3 governs the time evolution of a quantum state $|\Psi\rangle$ between two specified temporal points. According to Vukovic (2015) and Baaquie (2007), the corresponding differential equation in the Euclidean space can be reformulated in the Hilbert space, as shown in Equation 2.4, where the

special operator $\hat{H}_{BS}$ plays a role analogous to that of the Hamiltonian operator $\hat{H}$ in the Schrödinger equation. To align unitary Schrödinger evolution with pricing dynamics, we use the usual imaginary-time (Wick-rotated) representation, in which the factor i is removed. Economically, this replaces physical phase oscillations by the real diffusion-and-discount semigroup that propagates the terminal payoff back to its present value (Wróblewski (2014), Contreras et al. (2010)). For the financial application, we work in natural units and set $\hbar = 1$ throughout. The risk associated with the option's value is intrinsically related to the behaviour of the logarithm of the underlying stock price S. Therefore, it is advantageous to express the operator $\hat{H}_{BS}$ in terms of the x logarithm of the stock price, as illustrated in Equation 2.6. In addition, the Black-Scholes-Merton partial differential equation is posed with a terminal condition at maturity $t = T$ rather than an initial condition, as is typically specified in the Schrödinger initial-value formulation. It is therefore convenient to introduce the time to maturity variable $\tau = T - t$, so that the Hilbert space evolution (2.4) can be written as a backward time evolution with τ in the Schrödinger form (2.7).

$$\frac{\partial g(S,t)}{\partial t} = \left(r - rS\frac{\partial}{\partial S} - \frac{1}{2}\sigma^2 S^2 \frac{\partial^2}{\partial S^2}\right) g(S,t) \tag{2.2}$$

$$i\hbar \frac{\partial |\Psi\rangle}{\partial t} = \hat{H}|\Psi\rangle \tag{2.3}$$

$$\frac{\partial |g(S,t)\rangle}{\partial t} = \hat{H}_{BS}|g(S,t)\rangle \tag{2.4}$$

$$\hat{H}_{BS} = \left(r - rS\frac{\partial}{\partial S} - \frac{1}{2}\sigma^2 S^2 \frac{\partial^2}{\partial S^2}\right) \tag{2.5}$$

$$\hat{H}_{BS} = \left(r - \left(r - \frac{\sigma^2}{2}\right)\frac{d}{dx} - \frac{\sigma^2}{2}\frac{d^2}{dx^2}\right) \tag{2.6}$$

$$\frac{\partial}{\partial \tau}|g(S,\tau)\rangle = -\hat{H}_{BS}\,|g(S,\tau)\rangle \tag{2.7}$$

where, $g(S,t)$ represents the function of the option value, r is the constant risk-free interest rate, σ is the constant volatility of the underlying stock, $|\psi\rangle$ represents a specific quantum state, $|g(S,t)\rangle$ represents the quantum state of an option, x represents the logarithm of the stock price, τ is the time to maturity, $\hat{H}_{BS}$ is a Hamiltonian operator satisfying the Black-Scholes-Merton transformation (Baaquie (2007), Jana and Roy (2012)).

In quantum mechanics, the result of a measurement corresponds to real eigenvalues only if the associated operator is Hermitian. However, as noted by Vukovic (2015) and Baaquie (2007), the Black-Scholes-Merton Hamiltonian is non-Hermitian, which

allows the appearance of complex eigenvalues, although the non-Hermiticity does not automatically imply complex spectra. A broad framework underlying the emergence of real spectra in certain non-Hermitian settings is pseudo-Hermiticity, which means that an operator $\hat{H}$ is called η-pseudo-Hermitian if there exists an invertible Hermitian operator $\hat{\eta}$ such that shown in Equation (2.8) (Mostafazadeh (2002)). Jana and Roy (2012) demonstrate that the Black-Scholes-Merton Hamiltonian satisfies the $\hat{\eta}$-pseudo-Hermiticity condition. Using a similarity transformation (2.9), the operator $\hat{\rho}$ defined in (2.10) maps the non-Hermitian Hamiltonian $\hat{H}_{BS}$ to the Hermitian Hamiltonian $\hat{h}_{BS}$ (2.11). The metric operator $\hat{\eta}$, defined in Equation (2.12), ensures the η-pseudo Hermiticity of $\hat{H}_{BS}$ by modifying the scalar product and completeness relations in non-Hermitian systems. Economically, $\hat{\rho}$ is an invertible state-dependent reweighting that cancels the first-order risk-neutral log-return drift in the transformed representation. The drift's directional effect is retained in the endpoint factors $\rho(x_t)^{-1}\rho(x_T)$, while the constant term of $\hat{h}_{BS}$ carries discounting together with the drift correction, so the original Black-Scholes-Merton valuation is unchanged. One special path-dependent case is the double knock-out option, where L and U denote the lower and upper barrier levels of the underlying asset price, respectively. Within the pseudo-Hermitian similarity-transformed framework, the double knock-out option's barrier condition is naturally incorporated into the Hermitian problem by an infinite square-well potential, with survival restricted to paths that remain inside the corresponding interval $(\ln L, \ln U)$ up to maturity (Jana and Roy (2011)). Accordingly, the potential term in the transformed Black-Scholes-Merton Hamiltonian is infinite outside the barrier levels, as stated in Equation (2.13), which is equivalent to imposing absorbing Dirichlet boundary conditions at the barriers. The unbounded log-price coordinate takes values in the ordinary real line $\mathbb{R} = (-\infty, \infty)$, hence the symbols $\pm\infty$ in integration limits denote the unbounded extent of the domain, not points adjoined to it. In the double knock-out case, $+\infty$ denotes the hard-wall limit $V_0 \to +\infty$, whose operator-theoretic realisation is the self-adjoint Dirichlet Hamiltonian on $L^2((\ln L, \ln U))$; the barrier points belong to the boundary rather than to the open configuration domain. The same pseudo-Hermitian framework can also be extended to multiple correlated underlyings, as Léber (2025) provides an explicit multi-asset generalisation of the pseudo-Hermitian Black-Scholes-Merton Hamiltonian and applies the resulting operator setting to pricing path-dependent multi-asset contracts.

$$\hat{H}_{BS}^{\dagger} = \hat{\eta}\hat{H}_{BS}\hat{\eta}^{-1}. \tag{2.8}$$

$$\hat{\rho}\hat{H}_{BS}\hat{\rho}^{-1} = \hat{h}_{BS}, \tag{2.9}$$

$$\hat{\rho} = \exp\left[-\left(\frac{1}{2} - \frac{r}{\sigma^2}\right)x\right], \tag{2.10}$$

$$\hat{h}_{BS} = -\frac{\sigma^2}{2}\frac{d^2}{dx^2} + \frac{1}{2\sigma^2}\left(\frac{\sigma^2}{2} + r\right)^2, \tag{2.11}$$

$$\hat{\eta} = \hat{\rho}^2 = \exp\left[-\left(1 - \frac{2r}{\sigma^2}\right)x\right], \tag{2.12}$$

$$\hat{h}_{\mathrm{DKO}} = -\frac{\sigma^2}{2}\frac{d^2}{dx^2} + V_{\mathrm{DKO}}(x), \qquad V_{\mathrm{DKO}}(x) = \begin{cases} \frac{1}{2\sigma^2}\left(\frac{\sigma^2}{2} + r\right)^2, & \ln L < x < \ln U, \\ +\infty, & \text{otherwise}, \end{cases} \tag{2.13}$$

where, $\hat{H}_{BS}$ represents the non-Hermitian Black-Scholes-Merton Hamiltonian, $\hat{h}_{BS}$ is the transformed Hermitian Hamiltonian, $\hat{\rho}$ is the transformation operator used to map $\hat{H}_{BS}$ to $\hat{h}_{BS}$, $\hat{\eta}$ is the metric operator ensuring $\hat{\eta}$-pseudo Hermiticity, σ is the volatility, r is the risk-free interest rate, x represents the logarithmic transformation of the asset price and V_{DKO} is the potential operator of the Hamiltonian operator of the double knock out option. In particular, $\hat{H}^{\dagger}_{BS}$ denotes the Hermitian adjoint of $\hat{H}_{BS}$, while $\hat{h}_{\mathrm{DKO}}$ denotes the Hermitian Hamiltonian operator governing the double knock-out option in the similarity-transformed representation (Jana and Roy (2012, 2011)).

2.3 Arrow-Debreu Prices in Hilbert Space

Arrow-Debreu state-space pricing represents uncertainty by indexing all possible outcomes at the terminal date T by states of the world, so that any contingent claim can be decomposed into elementary state-contingent payoffs. An Arrow-Debreu security is an elementary claim that pays one unit of account in exactly one state at date T and zero otherwise (Arrow (1964)). In a finite state economy with states $s \in \{1, \ldots, S\}$, any payoff $X_T(s)$ can be written as a linear combination of the Arrow-Debreu payoffs, and at time t the no-arbitrage price of an asset is therefore a linear functional of the payoff as shown in Equation (2.14). When the relevant terminal state is continuous, such as $y = S_T$ or $y = x_T = \ln S_T$, the collection of Arrow-Debreu prices is represented by a state-price density $\pi_{t,T}(y)$, and the pricing functional takes the continuum form given by Equation (2.15), which is the counterpart of Equation (2.14). The state-price density assigns today's price to one unit of payoff delivered in an infinitesimal neighbourhood around state y at date T, and the non-negativity of Arrow-Debreu prices implies $\pi_{t,T}(y) \geq 0$ under no-arbitrage for standard state spaces and reference measures (Aït-Sahalia and Lo (1998)).

$$V_t = \sum_{s=1}^{S} \pi_{t,T}(s)\, X_T(s), \tag{2.14}$$

$$V_t = \int_{\mathcal{Y}} \pi_{t,T}(y)\, X_T(y)\, dy, \tag{2.15}$$

where t denotes the valuation time, T is the maturity, V_t is the time t no-arbitrage value of the contingent claim, S is the number of states in the finite-state economy, $X_T(s)$ is the payoff delivered at time T in state s, and $\pi_{t,T}(s)$ is the time t price of the Arrow-Debreu security that pays one unit of account at time T if and only if state s occurs. In the continuum-state representation, y denotes a terminal state variable such as S_T or $x_T = \ln S_T$, $\mathcal{Y}$ is the corresponding state space, $X_T(y)$ is the terminal payoff expressed as a function of y, and $\pi_{t,T}(y)$ is the state-price density over $\mathcal{Y}$ (Arrow (1964), Aït-Sahalia and Lo (1998)).

State-price densities in the continuum case and Arrow-Debreu state prices in a finite-state economy offer a direct route to valuing many derivatives because once the relevant state-price system is characterised, the value of any terminal state-contingent payoff is obtained by integrating or summing against those state prices. In practice, the main difficulty is constructing the required state prices, which is particularly challenging for path-dependent claims whose payoff depends on additional path functionals, such as running averages or extrema. In this case, the standard solution is therefore to enlarge the Markov state by introducing auxiliary variables that track the relevant path functionals so that the augmented process becomes Markov on a higher-dimensional state space, which restores a state-price representation but typically increases numerical complexity and motivates specialised partial differential equations and auxiliary-state algorithms for path-dependent derivatives (Ritchken et al. (1993), Rogers and Shi (1995), Lin and Ritchken (2006)). Barrier options provide a canonical illustration because their payoffs hinge on whether a barrier is breached prior to maturity, and therefore the terminal distribution recovered from plain-vanilla option prices does not, in general, identify the required first-hitting-time behaviour (Brown et al. (2001)). A complementary strand of the asset-pricing literature formulates Arrow-Debreu valuation in a Hilbert space of square-integrable payoffs. It represents the price map as a continuous positive linear functional on a payoff subspace, which is useful because it yields a norm-based geometric framework for analysing existence and non-uniqueness of admissible state-price representations in incomplete markets and for deriving sharp bounds and diagnostics implied by limited asset data, while the economically essential restriction of positivity rules out price assignments that would generate arbitrage for nonnegative claims (Harrison and Kreps (1979), Chamberlain and Rothschild (1983), Hansen and Richard (1987), Hansen and Jagannathan (1991, 1997)).

Within a Hilbert-space formulation of derivative pricing, it is useful to recall the standard spectral construction of the time-dependent Schrödinger equation from the stationary

eigenvalue problem of the corresponding time-independent Hamiltonian. For notational simplicity, we first write the discrete-spectrum case. If the stationary eigenstates satisfy Equation (2.16) and form a complete orthonormal system as in Equation (2.17), then an arbitrary state can be expanded according to Equation (2.18). Substituting this expansion into the time-dependent Schrödinger equation (2.3) yields Equation (2.19), and orthonormality reduces the dynamics to the decoupled coefficient equation together with its explicit solution in Equation (2.20). This in turn leads to the standard spectral representation (2.21). On unbounded domains, the corresponding discrete sum is understood in the usual spectral-integral sense, whereas in bounded domains, such as the double knock-out barrier problem, the spectrum is genuinely discrete (Griffiths and Schroeter (2018)).

$$\hat{H}\,|\phi_n\rangle = E_n\,|\phi_n\rangle\,, \tag{2.16}$$

$$\sum_n |\phi_n\rangle\,\langle\phi_n| = \hat{I}, \qquad \int_{\mathbb{R}} dx\ |x\rangle\,\langle x| = \hat{I}, \tag{2.17}$$

$$|\Psi(t)\rangle = \sum_n a_n(t)\ |\phi_n\rangle\,, \tag{2.18}$$

$$\sum_n \frac{da_n(t)}{dt}\,|\phi_n\rangle = -\frac{i}{\hbar}\sum_n a_n(t)\,\hat{H}\,|\phi_n\rangle = -\frac{i}{\hbar}\sum_n a_n(t)\,E_n\ |\phi_n\rangle\,, \tag{2.19}$$

$$\frac{da_n(t)}{dt} = -\frac{i}{\hbar}E_n a_n(t), \qquad a_n(t) = a_n(0)\,e^{-iE_n t/\hbar}, \qquad a_n(0) = \langle\phi_n|\,|\Psi(0)\rangle\,. \tag{2.20}$$

$$|\Psi(t)\rangle = \sum_n e^{-iE_n t/\hbar}\,|\phi_n\rangle\,\langle\phi_n|\,|\Psi(0)\rangle = e^{-i\hat{H}t/\hbar}\,|\Psi(0)\rangle\,. \tag{2.21}$$

where $\hat{H}$ is a Hamiltonian, $|\phi_n\rangle$ and E_n denote its stationary eigenvectors and eigenvalues, $|\Psi(t)\rangle$ is an arbitrary quantum state at time t, $|\Psi(0)\rangle$ is the initial state, $a_n(t)$ are the time-dependent expansion coefficients, $\hbar$ is Planck's constant, $\hat{I}$ is the identity operator, and $|x\rangle$ denotes the position basis state (Griffiths and Schroeter (2018)).

The same spectral argument carries over to the Black-Scholes-Merton setting once the pricing problem is written as a Euclidean-time evolution in the time-to-maturity variable $\tau = T - t$. Since the terminal payoff is prescribed at maturity rather than the initial state at $t = 0$, we denote the known maturity state by $|g_T\rangle = |g(\tau = 0)\rangle = |g(t = T)\rangle$. Because $\hat{H}_{BS}$ is η-pseudo-Hermitian, as discussed in Equations (2.8)-(2.12), its eigenstates satisfy the modified orthonormality and completeness relations in Equation (2.23). Accordingly, the option state can be expanded as in Equation (2.24), and substitution into Equation (2.7) gives Equation (2.25). Projection with the η-inner product then yields the decoupled

coefficient equation together with its explicit solution in Equation (2.26). This leads to the pseudo-Hermitian spectral representation (2.27) and the corresponding spectral decomposition (2.28) (Jana and Roy (2012)).

$$\hat{H}_{BS} \left|\psi_n\right\rangle = \epsilon_n \left|\psi_n\right\rangle, \tag{2.22}$$

$$\left\langle\psi_m\right| \hat{\eta} \left|\psi_n\right\rangle = \delta_{mn}, \qquad \sum_n \left|\psi_n\right\rangle \left\langle\psi_n\right| \hat{\eta} = \hat{I}, \tag{2.23}$$

$$\left|g(\tau)\right\rangle = \sum_n b_n(\tau) \left|\psi_n\right\rangle, \tag{2.24}$$

$$\sum_n \frac{db_n(\tau)}{d\tau} \left|\psi_n\right\rangle = -\sum_n b_n(\tau) \hat{H}_{BS} \left|\psi_n\right\rangle = -\sum_n b_n(\tau) \epsilon_n \left|\psi_n\right\rangle, \tag{2.25}$$

$$\frac{db_n(\tau)}{d\tau} = -\epsilon_n b_n(\tau), \qquad b_n(\tau) = b_n(0)\, e^{-\epsilon_n \tau}, \qquad b_n(0) = \left\langle\psi_n\right| \hat{\eta} \left|g_T\right\rangle. \tag{2.26}$$

$$\left|g(\tau)\right\rangle = \sum_n e^{-\epsilon_n \tau} \left|\psi_n\right\rangle \left\langle\psi_n\right| \hat{\eta} \left|g_T\right\rangle = e^{-\tau \hat{H}_{BS}} \left|g_T\right\rangle, \tag{2.27}$$

$$e^{-\tau \hat{H}_{BS}} = \sum_n e^{-\epsilon_n \tau} \left|\psi_n\right\rangle \left\langle\psi_n\right| \hat{\eta}. \tag{2.28}$$

where $x = \ln S$ is the log value of the spot price, $\tau = T - t$ is time to maturity, $|g_T\rangle$ is the payoff state prescribed at maturity, $|g(\tau)\rangle$ is the option state at time to maturity τ, $\hat{H}_{BS}$ is the non-Hermitian Black-Scholes-Merton Hamiltonian in log-price representation, $|\psi_n\rangle$ and ϵ_n are the corresponding eigenstates and eigenvalues, $\hat{\eta}$ is the metric operator, $b_n(\tau)$ are the corresponding expansion coefficients, and $\hat{I}$ is the identity operator (Jana and Roy (2012)).

To recover Arrow-Debreu state prices in the log-price representation, we project the solution onto the basis $|x\rangle$. This yields the log-price space representation (2.29) and suggests introducing the terminal payoff function in log-price form as $g_T(x_T) = \langle x_T | g_T \rangle$. We then define the Arrow-Debreu state-price density according to Equation (2.30). With this definition, the option value is obtained by integrating the payoff against the state-price kernel as in Equation (2.31). Using the similarity transform (2.9), the Arrow-Debreu state-price density can be expressed by the metric operators (2.32). Substituting the pseudo-Hermitian spectral decomposition (2.28) into Equation (2.30), the equivalent eigenfunction expansion (2.33) is expressed, in which the metric factor enters explicitly through the modified completeness relation. In this convention, the metric operator

transformation is part of the Arrow-Debreu state-price density, so it no longer needs to be taken into account once more in the pricing integral (Jana and Roy (2012, 2011)). In this operator formulation, the Arrow–Debreu state-price density computed in the simulations is the financial counterpart of the Green's function, or propagator kernel, used in physics. A Green's function describes how the effect of a unit impulse introduced at one state is propagated by the governing linear operator to another state over a given time interval (Baaquie (2007)).

$$g(x_t,\tau) = \langle x_t | g(\tau) \rangle = \langle x_t | \, e^{-\tau \hat{H}_{BS}} \, | g_T \rangle = \int_D dx_T \; \langle x_t | \, e^{-\tau \hat{H}_{BS}} \, | x_T \rangle \; \langle x_T | g_T \rangle \,, \tag{2.29}$$

$$\pi_{t,T}(x_t,x_T) = \langle x_t | \, e^{-\tau \hat{H}_{BS}} \, | x_T \rangle \,, \tag{2.30}$$

$$g(x_t,\tau) = \int_D dx_T \, \pi_{t,T}(x_t,x_T) \, g_T(x_T), \tag{2.31}$$

$$\pi_{t,T}(x_t,x_T) = \langle x_t | \, \hat{\rho}^{-1} e^{-\tau \hat{h}_{BS}} \hat{\rho} \, | x_T \rangle = \hat{\rho}(x_t)^{-1} \, \hat{\rho}(x_T) \; \langle x_t | \, e^{-\tau \hat{h}_{BS}} \, | x_T \rangle \,. \tag{2.32}$$

$$\pi_{t,T}(x_t,x_T) = \sum_n e^{-\epsilon_n \tau} \; \langle x_t | \psi_n \rangle \; \langle \psi_n | \, \hat{\eta} \, | x_T \rangle = \sum_n \hat{\eta}(x_T) \, e^{-\epsilon_n \tau} \, \psi_n(x_t) \, \psi_n^*(x_T), \tag{2.33}$$

where $x_t = \ln S_t$ is the log-price at valuation time t, $x_T = \ln S_T$ is the terminal log-price at maturity T, $\tau = T - t$ is time to maturity, $|x\rangle$ is the log-price basis, $g(x_t,\tau)$ is the option value in log-price representation, $g_T(x_T)$ is the terminal payoff function in the same representation, $\pi_{t,T}(x_t,x_T)$ is the corresponding Arrow-Debreu state-price density, $\psi_n(x) = \langle x | \psi_n \rangle$ are the eigenfunctions of $\hat{H}_{BS}$, $\psi_n^*(x) = \langle \psi_n | x \rangle$ denotes the complex conjugate of $\psi_n(x)$ and is the coordinate-space representation of the bra $\langle \psi_n |$, which is the Hermitian adjoint of the ket $|\psi_n\rangle$, ϵ_n are the associated eigenvalues, $\hat{\eta}$ is the metric operator, $\hat{h}_{BS}$ is the Hermitian Hamiltonian obtained from the similarity transformation.

2.4 Numerical Simulations of the State-Price Densities and Derived Quantities

This subsection compares numerical constructions of Arrow-Debreu state prices and derived quantities for plain-vanilla European and double knock-out barrier options. All experiments are implemented in Python and executed on an Intel(R) Core(TM) i7-9700K

3.6 GHz CPU without parallelisation. Across methods, the algorithms start from the same set of financial inputs, namely the constant risk-free rate r, the constant volatility σ, strike K, initial price S_0, maturity T, when relevant, the lower and upper barriers (L, U), and differ primarily in a sensitivity parameter that controls numerical granularity. To avoid ambiguity, we distinguish between price solvers used in the runtime/error comparisons and density solvers used in the state-price comparison. The resolution parameters are the number of simulated paths and monitoring times for Monte Carlo, the number of tree steps for the Cox-Ross-Rubinstein lattice, the spatial and temporal grid sizes for Crank–Nicolson, the log-price grid size n for the dense Hamiltonian spectral solver, and additionally the spectral-tail tolerance $\varepsilon_{\text{tail}}$ for the shift-invert Lánczos variant.

The classical baseline constructions follow the standard computational finance simulations. For Monte Carlo, the double knock-out state-price density is obtained by simulating discretely monitored geometric-Brownian paths, retaining only those paths that remain within the barrier region throughout the monitoring dates, constructing a histogram-based approximation of their terminal log-prices on the target interior grid, and discounting the resulting density. The Cox-Ross-Rubinstein lattice binomial tree baseline propagates risk-neutral node probabilities forward, removes nodes that leave (L, U) in the barrier case, and maps the surviving terminal state probabilities from the native tree log-grid to the common comparison grid, followed by a renormalisation to preserve total probability (Wilmott (2013), Hull and Basu (2016)). For Crank-Nicolson, the option valuation benchmark uses the standard backward pricing partial differential equation, whereas the state price density comparison uses a forward Kolmogorov equation. This means that time is advanced using a second-order implicit scheme on an interior log-price grid, yielding a tridiagonal linear system at each time step. In our implementation, these tridiagonal systems are solved using the Thomas algorithm, which exploits the banded structure to obtain a good numerical stability for the diagonally dominant matrices arising in this setting (Wilmott (2013), Wade et al. (2007), Zvan et al. (2000)).

Finally, we use the Hermitian Black-Scholes-Merton Hamiltonian (2.11) to construct the Arrow-Debreu state prices via the spectral expansion (2.32). First, we discretise the log-price coordinate $x = \ln S$ on a uniform grid and represent $\hat{h}_{BS}$ by a real symmetric tridiagonal sparse matrix obtained from a second-order finite-difference approximation of the kinetic operator plus the constant potential term. Second, we diagonalise the discrete operator by computing its eigenstates and eigenvalues, and construct the corresponding Arrow-Debreu prices for a given derivative. Algorithm 1 summarises the corresponding Hamiltonian simulation of the double knock-out option's Arrow Debreu prices, in which the lower and upper barriers are the log-domain $(x_L, x_U) = (\ln L, \ln U)$ and impose absorbing Dirichlet boundaries through the choice of interior grid points. The same

simulation applies to plain-vanilla options, except that (x_L, x_U) is chosen as a sufficiently wide truncation interval, so that boundary truncation error is negligible. The values of the different option contracts are then obtained as the inner product of the Arrow-Debreu state prices with the corresponding terminal payoff function (Léber (2021)). An important practical feature of this construction is that the Arrow-Debreu state prices can be stored once they have been computed on the chosen grid. After that, the valuation of alternative terminal payoffs, digital components, and several state-price-density-based risk measures can be estimated at low cost by post-processing the same saved state-price density, rather than rerunning the full numerical simulation. The same framework is also effective when results are needed for several maturities. The full density $\pi_{0,T}$ is maturity-specific, so it is not the same saved density reused across maturities. However, the discretised Hamiltonian and its eigenpairs are maturity-independent for fixed model parameters and grid. Once those eigenpairs have been computed and stored, moving from one maturity to another requires only updating the weights $e^{-\epsilon_n T}$, which is much cheaper than a new simulation (Davydov and Linetsky (2003), Linetsky (2007)).

Algorithm 1 Hamiltonian spectral simulation of the Arrow-Debreu state prices in the case of the double knock-out barrier option

1: $x_0 \leftarrow \ln S_0, \quad x_L \leftarrow \ln L, \quad x_U \leftarrow \ln U, \quad \Delta x \leftarrow \dfrac{x_U - x_L}{n+1}$

2: $x_i \leftarrow x_L + i\Delta x$ for $i = 1, \ldots, n$

3: $\alpha \leftarrow \dfrac{\frac{1}{2}\sigma^2 - r}{\sigma^2} \quad b \leftarrow \dfrac{\left(\frac{1}{2}\sigma^2 + r\right)^2}{2\sigma^2} \quad d \leftarrow \dfrac{\sigma^2}{(\Delta x)^2} + b, \quad e \leftarrow -\dfrac{\sigma^2}{2(\Delta x)^2}$

4: Compute all eigenpairs (ϵ_j, v_j) of $h_{BSM} = \text{tridiag}(e, d, e)$ for $j = 1, \ldots, n$

5: $V = [v_1, \ldots, v_n] \quad \epsilon = [\epsilon_1, \ldots, \epsilon_n]$

6: $j_0 \in \arg\min_{1 \le i \le n} |x_i - x_0|, \quad (e_0)_i \leftarrow \dfrac{1}{\Delta x}\,\delta_{ij_0}$ for $i = 1, \ldots, n$

7: $k \leftarrow V\Big(\exp(-T\epsilon) \circ (V^\top e_0)\Big)$

8: $\pi_{0,T}(x_i) \leftarrow \exp\big(\alpha(x_0 - x_i)\big)\, k_i$ for $i = 1, \ldots, n$

9: $\pi_{0,T}(x_i) \leftarrow \max(\pi_{0,T}(x_i), 0)$

10: **return** $\pi_{0,T}(x_i)$

The direct diagonalisation in Algorithm 1 is robust and convenient on moderate grids because H is tridiagonal and symmetric, but it becomes costly as the n sensitivity parameter grows because the algorithm should compute all eigenpairs. As emphasised by Linetsky (2007), Davydov and Linetsky (2003), accurate valuation typically does not require the full spectrum, since the contribution of higher modes is suppressed. In Algorithm 1, the state price density is weighted by Euclidean-time factors $\exp\big(-\epsilon_j T\big)$, so high-energy contributions are exponentially suppressed, and only a relatively small subset of the low-lying eigenpairs is needed in practice. Economically, the omitted high-ϵ_j modes represent fine-scale variations across nearby log-price states and are excluded only when

their contributions fall below the chosen spectral-tail tolerance, whereas the retained low-lying modes carry the smoother valuation structure that persists over the horizon T. This motivates the use of Krylov-subspace methods, which approximate the action of H on a much smaller invariant subspace. Starting from a vector v, one forms the Krylov subspace $\mathcal{K}_m(H, v) = \text{span}\{v, Hv, \ldots, H^{m-1}v\}$ and projects H onto $\mathcal{K}_m$, obtaining a reduced $m \times m$ matrix whose eigenvalues (Ritz values) approximate the spectral region of interest. For general non-Hermitian matrices, this construction can be implemented by the Arnoldi process, which generates an orthonormal Krylov basis together with an upper-Hessenberg projected matrix (Sorensen (1997), Saad (2011)). In the present framework, however, the pseudo-Hermitian similarity transform yields a Hermitian discretised Hamiltonian, so the Arnoldi reduction simplifies to the Lánczos case, which means the projected matrix becomes real symmetric tridiagonal, and the basis construction admits a short three-term recurrence (Sorensen (1997), Saad (2011)).

Lánczos (1950) employs a Liouville-Neumann expansion to develop a systematic procedure for evaluating the latent roots and principal axes of a matrix. The author subsequently refines this approach using the method of minimised iterations, which is intended to control the accumulation of rounding errors and prevent strongly dominant eigencomponents from obscuring smaller ones in highly dispersed spectra. In the Hermitian case, the construction reduces to a single orthogonal sequence. Lánczos (1950) shows that the Green operator is the reciprocal of the differential operator, but also highlights that closed-form Green functions are frequently unavailable. Consequently, he advocates carrying out the iterations directly with the given differential operator. In this sense, the present Hermitian operator-based pricing framework can follow the same procedure, and it is therefore natural that Lánczos proposes the method for linear differential and integral operators (Lánczos (1950)). Accordingly, we replace full diagonalisation by the so-called shift-invert Lánczos procedure presented in Algorithm 2, which targets only the spectral window required for the corresponding Euclidean-time propagation. In Algorithm 2, the Krylov-based eigensolver follows a shift-invert Lánczos strategy with zero shift. Specifically, we apply the Lánczos iteration to the inverse operator $\hat{\mathcal{M}} = \hat{h}_{\text{BSM}}^{-1}$, implemented through a single sparse LU factorisation and repeated triangular solves, rather than forming $\hat{\mathcal{M}}$ explicitly (Saad (2011), Sorensen (1997), Lehoucq et al. (1998)). This choice is well suited to our setting because we require only the low-lying eigenpairs of the Hermitian Hamiltonian and the inversion maps small eigenvalues of $\hat{h}_{\text{BSM}}$ to large-magnitude eigenvalues of $\hat{h}_{\text{BSM}}^{-1}$, thereby turning the computation of low-lying modes into an extremal eigenvalue problem where Krylov methods converge rapidly (Saad (2011), Sorensen (1997)). For the finite Dirichlet discretisation used here, $\hat{h}_{\text{BSM}}$ is positive definite, so $0 \notin \sigma(\hat{h}_{\text{BSM}})$ and the zero-shift inverse is well defined. Spectrally, $\epsilon_j \to +\infty$ implies $\epsilon_j^{-1} \to 0$, thus the informal notation $1/\infty = 0$ is only shorthand for this limit, not

an arithmetic identity or an indication that $+\infty$ is a spectral point. A nonzero shift σ_{SI} would be beneficial when targeting eigenvalues in the interior of the spectrum or within clustered regions, because the shift-invert procedure makes eigenvalues near σ_{SI} appear as extremal Ritz values of the transformed problem and can substantially accelerate convergence (Saad (2011), Sorensen (1997)). In practice, such a shift is most effective when one has prior information about the spectral region of interest.

Algorithm 2 Hamiltonian shift-invert Lánczos spectral simulation of the Arrow-Debreu state prices in the case of the double knock-out barrier option

$x_0 \leftarrow \ln S_0, \quad x_L \leftarrow \ln L, \quad x_U \leftarrow \ln U, \quad \Delta x \leftarrow \dfrac{x_U - x_L}{n+1}$

$x_i \leftarrow x_L + i\Delta x$ for $i = 1, \ldots, n$

$\alpha \leftarrow \dfrac{\frac{1}{2}\sigma^2 - r}{\sigma^2} \quad b \leftarrow \dfrac{\left(\frac{1}{2}\sigma^2 + r\right)^2}{2\sigma^2} \quad d \leftarrow \dfrac{\sigma^2}{(\Delta x)^2} + b, \quad e \leftarrow -\dfrac{\sigma^2}{2(\Delta x)^2}$

$\hat{h}_{BSM} \leftarrow \mathrm{tridiag}(e, d, e)$ in sparse format

$k_{\mathrm{eig}} \leftarrow \min\left(n-1, \left\lceil \frac{x_U - x_L}{\pi}\sqrt{\frac{2\log(1/\varepsilon_{\mathrm{tail}})}{\sigma^2 T}} \right\rceil\right)$

Sparse LU factorisation $h_{BSM} = LU$

$\hat{\mathcal{M}}(v) \leftarrow \hat{h}_{BSM}^{-1} v$ via the LU solves

Choose Krylov dimension m and an initial unit vector q_1; set $q_0 \leftarrow 0$, $\beta_0 \leftarrow 0$

repeat

  for $j = 1$ to m **do**

    $w \leftarrow \hat{\mathcal{M}}(q_j) - \beta_{j-1} q_{j-1}$

    $\alpha_j \leftarrow q_j^{\top} w$

    $w \leftarrow w - \alpha_j q_j$

    $\beta_j \leftarrow \|w\|_2$

    $q_{j+1} \leftarrow w/\beta_j$

  end for

  $Q_m \leftarrow [q_1, \ldots, q_m]$

  $T_m \leftarrow \mathrm{tridiag}(\beta_1, \ldots, \beta_{m-1}; \alpha_1, \ldots, \alpha_m)$

  Compute k_{eig} Ritz pairs (μ_i, y_i) of T_m with largest $|\mu_i|$

  $z_i \leftarrow Q_m y_i$ for $i = 1, \ldots, k_{\mathrm{eig}}$

  $\rho_i \leftarrow |\beta_m e_m^{\top} y_i|$ for $i = 1, \ldots, k_{\mathrm{eig}}$

  Apply implicit restart and continue if fewer than k_{eig} values satisfy $\rho_i \le \mathtt{tol}$

until $\rho_i \le \mathtt{tol}$ holds for $i = 1, \ldots, k_{\mathrm{eig}}$

$\epsilon_j \leftarrow 1/\mu_j$ for $j = 1, \ldots, k_{\mathrm{eig}}$ and reorder so that $\epsilon_1 \le \cdots \le \epsilon_{k_{\mathrm{eig}}}$ with matching vectors z_j

$V \leftarrow [z_1, \ldots, z_{k_{\mathrm{eig}}}], \quad \epsilon \leftarrow (\epsilon_1, \ldots, \epsilon_{k_{\mathrm{eig}}})$

$j_0 \in \arg\min_{1 \le i \le n} |x_i - x_0|, \quad (e_0)_i \leftarrow \dfrac{1}{\Delta x}\delta_{ij_0}$ for $i = 1, \ldots, n$

$k \leftarrow V\Big(\exp(-T\epsilon) \circ (V^{\top} e_0)\Big)$

$\pi_{0,T}(x_i) \leftarrow \exp\big(\alpha(x_0 - x_i)\big)\, k_i$ for $i = 1, \ldots, n$

$\pi_{0,T}(x_i) \leftarrow \max(\pi_{0,T}(x_i), 0)$

return $\pi_{0,T}(x_i)$

One convenient approach for choosing the truncation level k in Algorithm 2 is to use a spectral-tail truncation criterion based on the expected high-frequency eigenvalue growth. The relevant scale is primarily determined by the kinetic prefactor $\sigma^2/2$ and the effective log-domain length L_x. Under Dirichlet boundary conditions on an interval of length L_x, the Laplacian eigenvalues satisfy the asymptotic scaling $\lambda_j \sim (\pi j/L_x)^2$ as $j \to \infty$, and therefore the contribution of neglected modes decays roughly like $\exp\big(-\frac{\sigma^2}{2}(\pi j/L_x)^2 T\big)$. A proposed heuristical truncation recommendation is to choose k according to Equation (2.34) so that the first neglected term falls below a user tolerance $\varepsilon_{\text{tail}}$. The proposed recommendation makes explicit that longer maturities T and higher volatilities σ reduce the required number of eigenpairs due to stronger suppression, whereas a more distant lower and upper price limits, which means a larger effective log-domain length interval L_x, increase it because the eigenvalues become more closely spaced. For both the double knock-out and plain-vanilla option cases, the constant potential term shifts the spectrum but does not materially change the high-frequency spacing driven by the kinetic operator, whereas sharp payoffs can slow the decay of eigen-coefficients and therefore increase the number of eigenpairs needed for a given accuracy. The user-specified tolerance $\varepsilon_{\text{tail}}$ determines how small the leading neglected spectral contribution is required to be in the truncated expansion. For example, setting $\varepsilon_{\text{tail}} = 0.000001$ means that k is chosen so that the first omitted mode is expected to be suppressed to roughly $\times 10^{-6}$ or less, although this should be interpreted as a practical tail criterion rather than as a strict bound on the total truncation error. In practice, Equation (2.34) provides a useful starting point, but for other payoff classes it is advisable to also inspect nearby k values to verify that the quantity of interest has stabilised.

$$k \;\geq\; \left\lceil \frac{L_x}{\pi}\sqrt{\frac{2\log(1/\varepsilon_{\text{tail}})}{\sigma^2 T}} \right\rceil, \tag{2.34}$$

where k is the number of retained eigenpairs in the truncated spectral representation, L_x denotes the effective log-domain length, which is the difference between the log upper limit price $\ln U$ and log lower limit price $\ln L$, $\varepsilon_{\text{tail}}$ is the user-prescribed tolerance controlling of the the spectral tail, σ is the constant volatility parameter, and T is the maturity.

2.4.1 Comparison of the Estimated Arrow-Debreu State Prices

To compare the Arrow-Debreu state prices estimated by the different numerical methods, all metrics are estimated on a common uniform grid with $n = 400$ interior points. In the case of the Crank Nicolson and the different Hamiltonian simulations, the log-domain of the plain-vanilla option is selected by an arbitrary truncation rule as a symmetric interval of approximately $\pm 6\sigma\sqrt{T}$ around the risk-neutral mean of the terminal log price $\ln S_T$. For Monte Carlo state-price estimation, we use $N_{\text{paths}} = 30{,}000$ paths, and the binomial tree

method adopts a daily-style monitoring resolution, which means $n_{\text{steps}} = 252$ before mapping onto the common grid. To quantify differences in state-price densities produced by different numerical procedures beyond visual examination, the comparison uses complementary metrics that capture both absolute discrepancies and shape differences. First, the total integrated value consistency is evaluated through the discrete total integrated values of the two sampled state prices, and the associated relative total integrated value error is calculated by Equations (2.35). Second, absolute differences on the state-price arrays are summarised by discrete L^1, L^2, and L^∞ discrepancies, defined in Formulas (2.36). Third, to compare shapes independently of total integrated value, the arrays are normalised to probability densities (2.37) and several distributional distances are used, namely total variation distance $\mathrm{TV}(p, q)$, the one-dimensional 1-Wasserstein distance $W_1(p, q)$ computed via discrete cumulative distribution functions, and information-theoretic divergences based on Kullback-Leibler $D_{\mathrm{KL}}(p\|q)$ and Jensen-Shannon $D_{\mathrm{JS}}(p, q)$ functionals, as specified in Formulas (2.38)-(2.41) (Gibbs and Su (2002)).

$$m = \sum_{i=1}^{n} \pi_i \, \Delta x, \qquad m_{\text{ref}} = \sum_{i=1}^{n} \pi_i^{(\text{ref})} \, \Delta x, \quad \text{relIV} = \frac{|m - m_{\text{ref}}|}{m_{\text{ref}}} \tag{2.35}$$

$$L^1 = \sum_{i=1}^{n} \left|\pi_i - \pi_i^{(\text{ref})}\right| \Delta x, \quad L^2 = \left(\sum_{i=1}^{n} \left(\pi_i - \pi_i^{(\text{ref})}\right)^2 \Delta x\right)^{1/2}, L^\infty = \max_{1 \le i \le n} \left|\pi_i - \pi_i^{(\text{ref})}\right| \tag{2.36}$$

$$p_i = \frac{\pi_i}{m}, \qquad q_i = \frac{\pi_i^{(\text{ref})}}{m_{\text{ref}}}. \tag{2.37}$$

$$\mathrm{TV}(p, q) = \frac{1}{2} \sum_{i=1}^{n} |p_i - q_i| \, \Delta x. \tag{2.38}$$

$$W_1(p, q) = \int_{\mathbb{R}} \left|F_p(x) - F_q(x)\right| dx, \tag{2.39}$$

$$D_{\mathrm{KL}}(p\|q) = \sum_{i=1}^{n} p_i \log\left(\frac{p_i}{q_i}\right) \Delta x. \tag{2.40}$$

$$D_{\mathrm{JS}}(p, q) = \frac{1}{2} D_{\mathrm{KL}}(p\|r) + \frac{1}{2} D_{\mathrm{KL}}(q\|r), \qquad r = \frac{1}{2}(p + q). \tag{2.41}$$

where $\pi_{i\,i=1}^{n}$ and $\pi_i^{(\text{ref})\,n}_{\;\;i=1}$ are two state-price density arrays sampled on the common grid with uniform spacing, m and m_{ref} are the discrete total integrated value of the corresponding state price density arrays, relIV is the corresponding relative total integrated

value discrepancy, the quantities L^1, L^2, and L^∞ are discrete discrepancies computed directly from the state price arrays, the normalised arrays p_i and q_i define probability densities on the grid, $\mathrm{TV}(p, q)$ is the total variation distance, $W_1(p, q)$ is the one-dimensional 1-Wasserstein distance computed via discrete cumulative distribution functions F_p and F_q, D_{KL} and D_{JS} are the Kullback-Leibler and Jensen-Shannon divergences (Gibbs and Su (2002)).

Table 2.1: Model and numerical sensitivity parameters

Parameter	**Value**
Spot price	$S_0 = 100$
Strike price	$K = 125$
Time to maturity	$T = 1.0$
Risk-free interest rate	$r = 3\%$
Volatility	$\sigma = 20\%$
Upper limit price	$U = 140$
Lower limit price	$L = 50$
Monte Carlo sensitivity parameter	$N_{\mathrm{paths}} = 30{,}000$
Binomial tree sensitivity parameter	$n_{\mathrm{steps}} = 252$
Crank–Nicolson sensitivity parameter	$M = 400$
Hamiltonian spectral sensitivity parameter	$n = 400$
Hamiltonian Lánczos spectral sensitivity parameters	$n = 400,\ \varepsilon_{\mathrm{tail}} = 10^{-10}$

Table 2.1 reports the pricing parameters and the method-specific numerical sensitivity settings. We first compare the Arrow-Debreu state prices $\pi_{0,T}(x)$ for the plain-vanilla option across all standard and suggested methods. Table 2.2 summarises the resulting discrepancy measures on the common grid, using the analytical lognormal state price density as the reference distribution. The state price densities obtained from the Hamiltonian spectral simulations are the closest to the analytically derived benchmark density. In particular, the discrepancies produced by the Lánczos-based Hamiltonian approximation look identical to those of the full Hamiltonian spectral method across all reported metrics. This indicates that truncating the spectral expansion and neglecting the less significant eigenvalues introduces no material loss of accuracy in the state price density estimation, while still allowing the approximation to be constructed from a reduced set of eigenpairs. Among the remaining methods, the binomial tree provides the next most accurate approximation after the Hamiltonian simulation-based approaches on the common comparison grid. The Crank-Nicolson density still preserves the total integrated value very accurately, but its pointwise and distributional discrepancies are noticeably larger than those of the binomial tree in the plain-vanilla setting. This indicates that accurate integral conservation and strong price performance do not automatically imply the smallest local density error on the truncated comparison grid. Finally, the Monte Carlo estimator exhibits the largest deviations from the analytical state-price density, consistent with the presence of

finite-sample noise. Figure 2.1 visually shows that the deterministic algorithms, binomial tree, Crank-Nicolson, and the Hamiltonian spectral variants produce indistinguishable state-price densities on the interior grid, while the Monte Carlo estimator exhibits the expected sampling noise due to histogram binning and finite path count.

Table 2.2: Discrepancy measures of the Arrow-Debreu state prices for the plain vanilla option, the reference is the analytical Black-Scholes-Merton state price density

Metric	Monte Carlo	Binomial Tree	Crank–Nicolson	Hamiltonian	Hamiltonian Lánczos
mass	9.704×10^{-1}	9.704×10^{-1}	9.704×10^{-1}	9.704×10^{-1}	9.704×10^{-1}
relIV	2.160×10^{-9}	7.107×10^{-10}	2.015×10^{-9}	1.568×10^{-7}	1.568×10^{-7}
L^1	6.031×10^{-2}	1.949×10^{-3}	1.159×10^{-2}	6.973×10^{-5}	6.973×10^{-5}
L^2	7.745×10^{-2}	2.278×10^{-3}	1.220×10^{-2}	6.624×10^{-5}	6.624×10^{-5}
L^∞	2.716×10^{-1}	5.717×10^{-3}	1.767×10^{-2}	1.044×10^{-4}	1.044×10^{-4}
TV	3.107×10^{-2}	1.004×10^{-3}	5.972×10^{-3}	3.593×10^{-5}	3.593×10^{-5}
W_1	1.352×10^{-3}	2.695×10^{-4}	2.993×10^{-3}	8.332×10^{-6}	8.332×10^{-6}
KL	1.685×10^{-2}	3.422×10^{-6}	1.119×10^{-4}	2.126×10^{-8}	2.126×10^{-8}
JS	1.311×10^{-3}	8.524×10^{-7}	2.799×10^{-5}	5.305×10^{-9}	5.305×10^{-9}

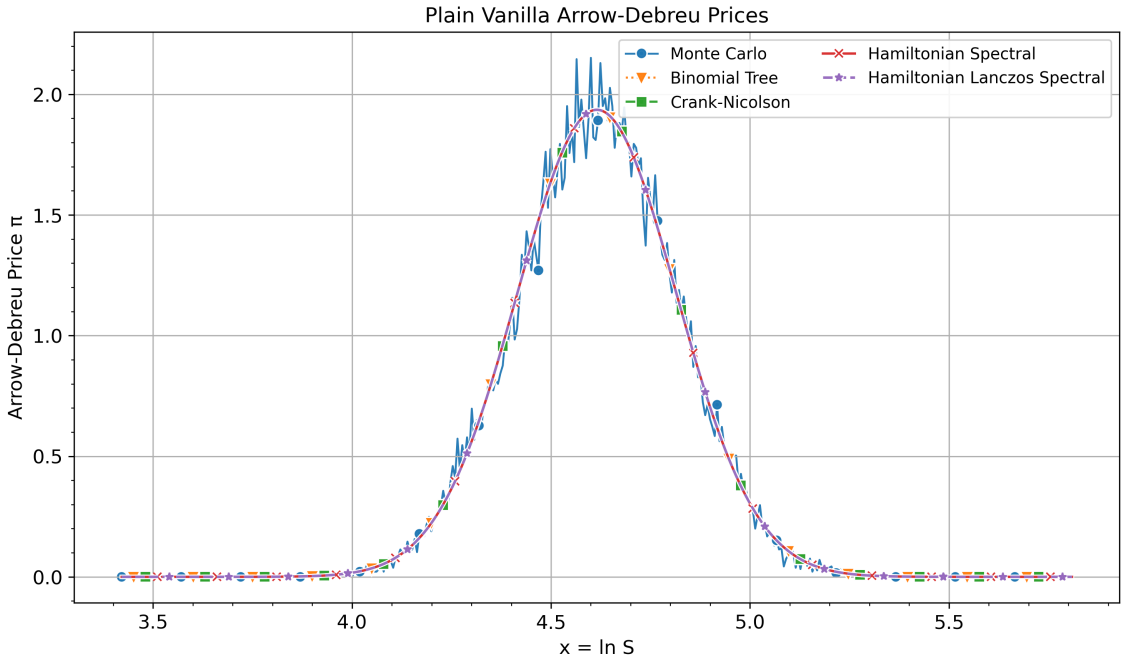

Figure 2.1: Estimation of the plain-vanilla Arrow-Debreu state prices by different algorithms

We also compare the Arrow-Debreu state prices for the double knock-out contract. In this case, no analytical benchmark is available, and we therefore adopt the simple Hamiltonian spectral simulation as the reference solution. Table 2.3 summarises the

corresponding quantitative discrepancy measures. The first observation from Table 2.3 is fully consistent with the conclusion already established in the plain-vanilla setting, namely that the state price densities generated by the Hamiltonian Lánczos simulation are extremely close to those obtained from the simple Hamiltonian spectral method. In the present case, the difference between the two Hamiltonian-based approximations is essentially at the level of machine precision. Moreover, this gap can be reduced even further by decreasing the user-controlled tolerance parameter $\varepsilon_{\text{tail}}$, thereby allowing the Lánczos approximation to recover the simple Hamiltonian solution with still greater accuracy. This again supports the view that, even in the path-dependent setting, the contribution of the less significant eigenpairs is negligible in practice. It is also evident that the Crank-Nicolson approximation remains very close to the Hamiltonian-based estimates. This indicates that numerical procedures based on the direct solution of the governing partial differential equation for pricing yield highly consistent state price densities, even in the double-knockout case. The binomial tree approximation performs reasonably well and is substantially closer to the Crank-Nicolson and Hamiltonian-based solutions than the Monte Carlo estimator. By contrast, the Monte Carlo method exhibits the largest discrepancy among all methods considered. In the path-dependent case, a plausible explanation is that the double knock-out payoff introduces strong path dependence through the absorbing barriers, making the state price density particularly sensitive to rare barrier-hitting events and to the accurate resolution of survival probabilities near the boundaries. Standard Monte Carlo estimation is less well suited to high precision exotic option pricing, since a finite sample may inadequately capture the frequency and timing of barrier crossings, especially when the density is reconstructed on a fixed grid. As a result, the estimator exhibits noticeably greater sampling noise than the partial differential equation-based methods, which handle the barrier structure directly through the underlying operator formulation. Figure 2.2 visually exhibits the different double knock-out state-price densities, in which the vertical dotted lines indicate the barrier locations in log-space. The Crank-Nicolson forward-density solver and the Hamiltonian spectral constructions again coincide almost perfectly across the interior region. The Monte Carlo and binomial kernel estimates are broadly consistent in shape but are visibly rougher due to sampling noise and the discrete nature of the histogram mapping.

2.4.2 Accuracy Against Analytical Solution of Option Price Valuation and Runtime-Error Diagnosis

The plain-vanilla European call option admits a closed-form analytical solution under the Black-Scholes-Merton framework, which is therefore used as the benchmark for assessing the numerical accuracy of option pricing. Based on the parameter configuration

Table 2.3: Discrepancy measures of the Arrow–Debreu state prices in the case of the double knock-out option, the reference is the Hamiltonian spectral estimation

Metric	Monte Carlo	Binomial Tree	Crank–Nicolson	Hamiltonian Lánczos
mass	8.802×10^{-1}	8.764×10^{-1}	8.725×10^{-1}	8.725×10^{-1}
relIV	8.913×10^{-3}	4.524×10^{-3}	1.138×10^{-7}	2.462×10^{-12}
L^1	7.316×10^{-2}	4.846×10^{-3}	1.668×10^{-7}	2.148×10^{-12}
L^2	1.049×10^{-1}	9.096×10^{-3}	2.231×10^{-7}	3.104×10^{-12}
L^∞	3.858×10^{-1}	6.606×10^{-2}	4.541×10^{-7}	6.518×10^{-12}
TV	4.116×10^{-2}	3.031×10^{-3}	8.868×10^{-8}	3.594×10^{-13}
W_1	3.077×10^{-3}	7.357×10^{-4}	3.229×10^{-8}	8.234×10^{-14}
KL	1.934×10^{-2}	9.917×10^{-3}	2.083×10^{-14}	9.445×10^{-18}
JS	1.919×10^{-3}	1.638×10^{-4}	5.196×10^{-15}	1.891×10^{-18}

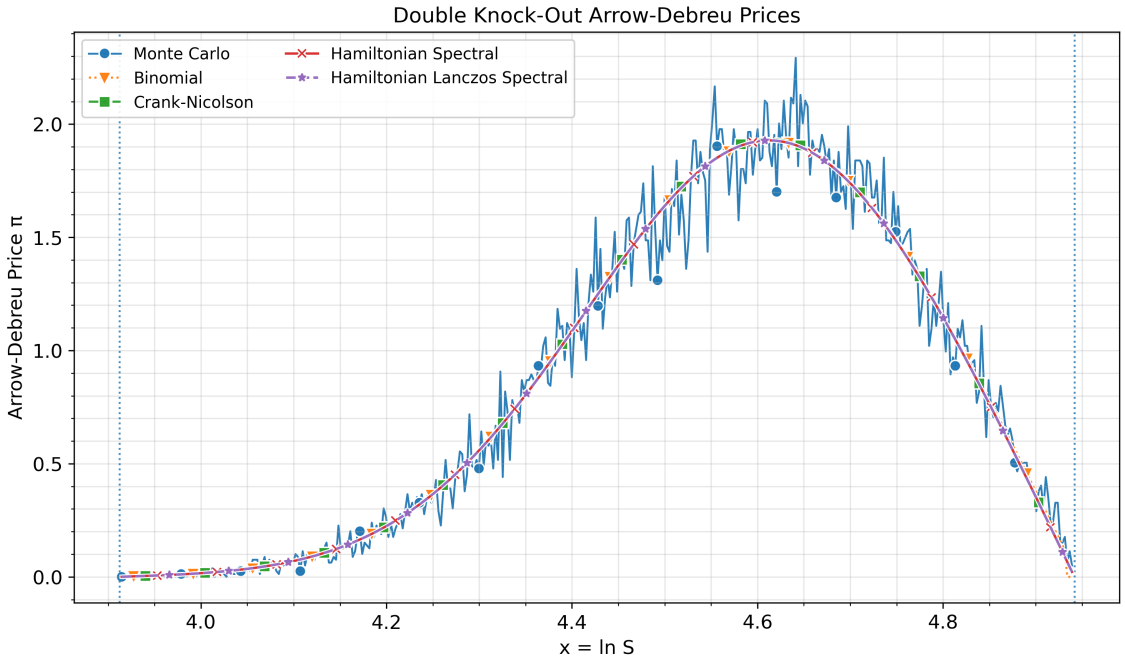

Figure 2.2: Double knock-out Arrow-Debreu state prices

reported in Table 2.1, the plain-vanilla call prices obtained from the different numerical methods are compared in Table 2.4, which summarises the option values produced by each simulation together with the corresponding absolute and relative errors with respect to the analytical Black-Scholes-Merton price. Among all methods considered, the Hamiltonian spectral approach achieves the smallest absolute error, followed very closely by the Hamiltonian Lánczos spectral method. The additional discrepancy introduced by the Lánczos approximation is again negligible and remains essentially at the level of machine precision relative to the simple Hamiltonian spectral solution. As in the state price density comparison, the next most accurate result is provided by the Crank-Nicolson scheme, followed by the binomial tree method, while the Monte Carlo simulation exhibits the largest

deviation from the analytical benchmark. Overall, the results confirm the high accuracy of the Hamiltonian-based pricing framework, and show that the Lánczos reduction preserves this accuracy to an almost indistinguishable extent.

Table 2.4: Plain-vanilla European call price comparison against the analytical Black-Scholes-Merton value

Method	Price	$\lvert V - V_{\mathrm{BSM}} \rvert$	Relative error
Monte Carlo	1.9647231330	1.262382×10^{-2}	6.4668×10^{-3}
Binomial Tree	1.9491787869	2.920522×10^{-3}	1.4961×10^{-3}
Crank–Nicolson	1.9526430849	5.437764×10^{-4}	2.7856×10^{-4}
Hamiltonian Spectral	1.9524891266	3.898181×10^{-4}	1.9969×10^{-4}
Hamiltonian Lánczos Spectral	1.9524891267	3.898182×10^{-4}	1.9969×10^{-4}

A more informative perspective is provided by the runtime-accuracy diagram shown in Figure 2.3, which plots the absolute pricing error against runtime on logarithmic axes. The figure clearly illustrates the expected trade-off between computational cost and numerical accuracy across the different methods. Monte Carlo simulation exhibits only moderate runtimes in this experiment, but its error remains comparatively large and decreases rather slowly, reflecting the persistent effect of sampling noise. In contrast, the binomial tree achieves millisecond-scale runtimes and therefore performs attractively in the very low-cost regime, although its convergence in accuracy is slower and less uniform than the partial differential equation solutions and spectral approaches. The Crank-Nicolson finite-difference method attains very small errors, but this improvement comes at the cost of a substantial increase in runtime as the discretisation is refined. The simple Hamiltonian spectral methods achieve a particularly favourable balance between runtime and accuracy. The conventional dense spectral implementation already yields low errors at comparatively short runtimes, whereas the Hamiltonian Lánczos spectral method becomes increasingly advantageous as the resolution increases. At coarse discretisations, the Lánczos variant is initially less efficient because of its additional setup overhead. However, beyond the first low-resolution cases, it becomes increasingly favourable and eventually attains essentially the same accuracy as the full spectral method, while requiring less computational time. In other words, the two Hamiltonian implementations are virtually indistinguishable in accuracy, differing only at the level of machine precision, whereas their runtime profiles diverge significantly as the problem size increases. This behaviour reflects that the dense spectral solver performs a full diagonalisation and therefore becomes increasingly expensive as the matrix dimension grows, whereas the shift-invert Lánczos approach constructs only a truncated subset of the low-lying eigenpairs. As a result, the Lánczos method offers a more attractive asymptotic trade-off between cost and accuracy once the grid is sufficiently refined. At the same time, this introduces an explicit trade-off between runtime and accuracy through the user-controlled tolerance parameter $\varepsilon_{\mathrm{tail}}$, which determines how

many eigenpairs are retained in the truncated expansion. Another noteworthy feature of the Figure 2.3 is that the Hamiltonian methods dominate much of the practically relevant accuracy range. They substantially outperform Monte Carlo in accuracy for comparable runtimes, and they achieve errors competitive with or smaller than those of Crank-Nicolson at a markedly lower computational cost. The run times correspond to a pure Python implementation without parallelisation and should therefore be interpreted primarily as a relative comparison of the numerical methods under a common computational setup, rather than as absolute performance claims.

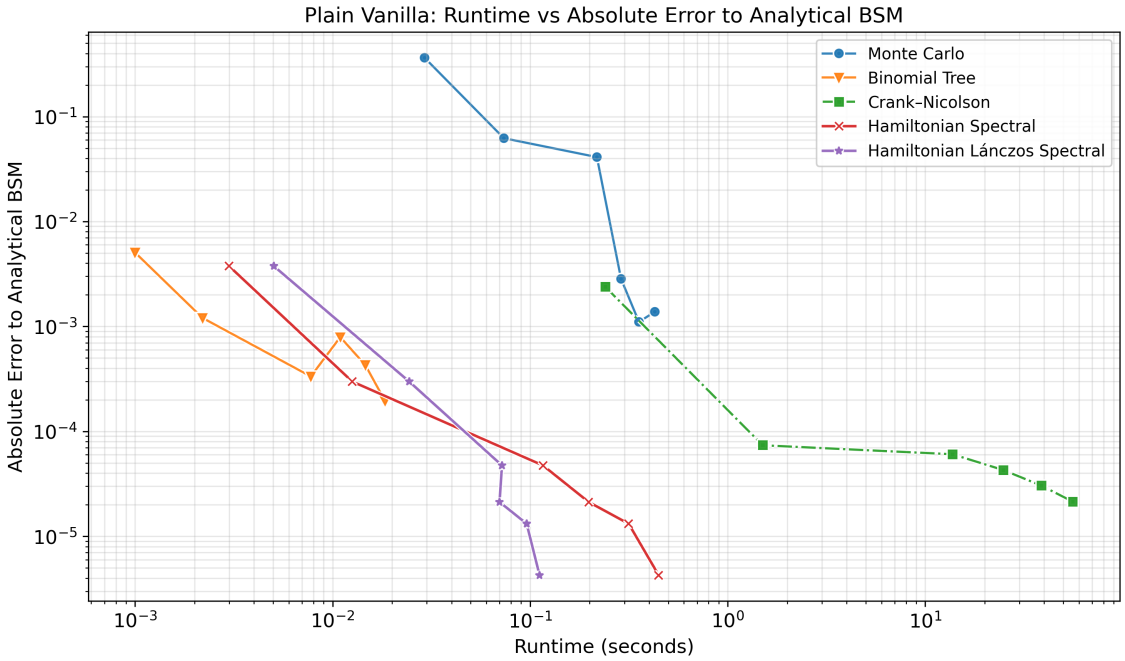

Figure 2.3: Runtime versus absolute error against the analytical Black-Scholes-Merton price

2.4.3 The Double Knock-Out Runtime-Convergence Comparison

A comparable runtime-accuracy analysis cannot be performed for the double knock-out option, since path-dependent options typically do not admit a closed-form analytical pricing formula that can serve as an exact benchmark. Nevertheless, the numerical behaviour of the different methods can still be assessed by examining their convergence patterns and computational runtimes as the sensitivity parameter increases. Figure 2.4 presents the convergence of double knock-out price estimates for all examined algorithms. The Monte Carlo curve exhibits substantial fluctuations at relatively low sensitivity levels. Although the estimates become more stable as the sensitivity parameter increases, they remain noticeably noisier than those from deterministic methods across the range considered. To better resolve the behaviour of the deterministic competitors, Figure 2.5

excludes the Monte Carlo results. On this scale, the agreement between the conventional Hamiltonian spectral method and the Hamiltonian Lánczos spectral method is essentially exact, confirming that the Lánczos eigenpair selection strategy is fully adequate for the grid sizes examined. A similar convergence pattern for the double knock-out barrier case has already been examined by Léber (2021) in the context of a simple Hamiltonian simulation framework, although without a comparative analysis across alternative numerical methods. The Crank-Nicolson and Hamiltonian spectral curves both stabilise rapidly and remain very close to each other over the entire sensitivity range. At the same time, the fact that these curves intersect at several points indicates that there is no uniformly clear ordering between them in terms of which method approaches the solution value more quickly. Rather, both methods display fast and highly comparable convergence, suggesting that direct partial differential equation solutions yield stable, high-accuracy price estimates with substantially fewer refinement steps than either Monte Carlo or the binomial tree method. By contrast, the binomial tree estimates remain systematically farther from the partial differential equation solutions over the same sensitivity range and continue to exhibit an oscillatory pattern as the number of steps varies. The barrier and the discontinuous knock-out condition interact with the underlying time-space discretisation in such a way that even small changes in the step count may alter the relative alignment of the lattice with the barrier, thereby producing alternating over- and under-estimation. This type of erratic or oscillatory convergence of barrier options is well studied in the literature (Boyle and Lau (1994), Lin and Palmer (2013)).

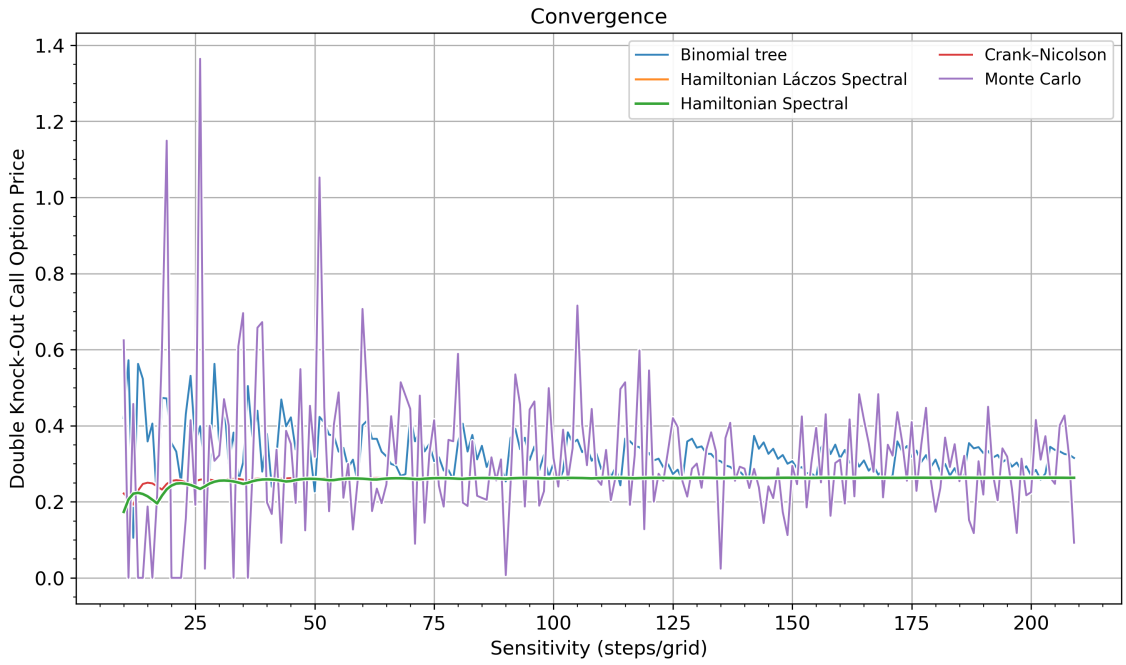

Figure 2.4: Double Knock-Out price convergence as a function of the sensitivity parameter

We also compare the computational runtimes of the examined algorithms in Figure 2.6.

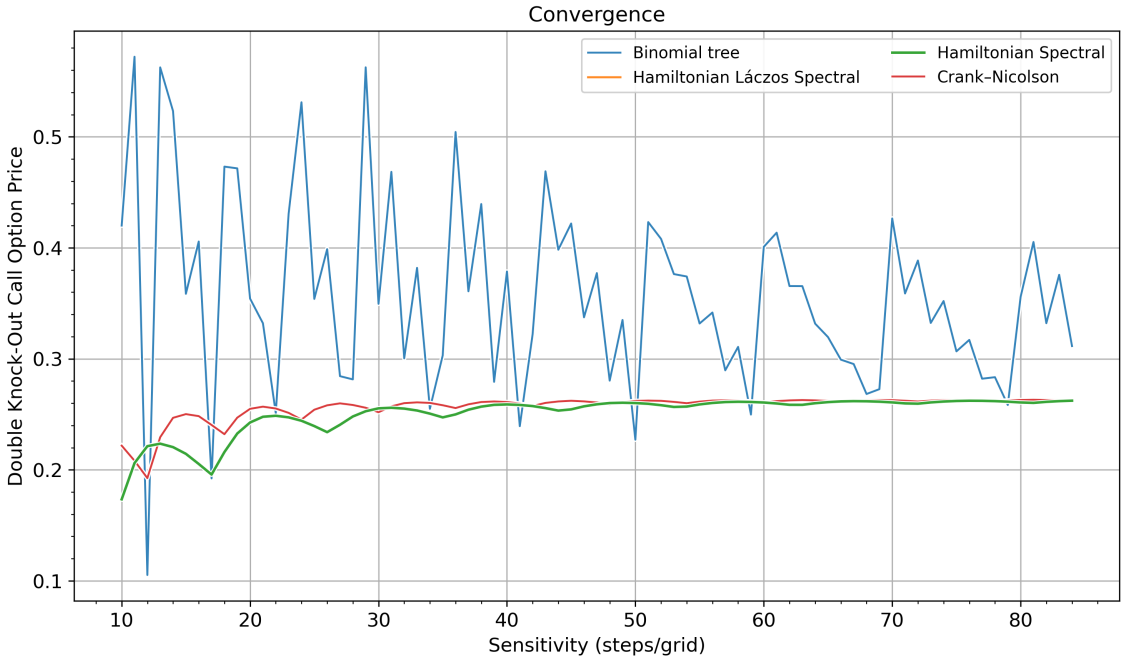

Figure 2.5: Double Knock-Out price convergence as a function of the sensitivity parameter without Monte Carlo

The figure shows that, in the present implementation, the main sources of runtime growth are the Monte Carlo and Crank-Nicolson methods. Among these two, Monte Carlo displays an approximately linear increase in runtime as the sensitivity parameter grows, whereas Crank-Nicolson shows an exponential increase in runtime, and it is cheaper under moderate sensitivity, but after that, it becomes the most time-consuming method among the examined algorithms. By contrast, the Hamiltonian spectral and binomial tree methods remain substantially less expensive on the same scale, while the Hamiltonian Lánczos spectral curve is so close to the Hamiltonian spectral one that their difference is not visually informative in the full comparison. To obtain a clearer comparison among the computationally cheaper methods, Figure 2.7 omits the Monte Carlo and Crank-Nicolson and focuses on the other deterministic methods. First, at very low sensitivity levels, the binomial tree is the fastest of the three methods. However, this advantage can be observed for very small matrices. As the sensitivity of the binomial tree increases, its runtime grows steadily and eventually exceeds the computational cost of Hamiltonian spectral methods. Second, the Hamiltonian Lánczos spectral method is initially slower than the conventional Hamiltonian spectral implementation up to a sensitivity of approximately 120-130. This behaviour is consistent with the additional setup overhead associated with the Lánczos procedure, including the shift-invert preprocessing and the construction of the Krylov-subspace approximation, which can outweigh the benefits of truncation. Beyond this crossover region, however, the Hamiltonian Lánczos spectral method becomes clearly more efficient. Its runtime remains nearly flat over the rest of the tested interval and stays uniformly low relative to the other deterministic competitors. This is a particularly

important observation, since it shows that, for the double knock-out option, the Lánczos-based approach not only preserves the spectral accuracy established earlier but also offers a markedly more favourable runtime profile when high-precision results are required. The dense Hamiltonian spectral method, by contrast, exhibits a more pronounced upward trend because it computes the full spectrum, thereby increasing its computational cost more directly with the matrix dimension. These results indicate that the path-dependent case exhibits the same fundamental trade-off already observed in the plain-vanilla setting. The reported runtimes should therefore be interpreted primarily as relative comparisons under a common computational setup, since their absolute values may vary across hardware, software environments, and implementations, even though the qualitative differences in scaling behaviour are expected to remain similar.

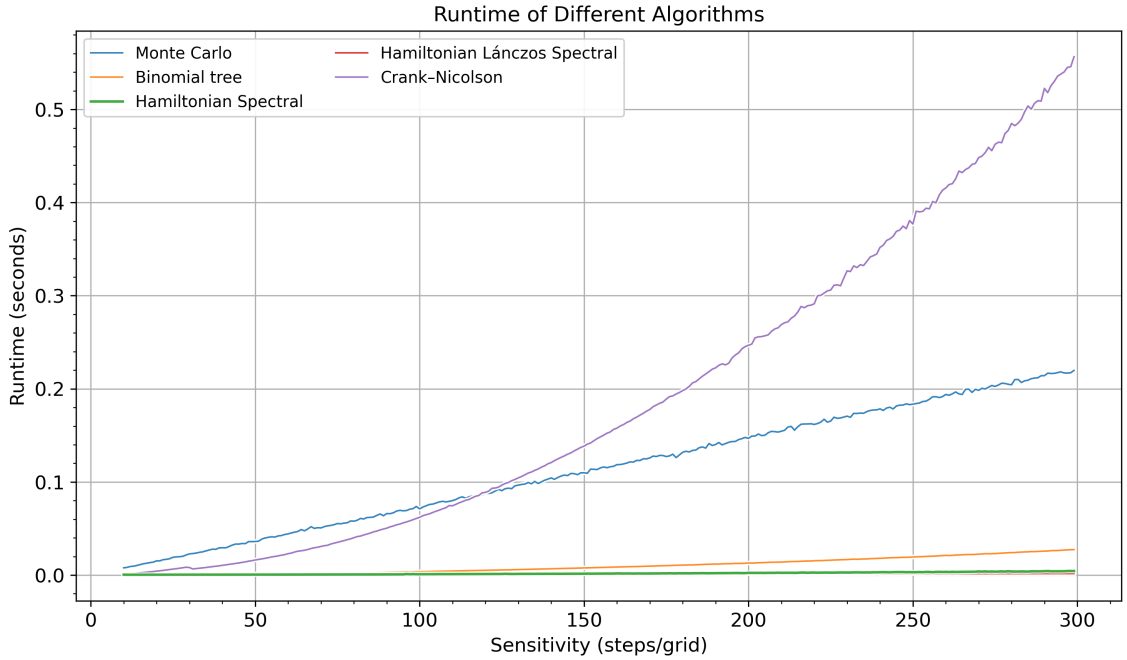

Figure 2.6: Runtime as a function of sensitivity for double knock-out pricing

If even higher accuracy is required than that presented in the numerical examples, a natural strategy is to increase the sensitivity and/or decrease the tolerance parameter $\varepsilon_{\text{tail}}$ in the Hamiltonian Lánczos spectral method so that a larger number of eigenpairs is retained in the truncated spectral expansion. Large changes in either the sensitivity or the tolerance parameter, however, can substantially increase the computational cost. A possible way to mitigate part of this additional cost is to execute the shift-invert Lánczos procedure on GPU hardware rather than on a CPU. This is particularly plausible because the dominant computational kernels in Lánczos-type sparse eigensolvers are sparse linear-algebra operations, whereas in shift-invert procedures, the repeated solution of shifted linear systems is the principal bottleneck (Woo et al. (2024)). In the GPU implementation of the filtered Lánczos procedure by Aurentz et al. (2017), the computation time is reduced

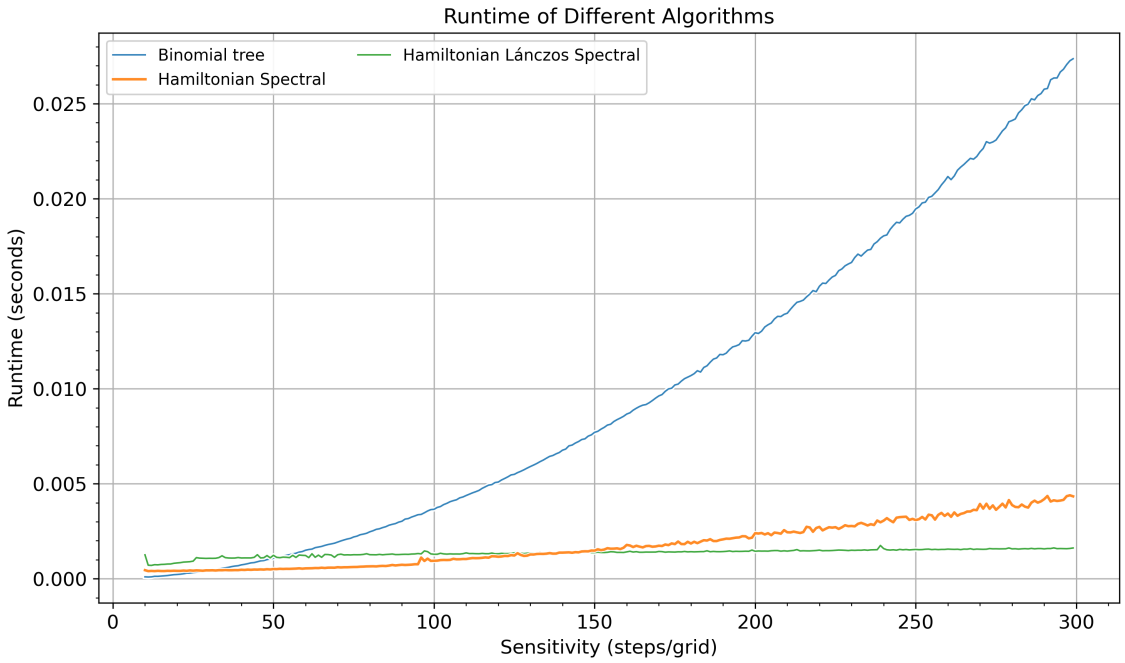

Figure 2.7: Runtime as a function of sensitivity for double knock-out pricing without Monte Carlo and Crank-Nicolson algorithms

by more than a factor of 10 relative to an equivalent CPU implementation, showing that order-of-magnitude gains are realistic when the runtime is dominated by sparse linear algebra. Looking further ahead, recent advances in quantum Lánczos and quantum Krylov subspace diagonalisation suggest that, for substantially larger Hamiltonians and higher-dimensional settings, quantum Krylov-type approaches may eventually provide an attractive route for accelerating the estimation of low-lying eigenvalues (Kirby et al. (2023), Epperly et al. (2022)). In particular, Cortes and Gray (2022) show that quantum Krylov subspace diagonalisation can provide a low-cost alternative to conventional quantum phase estimation for estimating ground and excited state energies, which further supports its relevance as a promising future direction for large-scale spectral computations.

2.4.4 Risk Sensitivity Diagnosis and Digital Option Valuation

From a risk-management perspective, estimating risk sensitivities is among the most accuracy-sensitive tasks in derivative valuation, because hedge ratios, rebalancing decisions, and trading limits depend directly on local price sensitivities rather than on prices alone. Consequently, even relatively small local errors in an estimated state price density may propagate into materially larger errors when translated into risk measures or special payoff values, such as digital options. This point is especially important for Greeks and digital-style payoffs because the economically relevant difference between competing numerical schemes may be apparent only in several decimal places. A cash-or-nothing claim is discontinuous at the strike, so its value is determined by the integral of the

state-price density. As a result, cash-or-nothing prices and Greek estimates are more accuracy-sensitive than the option price itself and therefore require high-precision density estimation and stable numerical differentiation (Hull (2023), Hull and Basu (2016), Wilmott (2013)). To illustrate how the examined simulations are treated, we present representative results for both the plain-vanilla and the double-knockout contracts. The plain-vanilla Delta can be recovered from the Arrow-Debreu state-price density according to Equation (2.42), while the corresponding cash-or-nothing digital call value is obtained from Equation (2.43). For the double knock-out contract, the same digital quantity is evaluated with the barrier-conditioned surviving density in Equation (2.44), the corresponding Delta is computed by the centred finite-difference approximation in Equation (2.45), and in this case the survival and knock-out probabilities follow from Equations (2.46) and (2.47). In the plain-vanilla case, the cash-or-nothing digital call value coincides with the standard digital call value and in the double knock-out case, it represents the price of a surviving digital payoff (Hull and Basu (2016), Wilmott (2013)). Table 2.5 shows that the Hamiltonian spectral and Hamiltonian Lánczos spectral methods provide the most accurate plain-vanilla risk sensitivity estimate and digital option valuation, with relative errors of order 10^{-4} for both Delta and the digital-call value. The Lánczos reduction is practically indistinguishable from the full spectral diagonalisation, which again indicates that the truncated eigenpair set is sufficient for this risk measure and digital call pricing estimation. Monte Carlo and the binomial tree are less accurate than the Hamiltonian-based methods but remain reasonably stable, whereas Crank-Nicolson is materially less accurate for these plain-vanilla strike-sensitive functionals in the present test. This indicates that even when Crank-Nicolson remains highly competitive for option-price estimation, its density reconstruction on the common truncated grid may still generate larger local errors near the strike, and these are amplified in Delta and digital-call estimates.

$$\Delta_{\text{van}}(S_0) = \frac{1}{S_0} \int_{\ln K}^{\infty} e^x \, \pi_{0,T}^{\text{van}}(\ln S_0, x)\, dx \tag{2.42}$$

$$D_0^{\text{CoN,van}}(S_0) = \int_{\ln K}^{\infty} \pi_{0,T}^{\text{van}}(\ln S_0, x)\, dx \tag{2.43}$$

$$D_0^{\text{CoN,DKO}}(S_0) = \int_{\ln K}^{\infty} \pi_{0,T}^{\text{DKO}}(\ln S_0, x)\, dx \tag{2.44}$$

$$\begin{aligned}\Delta_{\text{DKO}}(S_0) \approx \frac{1}{2\,\delta S} \Bigg[&\int_{\ln K}^{\ln U} (e^x - K)\, \pi_{0,T}^{\text{DKO}}(\ln(S_0 + \delta S), x)\, dx \\ &- \int_{\ln K}^{\ln U} (e^x - K)\, \pi_{0,T}^{\text{DKO}}(\ln(S_0 - \delta S), x)\, dx \Bigg]\end{aligned} \tag{2.45}$$

$$P_{\text{surv}}^{\text{DKO}}(0,T) = e^{rT} \int_{\ln L}^{\ln U} \pi_{0,T}^{\text{DKO}}(\ln S_0, x)\, dx \tag{2.46}$$

$$P_{\mathrm{KO}}^{\mathrm{DKO}}(0,T) = 1 - P_{\mathrm{surv}}^{\mathrm{DKO}}(0,T) \tag{2.47}$$

where $\pi_{0,T}^{\mathrm{van}}$ denotes the plain-vanilla Arrow-Debreu state-price density, $\pi_{0,T}^{\mathrm{DKO}}$ denotes the Arrow-Debreu state-price density for the double knock-out contract, $\Delta_{\mathrm{van}}(S_0)$ is the plain-vanilla Delta reconstructed from the state-price density, $D_0^{\mathrm{CoN,van}}(S_0)$ is the plain-vanilla cash-or-nothing digital call value, $D_0^{\mathrm{CoN,DKO}}(S_0)$ is the corresponding digital-call value under the double knock-out barrier structure, $\Delta_{\mathrm{DKO}}(S_0)$ is the double knock-out Delta estimated by centred finite differences, $P_{\mathrm{surv}}^{\mathrm{DKO}}(0,T)$ and $P_{\mathrm{KO}}^{\mathrm{DKO}}(0,T)$ are the survival and knock-out probabilities of the double knock-out contract up to maturity, S_0 is the spot price at valuation time, K is the strike price, L and U are the lower and upper limit prices, and $x = \ln S_T$ is the terminal log-price (Linetsky (2007), Davydov and Linetsky (2003), Hull and Basu (2016), Wilmott (2013)).

Table 2.5: Plain-vanilla delta and digital-call estimations from the state-price densities, the relative errors are estimated against the analytical Black–Scholes–Merton value

	Monte Carlo	Binomial Tree	Crank–Nicolson	Hamiltonian	Hamiltonian Lánczos
Delta	1.9367×10^{-1}	1.9397×10^{-1}	1.8872×10^{-1}	1.9335×10^{-1}	1.9335×10^{-1}
rel.err	1.8172×10^{-3}	3.3977×10^{-3}	2.3812×10^{-2}	1.4703×10^{-4}	1.4703×10^{-4}
Digital	1.3935×10^{-1}	1.3950×10^{-1}	1.3581×10^{-1}	1.3905×10^{-1}	1.3905×10^{-1}
rel.err	2.2644×10^{-3}	3.3578×10^{-3}	2.3254×10^{-2}	1.1820×10^{-4}	1.1820×10^{-4}

For the double knock-out option, the barrier-conditioned state-price density enables inference not only of sensitivity measures but also of the survival and knock-out probabilities up to maturity. Since the surviving digital-call value and the finite-difference Delta are both reconstructed from the same state price density, it is natural to examine these probability quantities jointly with the corresponding sensitivity estimates. Table 2.6 shows that the Hamiltonian spectral, Hamiltonian Lánczos spectral, and Crank-Nicolson methods imply the same knock-out probability to the reported four-decimal precision, namely about 10.10%. This is consistent with the fact that the Hamiltonian and Crank-Nicolson approaches impose absorbing barriers on the same continuous log-domain and therefore approximate the same continuous-time barrier event within the implemented discretisation. By contrast, the Monte Carlo and binomial estimates are lower in this experiment, which is consistent with discrete monitoring conventions, under which barrier crossings between observation times may be missed, and, in the Monte Carlo case, with residual sampling noise. Table 2.6 also shows that the Hamiltonian Lánczos spectral method matches the dense Hamiltonian spectral reference up to machine precision for both Delta and the surviving digital-call value again. Crank-Nicolson also matches the Hamiltonian benchmark extremely closely in this experiment, its relative error is of order 10^{-6} for Delta and of order 10^{-7} for the surviving digital-call value, while its implied

knock-out and survival probabilities coincide with the Hamiltonian values to the reported table precision. By contrast, the binomial tree produces a severely distorted Delta and even the wrong sign in this experiment, while Monte Carlo remains materially noisier than the deterministic solvers. This difference in behaviour is not surprising, since the surviving digital value is obtained as a direct integral of the barrier-conditioned state-price density, whereas the double knock-out Delta is computed by centred finite differences applied to the state-price density and therefore amplifies both lattice-alignment effects in the tree and sampling noise in Monte Carlo more strongly than the digital payoff itself.

Table 2.6: Double knock-out survival probabilities, delta and digital-call estimations, the relative errors are estimated against the Hamiltonian spectral estimation

	Hamiltonian	Hamiltonian Lánczos	Crank-Nicolson	Binomial Tree	Monte Carlo
Delta	1.2976×10^{-2}	1.2976×10^{-2}	1.2976×10^{-2}	-2.3734×10^{-3}	1.4418×10^{-2}
rel.err	0	2.2139×10^{-12}	7.8467×10^{-6}	1.1829×10^{0}	1.1113×10^{-1}
Digital	5.3827×10^{-2}	5.3827×10^{-2}	5.3827×10^{-2}	5.5784×10^{-2}	5.8556×10^{-2}
rel.err	0	7.3866×10^{-13}	1.9978×10^{-7}	3.6364×10^{-2}	8.7851×10^{-2}
KO prob.	10.0979%	10.0979%	10.0979%	9.6912%	9.2967%
Surv. prob.	89.9021%	89.9021%	89.9021%	90.3088%	90.7033%

2.5 Conclusion

This study develops a practical numerical implementation of the pseudo-Hermitian Black-Scholes-Merton framework and benchmarks it against several of the most widely used classical pricing algorithms in computational finance. Within this framework, Arrow-Debreu state prices arise naturally as a representation of the Euclidean-time propagator, and this representation can be exploited to construct high-accuracy Hamiltonian simulations not only for option values but also for state-price densities, risk sensitivities, digital option values, and knock-out event probabilities. Building on the similarity transformation proposed by Jana and Roy (2012), we propose two computational realisations of the framework, a conventional dense Hamiltonian spectral decomposition and a Krylov-subspace-based shift-invert Lánczos Hamiltonian spectral approximation. The numerical results obtained for the plain-vanilla and double knock-out case studies show that the Hamiltonian-based methods are fully competitive with, and in several respects superior to, the classical benchmark methods considered here. In the plain-vanilla setting, where an analytical Black-Scholes-Merton solution is available, the Hamiltonian spectral approach delivers the smallest pricing error and the closest state-price-density approximation to the

analytical benchmark. In the path-dependent double knock-out setting, where no closed-form benchmark exists, the Hamiltonian spectral and Crank-Nicolson methods produce highly consistent state-price densities, while the Hamiltonian Lánczos approximation remains almost indistinguishable from the full spectral solution up to machine precision. This confirms the study of Davydov and Linetsky (2003), in Euclidean-time propagation, the dominant pricing information is concentrated in a relatively small subset of low-lying eigenpairs, so that the neglected high-frequency part of the spectrum is of negligible practical importance.

The study also shows that the Hamiltonian formulation is particularly attractive when the objective goes beyond price estimation alone. Since the state-price density is obtained directly, the same computational procedure can be reused to recover derived quantities such as digital call values, Delta estimates, and knock-out probabilities. In liquid markets, this state-price framework also provides a natural basis for Vanna-Volga-type volatility smile option price adjustments. Vanna-Volga corrects a reference Black-Scholes-Merton price by matching the vega, vanna, and volga of the target payoff to liquid benchmark quotes, thereby incorporating observed smile effects with minimal recalibration (Fisher (2007)). Within the present framework, once the baseline spectral solution is available, the model price and related sensitivities can be obtained efficiently within the same Hamiltonian setup, making such corrections computationally inexpensive. In the numerical experiments, spectral methods exhibit particularly strong performance for quantities that are locally sensitive to the density structure, such as cash-or-nothing digital payoffs and Delta-type risk sensitivities. This is important from a risk-management perspective because such sensitivities are typically more demanding numerically than prices themselves and are often the relevant inputs in hedging and exposure control. The results therefore indicate that the pseudo-Hermitian Hamiltonian approach is not merely a pricing alternative but a unified framework for the simultaneous estimation of several practically relevant risk measures.

From a computational perspective, the results clarify the roles of the two proposed spectral implementations. Full diagonalisation is attractive on small grids because the discretised Hamiltonian is symmetric and tridiagonal, which makes the dense spectral solver robust and straightforward to use. However, its cost increases unfavourably as the desired accuracy grows and the grid is refined. By contrast, the shift-invert Lánczos variant exploits the exponential suppression of high-energy modes under Euclidean-time propagation and therefore targets only the low-lying spectral window required for accurate valuation. In the numerical experiments, this yields a clear break-even point for modest grid sizes, after which the Lánczos approach delivers essentially the same accuracy as full diagonalisation, with substantially improved runtime scaling. The framework is not

restricted to the specific contracts examined here. Accordingly, the numerical Hamiltonian framework is expected to derive its main practical relevance from applications that move beyond the analytically solvable one-factor spectral class, including models with stochastic-volatility dynamics and higher-dimensional settings in which semi-analytical approaches are either unavailable or computationally burdensome (Linetsky (2007), Lin and Ritchken (2006)). The additional value of the proposed methods lies not in competing with the closed-form spectral formulas available for a limited class of analytically tractable one-factor models, but in providing a reusable numerical operator framework for settings where such formulas are unavailable or cumbersome to derive. The proposed methodology can be extended to other diffusion models studied in the spectral-pricing literature, including CIR, CEV, Vasicek-type, and related settings, as well as to further path-dependent exotic derivatives, such as Asian options (Linetsky (2007)).

A further natural future direction is the higher-dimensional extension suggested by the multi-asset pseudo-Hermitian formulation of Léber (2025). In that setting, the relevance of Lánczos-type Krylov reduction is likely to become even greater, although it would need to be combined with additional structure-exploiting techniques, such as sparse-grid, tensor, or low-rank approximations, in order to mitigate the curse of dimensionality. From a computational perspective, the proposed framework is also well aligned with future acceleration strategies. Since the shift-invert Lánczos procedure is dominated by sparse linear-algebra kernels and repeated shifted linear solves, it is naturally compatible with GPU-based implementations in larger-scale settings (Woo et al. (2024), Aurentz et al. (2017)). More broadly, because the methodology is based on operator evolution and the estimation of low-lying eigenpairs, it is also conceptually connected to emerging quantum Lánczos and quantum Krylov-subspace diagonalisation algorithms, which may offer further opportunities for large-scale spectral pricing in the future (Kirby et al. (2023), Epperly et al. (2022), Cortes and Gray (2022)). These considerations further support the pseudo-Hermitian Hamiltonian framework as a scalable direction for both classical and future hardware-oriented derivative pricing.

In the next chapter, we broaden this framework to multiple correlated assets, exploring the valuation of multi-asset European claims and path-dependent rainbow options.

Chapter 3

Quantum Mechanical Approach to Pricing Multi-Asset Path-Dependent Options

Building on the framework introduced in Chapter 2, this chapter generalises the framework to valuate derivatives driven by multiple correlated underlying assets. We derive a pseudo-Hermitian representation of the multi-asset Black–Scholes–Merton Hamiltonian, develop a semi-analytic and numerical solutions for path-dependent multi-asset options, and illustrate the approach on a multi-asset rainbow option with asset-specific barriers, thereby demonstrating that the generalised framework can provide fast and accurate valuation. The results presented here build on the work published in the journal Physica A: Statistical Mechanics and its Applications (Léber (2025)).

3.1 Introduction

In this chapter, we focus on the pricing of financial derivatives, especially multi-asset options. Various types of multi-asset options can be distinguished on the basis of their payoff functions (Hull and Basu (2016)). Basket options, characterised by payoffs derived from the weighted average of multiple correlated assets, are particularly notable for providing diversified exposure in financial markets. Quanto options address exchange rate risk by providing payouts in a currency different from the underlying assets, thereby offering an effective hedge against currency exchange rate risk. Spread options further expand the range of possibilities by basing their value on the price difference between two or more assets, making them an attractive choice for those seeking to profit from relative price movements. The value of the rainbow options depends on the conditions specific to each underlying asset. A rainbow option with a worst-of feature determines its payout on the basis of the performance of the weakest asset in the basket, while a best-of

rainbow option derives its value from the strongest performing asset (Rubinstein (1991), Stulz (1982), Johnson (1987)). Together, these options allow investors to implement sophisticated risk management strategies and take advantage of complex market dynamics. However, pricing such instruments accurately requires adjustments to the standard multi-asset Black-Scholes-Merton (BSM) framework to account for their path-dependent structures. Although the BSM framework has well-known limitations, such as assumptions of constant volatility and constant interest rates, it remains one of the cornerstone of option pricing in practice, where its shortcomings can be corrected through various methodologies, including volatility adjustment techniques like the Vanna-Volga pricing method (Fisher (2007)).

The field of advanced option pricing models has seen substantial development, incorporating a wide range of mathematical frameworks, numerical techniques, and quantum-inspired methodologies. The literature spans diverse approaches, including pseudo-Hermitian formulations, supersymmetry approaches, Feynman path integrals, and various numerical techniques, to explore alternative methodologies for solving complex option pricing models (Jana and Roy (2012, 2011), Linetsky (1997), Araneda and Villena (2021), Baaquie (2007)). Vukovic (2015) highlights parallels between the Schrödinger equation and the BSM equation, the research explores the deep connections between these two fundamental partial differential equations, particularly focusing on the nature of their respective Hamiltonians. The study demonstrates that while the Schrödinger Hamiltonian is Hermitian, the BSM Hamiltonian is non-Hermitian. Jana and Roy (2012) similarly conclude that the BSM Hamiltonian operator is non-Hermitian but demonstrate that it is pseudo-Hermitian. They construct the metric operator $\hat{\eta}$ for the original BSM Hamiltonian, providing a pseudo-Hermiticity of the BSM Hamilton operator. In a subsequent study, Jana and Roy (2011) use supersymmetric techniques to identify the isospectral partners of BSM Hamiltonians, including those with a double knock-out barrier potential. The double knock-out barrier option is modelled via a square-well potential, and supersymmetric techniques are employed to derive partner Hamiltonians with identical boundary conditions and comparable spectra. Baaquie et al. (2004) present a comprehensive framework for pricing path-dependent options by integrating Hamiltonian mechanics, the Feynman path integral formulation, and simulation techniques such as the Langevin and Monte Carlo methods. To model barrier options accurately, this chapter incorporates a potential function within the Hamiltonian framework, which constrains the stock price's stochastic process to reflect the specific conditions governing barrier options. Arbitrage is traditionally excluded from the BSM framework, but Haven (2002) introduces a specific arbitrage potential function, which is linked to the concept of a belief parameter. The chapter further interprets the Planck constant as a proxy for the degree of belief or limited rationality in financial markets. Yeşiltaş (2023) solves the Black-Scholes-Merton

equation via exceptional Laguerre polynomials and Darboux transformations, which are extended to higher-order supersymmetric cases. The study also identifies a new potential via the Darboux transformation of the heat equation, offering quantum-inspired insights into financial market analysis. Hicks (2020) introduces a quantum framework for financial derivatives, in which the Heisenberg framework is utilised instead of the more popular Schrödinger approach. Governed by the Heisenberg equation of motion, the model incorporates quantum interference effects with significant implications for derivative pricing. The study highlights how non-trivial phase factors can influence time value, while quantum measurement processes add complexity to traditional pricing models, even when the market wave function is real valued.

At the time of writing, to the best of our knowledge, the literature has not yet addressed the pricing of options on the basis of multiple correlated underlying assets inspired by quantum mechanical approaches. One of the most widely used approaches for pricing such options in the literature is the application of various numerical methods. Peherstorfer et al. (2015) address the curse of dimensionality by using sparse grid discretisation and reduced models via proper orthogonal decomposition for multi-asset basket options under the Black-Scholes-Merton and Heston models. Khodayari and Ranjbar (2019) propose a computationally efficient method that combines radial basis functions (RBFs) for position derivatives and finite differences for time derivatives to solve the multidimensional BSM equation. Their approach avoids solving dense linear systems and demonstrates high accuracy for complex multi-asset pricing problems. Another widely applied numerical approach for pricing multi-asset options is Monte Carlo simulation, as demonstrated by Rasulov et al. (2016) in their evaluation of rainbow options. In this work, we contribute to the literature by generalising the framework of the pseudo-Hermitian Black-Scholes-Merton Hamiltonian, as introduced by Jana and Roy (2012), to accommodate multiple correlated underlying assets. We identify the corresponding $\hat{\eta}$-metric operator, which allows the construction of a Hermitian multi-asset BSM Hamiltonian operator. Using Fourier transformation, we seek the solution for the general multi-asset option in the form of a momentum-space wave function, analogous to determining the wave function of a free particle in quantum mechanics. Furthermore, leveraging the analogies of potential energy as outlined by Baaquie et al. (2004), we provide a solution to price a double knock-out rainbow option via widely recognised multiple particles in a box framework traditionally employed in quantum mechanics.

Quantum computing has demonstrated remarkable potential to revolutionise computational finance by optimising numerical algorithms specifically tailored for quantum hardware. Two popular approaches to quantum computational methods for option pricing are quantum-accelerated Monte Carlo simulations and quantum Hamiltonian simulations.

Quantum accelerated Monte Carlo simulations, as developed by Stamatopoulos et al. (2020), utilise amplitude estimation (AE) to achieve quadratic speed-ups. These methods have been extended to handle path-dependent derivatives and incorporate error mitigation techniques suitable for current quantum processors. However, one limitation of the proposed quantum Monte Carlo methods is that they require explicit knowledge of the underlying probability distribution. Quantum Hamiltonian simulations explicitly simulate the time evolution of a Hamiltonian operator, requiring the operator to be Hermitian. Gonzalez-Conde et al. (2023) employ the quantum signal processing method which, although effective, requires additional ancillary qubits to handle the non-Hermitian nature of the Black–Scholes–Merton operator, leading to noisier results compared to the case where a similar Hermitian operator would have been used. Beyond the analytical solutions presented in this chapter, the pseudo-Hermiticity of the generalised Black-Scholes-Merton Hamiltonian introduced here may also be leveraged in the future to develop more efficient quantum simulations for option pricing.

3.2 Temporal evolution of multi-asset European options as a quantum state

For a detailed and insightful introduction to the analogies, methodologies, and conceptual foundations used in the following context, the reader is referred to the comprehensive discussion provided by Baaquie (2007). The generalisation of the BSM partial differential equation on the basis of the price of multiple correlated assets is presented in Equation (3.1) (Wilmott (2013)). In financial theory, the Black-Scholes-Merton (BSM) equation governs the temporal evolution of the value of a European option $g(S,t)$ from the moment of its issuance to its expiration. Similarly, in quantum mechanics, the Schrödinger differential equation (3.2) governs the time evolution of a quantum state $|\Psi\rangle$ between two specified temporal points. Consequently, the corresponding differential equation in Euclidean space can be reformulated in Hilbert space, as shown in Equation (3.3), where the corresponding $\hat{H}_{BSM}$ operator serves a role analogous to that of the Hamiltonian operator in the Schrödinger Equation (3.2). The risk associated with the value of the option is intrinsically related to the behaviour of the logarithm of the stock price. Therefore, it is advantageous to express the $\hat{H}_{BSM}$ operator (3.4) in terms of the derivative operator by the x_i logarithm of the i-th stock price (3.5).

$$\frac{\partial g(S,t)}{\partial t} = rg(S,t) - \sum_{i=1}^{N} rS_i \frac{\partial g(S,t)}{\partial S_i} - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} \rho_{ij}\sigma_i\sigma_j S_i S_j \frac{\partial^2 g(S,t)}{\partial S_i \partial S_j} \tag{3.1}$$

$$\frac{d|\Psi\rangle}{dt} = -i\hat{H}|\Psi\rangle \tag{3.2}$$

$$\frac{d|g\rangle}{dt} = \hat{H}_{BS}|g\rangle \tag{3.3}$$

$$\hat{H}_{BSM} = r - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}\rho_{ij}\sigma_i\sigma_j S_i S_j \frac{\partial^2}{\partial S_i \partial S_j} - \sum_{i=1}^{N} rS_i \frac{\partial}{\partial S_i} \tag{3.4}$$

$$\hat{H}_{BSM} = r - \sum_{i=1}^{N}\left(r - \frac{1}{2}\sigma_i^2\right)\frac{\partial}{\partial x_i} - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}\rho_{ij}\sigma_i\sigma_j \frac{\partial^2}{\partial x_i \partial x_j} \tag{3.5}$$

where, $g(S,t)$ represents the function of the option value, r is a constant risk-free interest rate, σ_i is the i-th stock's constant volatility, ρ_{ij} is the constant Pearson correlation between the i-th and the j-th stocks, N is the number of underlying assets, $|\Psi\rangle$ represents a specific quantum state, $|g\rangle$ represents an option's quantum state, x_i represents the logarithm of the i-th stock price, and $\hat{H}_{BSM}$ is an operator satisfying the Black-Scholes-Merton transformation.

The wave function $|g\rangle$ represents the value of an option that is not subject to normalisation condition (3.6) because the discounted states of probabilities are not necessarily summed up to one, unlike the integral of the absolute square of a particle's wave function $|\Psi\rangle$, which must be normalised to one. In quantum mechanics, the result of a measurement can correspond only to real eigenvalues. However, according to Vukovic (2015) and Baaquie (2007), since the BSM Hamiltonian is non-Hermitian, complex eigenvalues may also arise as a consequence. Therefore, the Hamilton operator does not satisfy the Hermiticity condition outlined in Eq. (3.7), because the representation of the BSM Hamilton operator is inherently non-symmetric due to the presence of first-order derivative operators. According to Jana and Roy (2012), the single-asset BSM Hamiltonian operator is pseudo-Hermitian, as it satisfies the parity-time ($\mathcal{PT}$) symmetry conditions. Bender and Boettcher (1998) propose parity-time symmetry, which extends traditional quantum mechanics by exploring non-Hermitian Hamiltonians that exhibit real eigenvalues under certain conditions. These specific Hamiltonians, although they are not Hermitian, are $\mathcal{PT}$ symmetric, which means that they are invariant under the combined operations of parity and time reversal. An operator is said to be pseudo-Hermitian if there exists a Hermitian invertible metric operator $\hat{\eta}$ that satisfies Equation (3.8). $\mathcal{PT}$ symmetry implies pseudo-Hermiticity because there exists an indefinite metric operator that renders the Hamiltonian pseudo-Hermitian according to Eq. (3.9). Our aim is to identify $f(\mathbf{x})$ as the target function for the construction of the $\hat{\eta}$ and $\hat{\gamma}$ operators. This function is aligned with the multi-asset Black-Scholes-Merton framework and facilitates the formulation of

the Hermitian representation of the BSM Hamiltonian operator.

$$\int_{-\infty}^{\infty} |\Psi|^2 dx = 1 \neq \int_{-\infty}^{\infty} |g|^2 dx \tag{3.6}$$

$$\hat{M}^{\dagger} = \left[\hat{M}^{T}\right]^{*} \tag{3.7}$$

$$\hat{H}^{\dagger} = \hat{\eta}\hat{H}\hat{\eta}^{-1} \tag{3.8}$$

$$\begin{aligned} \hat{h}^{\dagger} &= \hat{\gamma}\hat{H}\hat{\gamma}^{-1} \\ \hat{\gamma} &= \hat{\eta}^{1/2} = e^{f(\mathbf{x})} \end{aligned} \tag{3.9}$$

$$f(\mathbf{x}) = \sum_{i=1}^{N} a_i x_i \tag{3.10}$$

where, Ψ represents a specific quantum wave function, g represents an option's wave function, $\hat{H}$ is a non-Hermitian operator, $\hat{M}$, $\hat{\eta}$, $\hat{\gamma}$ and $\hat{h}^{\dagger}$ are Hermitian operators, N is the number of underlying assets, $f(\mathbf{x})$ is a desired function that constitutes part of the operator $\hat{\eta}$ and $\hat{\gamma}$, and a_i represents its coefficients.

The effect of the desired $\hat{\gamma}$ operators on an arbitrary non-Hermitian first-order (3.11) and a Hermitian second-order derivative operator (3.12) can be expressed. To ensure that the BSM Hamilton operator (3.13) becomes Hermitian, the coefficients of the corresponding first-order derivatives should be zero. Economically, $\hat{\eta}$ changes the weighting of log-price states and its square-root similarity map $\hat{\gamma}$ cancels the vector of risk-neutral log-return drifts only from the first-order terms. These drifts are retained in the endpoint $\hat{\gamma}$ factors, while the zero-order term $\hat{V}$ carries discounting together with the drift correction, so the original BSM valuation is unchanged. The preferred form of the $f(\mathbf{x})$ function can be determined if the system of differential Equations (3.14) is solved, where $\mathbf{\Sigma}$ denotes the constant instantaneous covariance matrix of log returns and is assumed to be positive definite, so that $\mathbf{\Sigma}^{-1}$ exists. The general solutions of the differential equation system are represented by Eq. (3.15). One of the possible representations of the $\hat{\gamma}$ metric operator, which we subsequently use, can be expressed in the form (3.16).

$$\hat{\gamma}\frac{\partial}{\partial x_i}\hat{\gamma}^{-1} = \frac{\partial}{\partial x_i} - \frac{\partial f(\mathbf{x})}{\partial x_i} \tag{3.11}$$

$$\hat{\gamma}\frac{\partial^2}{\partial x_i \partial x_j}\hat{\gamma}^{-1} = \frac{\partial^2}{\partial x_i \partial x_j} - \frac{\partial f(\mathbf{x})}{\partial x_j}\frac{\partial}{\partial x_i} - \frac{\partial f(\mathbf{x})}{\partial x_i}\frac{\partial}{\partial x_j} + \frac{\partial f(\mathbf{x})}{\partial x_i}\frac{\partial f(\mathbf{x})}{\partial x_j} \tag{3.12}$$

$$\hat{h}^{\dagger} = r - \sum_{i=1}^{N} \left(r - \frac{\sigma_i^2}{2} \right) \left(\frac{\partial}{\partial x_i} - \frac{\partial f(\mathbf{x})}{\partial x_i} \right) - \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sigma_i \sigma_j \rho_{ij} \left(\frac{\partial^2}{\partial x_i \partial x_j} - \frac{\partial f(\mathbf{x})}{\partial x_j} \frac{\partial}{\partial x_i} - \frac{\partial f(\mathbf{x})}{\partial x_i} \frac{\partial}{\partial x_j} + \frac{\partial f(\mathbf{x})}{\partial x_i} \frac{\partial f(\mathbf{x})}{\partial x_j} \right) \tag{3.13}$$

$$\boldsymbol{\Sigma} \begin{pmatrix} \frac{\partial f(\mathbf{x})}{\partial x_1} \\ \frac{\partial f(\mathbf{x})}{\partial x_2} \\ \vdots \\ \frac{\partial f(\mathbf{x})}{\partial x_n} \end{pmatrix} = \begin{pmatrix} r - \frac{\sigma_1^2}{2} \\ r - \frac{\sigma_2^2}{2} \\ \vdots \\ r - \frac{\sigma_n^2}{2} \end{pmatrix} \tag{3.14}$$

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_1^2 & \sigma_1 \sigma_2 \rho_{12} & \cdots & \sigma_1 \sigma_n \rho_{1n} \\ \sigma_2 \sigma_1 \rho_{21} & \sigma_2^2 & \cdots & \sigma_2 \sigma_n \rho_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_n \sigma_1 \rho_{n1} & \sigma_n \sigma_2 \rho_{n2} & \cdots & \sigma_n^2 \end{pmatrix}$$

$$f(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\Sigma}^{-1} \left(r - \frac{1}{2} \sigma^2 \right) \tag{3.15}$$

$$\hat{\gamma} = e^{\mathbf{x}^T \boldsymbol{\Sigma}^{-1} \left(r - \frac{1}{2} \sigma^2 \right)} \tag{3.16}$$

where $\hat{\gamma}$ and $\hat{h}^{\dagger}$ are Hermitian operators, $f(\mathbf{x})$ is a desired function that constitutes part of the $\hat{\gamma}$ operator, $\mathbf{x}$ is a vector of logarithm of prices, r is a constant risk-free interest rate, σ is a vector of the constant volatility of the underlying stocks, ρ_{ij} is the constant Pearson correlation between the i-th and the j-th stocks, N is the number of underlying assets, and $\boldsymbol{\Sigma}$ is the covariance matrix of the underlying assets.

In the case of two underlying assets, the inverse of the covariance matrix can be formulated as Eq. (3.17). Using the corresponding $\hat{\gamma}$ operator, the general multi-asset extension of the Hermitian BSM Hamilton operator can be determined via Eq. (3.18). In the single-asset case, both the $\hat{\gamma}$ operator and the $\hat{h}^{\dagger}$ Hermitian BSM Hamilton operator coincide with the operators defined by Jana and Roy (2012). However, in the multi-asset case, there are also partial cross-derivatives, which, according to Young's theorem, result in a symmetric operator. According to Fernández (2020), if the coefficients of partial cross-derivatives are real in Hamilton operators, then these operators are also Hermitian. In Equation (3.19), μ (3.21) represents the vector of the adjusted drift rate of the underlying asset's price process, and $\hat{V}$ (3.20) represents the constant potential energy operator.

$$\boldsymbol{\Sigma}^{-1} = \frac{1}{\sigma_1^2\sigma_2^2(1-\rho_{12}^2)} \begin{pmatrix} \sigma_2^2 & -\rho_{12}\sigma_1\sigma_2 \\ -\rho_{12}\sigma_1\sigma_2 & \sigma_1^2 \end{pmatrix} \tag{3.17}$$

$$\hat{h}^{\dagger} = -\frac{1}{2}\hat{\nabla}^{\top}\boldsymbol{\Sigma}\hat{\nabla} + \frac{1}{2}\mu^{\top}\boldsymbol{\Sigma}^{-1}\mu + r \tag{3.18}$$

$$\hat{h}^{\dagger} = -\frac{1}{2}\hat{\nabla}^{\top}\boldsymbol{\Sigma}\hat{\nabla} + \hat{V} \tag{3.19}$$

$$\hat{V} = \frac{1}{2}\mu^{\top}\boldsymbol{\Sigma}^{-1}\mu + r \tag{3.20}$$

$$\mu = \begin{bmatrix} r - \frac{\sigma_1^2}{2} \\ r - \frac{\sigma_2^2}{2} \\ \vdots \\ r - \frac{\sigma_n^2}{2} \end{bmatrix} \tag{3.21}$$

where $\hat{h}^{\dagger}$ is the Hermitian representation of the original BSM Hamilton operator, μ represents the vector of the adjusted drift rate of the underlying asset's price process, r is a constant risk-free interest rate, σ_i is a i-th stock's constant volatility, ρ_{ij} is the constant Pearson correlation between the i-th and the j-th stocks, $\boldsymbol{\Sigma}$ is the covariance matrix of the underlying assets, $\hat{V}$ is a constant potential operator, and $\hat{\nabla}$ incorporates differential operators corresponding to the different underlying assets.

3.3 Solution of Multi-Asset European Options and Free Particles in Space

In the case of the multi-asset BSM Hamiltonian operator, the potential term remains constant. In accordance with Vukovic (2015) and Araneda and Villena (2021), if the $\hbar$ Planck constant is set to 1, then the mass and momentum operator can be expressed according to Eq. (3.22). The interaction between two underlying assets in the multi-asset BSM framework can be effectively captured by the product of their respective momentum operators shown in Eq. (3.23), while the magnitude and direction of the interaction are governed by the ρ constant Pearson correlation coefficient.

$$
\begin{aligned}
m_1 &= \frac{1}{\sigma_1^2} \\
p_1 &= -i\frac{\partial}{\partial x_1}
\end{aligned}
\tag{3.22}
$$

$$
\hat{h}^\dagger = \frac{p_1^2}{2m_1} + \frac{\rho_{1,2}}{\sqrt{m_1 m_2}} p_1 p_2 + \frac{p_2^2}{2m_2} + \hat{V} \tag{3.23}
$$

where m_1 and the p_1 operator are derived from the BSM that corresponds to the equivalent mass and momentum operator related to the first underlying asset, $\hat{h}^\dagger$ is a two-asset Hermitian BSM Hamilton operator, ρ is a constant Pearson correlation coefficient, and $\hat{V}$ is the corresponding potential operator.

A free particle is characterised by the absence of a confining potential; therefore, the particle can freely move through space without any boundary constraints. Common boundary conditions such as Dirichlet and Neumann conditions do not apply here (Riley et al. (2006)). The Dirichlet boundary condition requires the wave function to vanish at a specific location. This condition arises in systems with infinite potential barriers, such as a particle in a box framework, but since a free particle moves without boundaries, there is no reason for the wave function to be zero at any point. Similarly, the Neumann boundary condition, which states that the derivative of the wave function is zero at a given boundary, does not bind. This condition typically appears in symmetric systems, but a free particle does not have a boundary that requires such symmetry. The time-independent Schrödinger equation can be expressed according to Equation (3.24). The eigenstates of the Hamilton operator are generally expressed in the form (3.25) that must satisfy the normalisation condition (3.6), which ensures that the total probability of finding the particle within the allowed region of space is equal to one. However, in the case of a free particle, the eigenstates are described by plane waves that extend infinitely across space. These states cannot be normalised in the conventional L^2 sense. Here the configuration space is the ordinary, non-extended space $\mathbb{R}^N$, so the symbols $\pm\infty$ in normalisation and Fourier integrals indicate integration over an unbounded domain rather than admissible coordinate values. The plane waves are therefore understood as Dirac-normalised generalised eigenstates in the standard Hilbert-space sense, whereas normalisable physical states belong to $L^2(\mathbb{R}^N)$.

$$
\hat{H}\phi(x) = E\phi(x) \tag{3.24}
$$

$$
\phi(x) = Ae^{ikx} \tag{3.25}
$$

where $\hat{H}$ is a specific Hamilton operator, E is the value of energy, $\phi(x)$ is the eigenstate of a given Hamilton operator, and A and k are constants.

The Fourier transform is an ideal choice for free particle problems because it naturally represents the system in terms of momentum spectra, which are more appropriate than position spectra due to the particle's delocalisation in space (Griffiths and Schroeter (2018)). Additionally, it circumvents the normalisation issue by using the Dirac delta function to ensure the orthogonality of the momentum eigenstates. The corresponding Schrödinger equation for the momentum-space wave function of the two-asset option is given by Equation (3.26). As the BSM Schrödinger equation is not formulated as an initial value problem but rather assumes a known terminal state, we introduce the τ variable to ensure the reversal of time. The $\tau = T - t$ affects only the time derivative term, resulting in a two-asset BSM Schrödinger Equation (3.27).

$$\frac{d\Phi}{dt} = \left(\frac{p_1^2}{2m_1} + \frac{\rho_{1,2}}{\sqrt{m_1 m_2}} p_1 p_2 + \frac{p_2^2}{2m_2} \right) \Phi + \hat{V}\Phi \tag{3.26}$$

$$\frac{d\Phi}{d\tau} = -\left(\frac{p_1^2}{2m_1} + \frac{\rho_{1,2}}{\sqrt{m_1 m_2}} p_1 p_2 + \frac{p_2^2}{2m_2} \right) \Phi - \hat{V}\Phi \tag{3.27}$$

where Φ is the momentum-space wave function of the option price, m_1 and p_1 are derived from the BSM that corresponds to the equivalent mass and the momentum operator related to the first underlying asset, ρ is a constant Pearson correlation coefficient, $\hat{V}$ is a corresponding potential operator, and τ expresses the time remaining until expiration of the option.

To solve the two-asset Black-Scholes-Merton Schrödinger equation, with the Planck constant $\hbar$ set to 1, the Fourier transform is used to convert the position-wave function $\Psi(\mathbf{x})$ into the momentum-wave function $\Phi(\mathbf{p})$. Conversely, the inverse Fourier transform allows the momentum-wave function to be transformed back into the position-wave function. These transformations, which are based on Fourier analysis principles, follow the approaches of Riley et al. (2006), Griffiths and Schroeter (2018), and use the properties outlined in Equations (3.28), (3.29) and (3.31).

$$\begin{aligned} \Phi(\mathbf{p}) &= \frac{1}{(2\pi)^{(N/2)}} \int_{\mathbb{R}^N} \Psi(\mathbf{x}) e^{-i\mathbf{p}\mathbf{x}} \, d^N\mathbf{x} \\ \Psi(\mathbf{x}) &= \frac{1}{(2\pi)^{(N/2)}} \int_{\mathbb{R}^N} \Phi(\mathbf{p}) e^{i\mathbf{p}\mathbf{x}} \, d^N\mathbf{p} \end{aligned} \tag{3.28}$$

$$\begin{aligned}\mathcal{F}\left[\frac{\partial^n \Psi(x)}{\partial x^n}\right] &= (ip)^n \Phi(p) \\ \mathcal{F}[\Psi_1(\mathbf{x}) * \Psi_2(\mathbf{x})] &= \frac{1}{(2\pi)^{(N/2)}} \mathcal{F}[\Psi_1(\mathbf{x})]\mathcal{F}[\Psi_2(\mathbf{x})]\end{aligned} \tag{3.29}$$

where the convolution means that $(\Psi_1 * \Psi_2) = \int_{-\infty}^{\infty} \Psi_1(x)\Psi_2(y-x)dx$, $\mathcal{F}[\Psi(\mathbf{x})]$ is a Fourier transform of the position-space wave function $\Psi(\mathbf{x})$, p is a momentum operator and $\Psi(\mathbf{p})$ is a momentum-space wave function.

First, we solve the ordinary differential Equation (3.27) according to Eq. (3.30). Using the covariance matrix of the underlying assets, the exponential term in the solution can also be expressed as the Fourier transform of the multivariate normal probability distribution function. The Fourier transform of the general multivariate normal probability distribution function is shown in Eq. (3.31), where N denotes the number of variables. If the covariance matrix is scaled by the τ scalar (3.32), then the solution of $\Phi(\mathbf{p})$ can be expressed as Eq. (3.33).

$$\Phi(p_1, p_2) = \Phi(p_T)e^{-\hat{V}\tau} \exp\left\{\left(-\frac{p_1^2}{2m_1} - \frac{\rho_{1,2}}{\sqrt{m_1 m_2}} p_1 p_2 - \frac{p_2^2}{2m_2}\right)\tau\right\} \tag{3.30}$$

$$\mathcal{F}\left[\frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}} \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{\Sigma}^{-1}\mathbf{x}\right)\right] = \frac{1}{(2\pi)^{(N/2)}} \exp\left(-\frac{1}{2}\left(\mathbf{p}^T\mathbf{\Sigma}^{-1}\mathbf{p}\right)\tau\right) \tag{3.31}$$

$$\mathbf{\Sigma}^* = \mathbf{\Sigma}\tau \tag{3.32}$$

$$\Phi(\mathbf{p}) = \Phi(\mathbf{p_T})e^{-\hat{V}\tau}(2\pi)^{(N/2)}\mathcal{F}\left[\frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}} \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{\Sigma}^{*-1}\mathbf{x}\right)\right] \tag{3.33}$$

where $\Phi(\mathbf{p})$ is the momentum-space wave function of the option price, $\Phi(\mathbf{p_T})$ is the momentum-space wave function of the option price at maturity, m and p are derived from the BSM that corresponds to the equivalent mass and the momentum operator, ρ is a constant Pearson correlation coefficient, $\hat{V}$ is the corresponding potential operator, $\mathbf{\Sigma}$ is a covariance matrix of the underlying assets, $|\mathbf{\Sigma}|$ is the determinant of the covariance matrix, N is the number of underlying assets, and τ expresses the time remaining until expiration.

By taking the Fourier transform of $\Phi(\mathbf{p_T})$ the momentum-space wave function of the option price at maturity (3.34); and subsequently applying the inverse Fourier transform along with the convolution theorem, the Hermitian partial solution of the multi-asset BSM

Schrödinger equation (3.35) is obtained.

$$\Phi(\mathbf{p}) = e^{-\hat{V}\tau}(2\pi)^{(N/2)}\mathcal{F}\left[\Psi(\mathbf{x_T}) * \frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}}\exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{\Sigma}^{*-1}\mathbf{x}\right)\right] \tag{3.34}$$

$$\Psi(\mathbf{x}) = e^{-\hat{V}\tau}\frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}}\int_{\mathbb{R}^N}\Psi(\mathbf{x_T})\exp\left(-\frac{1}{2}\mathbf{y}^T\mathbf{\Sigma}^{*-1}\mathbf{y}\right)d^N\mathbf{x_T} \tag{3.35}$$

where:

$$\mathbf{y} = \mathbf{x} - \mathbf{x_T} \tag{3.36}$$

$\Phi(\mathbf{p})$ is the momentum-space wave function of the option price, $\Psi(\mathbf{x})$ is the position-space wave function of the option price, $\Psi(\mathbf{x_T})$ is the position space wave function of the option price at maturity, $\hat{V}$ is the corresponding potential operator, $\mathbf{\Sigma}$ is a covariance matrix of the underlying assets, $|\mathbf{\Sigma}|$ is the determinant of the covariance matrix, N is the number of underlying assets, and τ expresses the time remaining until expiration.

We identify the states associated with the original non-Hermitian BSM Hamilton operator to investigate the solution of the original multi-asset BSM partial differential equation. By applying the metric operator $\hat{\gamma}$, the payoff function $\Psi(\mathbf{x_T})$ of the Hermitian system is transformed, and the final wave function $\Psi(\mathbf{x})$ must be subsequently transformed, according to Equations (3.37). Consequently, the solution to the general multi-asset BSM equation is expressed in Equation (3.38). Following the decomposition and simplification of the potential operator $\hat{V}$, the discount factor can be factored out of the integral, leading to the final result (3.39).

$$\begin{aligned}\Psi(\mathbf{x_T}) &= \hat{\gamma}\Psi_{BSM}(\mathbf{x_T})\\ \Psi_{BSM}(\mathbf{x}) &= \hat{\gamma}^{-1}\Psi(\mathbf{x})\end{aligned} \tag{3.37}$$

$$\Psi_{BSM}(\mathbf{x}) = e^{-\hat{V}\tau}\frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}}\int_{\mathbb{R}^N}\Psi_{BSM}(\mathbf{x_T})\hat{\gamma}(-\mathbf{y})\exp\left(-\frac{1}{2}\mathbf{y}^T\mathbf{\Sigma}^{*-1}\mathbf{y}\right)d^N\mathbf{x_T} \tag{3.38}$$

$$\Psi_{BSM}(\mathbf{x}) = \frac{1}{(2\pi\tau)^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}}e^{-r\tau}\int_{\mathbb{R}^N}\Psi_{BSM}(\mathbf{x_T})\exp\left(-\frac{1}{2}\left(\mathbf{y}+\mu^*\right)^T\mathbf{\Sigma}^{*-1}\left(\mathbf{y}+\mu^*\right)\right)d^N\mathbf{x_T} \tag{3.39}$$

where:

$$\mu^* = \mu\tau \tag{3.40}$$

$\Psi_{BSM}(\mathbf{x})$ and $\Psi(\mathbf{x})$ are the original and Hermitian solutions of the multi-asset option price, $\Psi_{BSM}(\mathbf{x_T})$ and $\Psi(\mathbf{x_T})$ are the original and Hermitian payoff functions of the multi-

asset option at maturity, $\hat{V}$ is the potential operator, $\mathbf{\Sigma}$ is a covariance matrix of the underlying assets, $|\mathbf{\Sigma}|$ is the determinant of the covariance matrix, N is the number of underlying assets, μ represents the vector of adjusted drift rates of the underlying assets, r is a constant risk-free interest rate, and τ expresses the time remaining until expiration.

The final solution can alternatively be expressed in terms of the stock price S, by reversing the transformation $x = log(S)$. Since $x_i = \log S_i$ maps the open half-line $S_i \in (0, \infty)$ bijectively onto $\mathbb{R}$, the corresponding stock-price domain is the open orthant $(0, \infty)^N$, with neither 0 nor $+\infty$ included as a coordinate value. In this case, the final solution of the multi-asset Black-Scholes-Merton PDE is given by Formula (3.41), which aligns with the solution presented by Wilmott (2013) and Contreras G. et al. (2016).

$$\begin{aligned}\Psi_{BSM}(\mathbf{S}) = &\frac{e^{-r(T-t)}}{(2\pi(T-t))^{\frac{N}{2}}\sqrt{|\mathbf{\Sigma}|}} \\ &\times \int_{\mathbb{R}^N} \frac{\Psi_{BSM}(\mathbf{S_T})}{\mathbf{S_T}} \exp\left[-\frac{1}{2}(\mathbf{u}+\mu^*)^T \mathbf{\Sigma}^{*-1}(\mathbf{u}+\mu^*)\right] d^N\mathbf{S_T}.\end{aligned} \tag{3.41}$$

where:

$$\mathbf{u} = \log \frac{\mathbf{S}}{\mathbf{S_T}} \tag{3.42}$$

$$\mu^* = (\mathbf{r} - \frac{\sigma^2}{2})(T-t) \tag{3.43}$$

$\Psi_{BSM}(\mathbf{S})$ is the general solution of the multi-asset European plain vanilla option, $\Psi_{BSM}(\mathbf{S_T})$ is the payoff function of the option, $\mathbf{\Sigma}$ is a covariance matrix of the underlying assets, $|\mathbf{\Sigma}|$ is the determinant of the covariance matrix, N is the number of underlying assets, μ represents the vector of adjusted drift rates of the underlying assets, r is the constant risk-free interest rate, σ is the vector of the volatility of underlying assets, t is the observation time, and T is the expiration date.

3.4 Solution of Multi-Asset Double Knock-Out Options and Particles in a Box

A significant class of financial derivatives consists of path-dependent, limit-price options that include instruments on the basis of multiple underlying assets. These options often incorporate specific price thresholds, leading to two main categories of limit-price options. The first type, knock-out options, loses their value entirely if the price of the underlying asset crosses a predetermined threshold over time. In contrast, knock-in options start with no intrinsic value and become active only if the price of the underlying asset

reaches a specified level within a given time frame. In this example, we present the pricing of a two-asset double knock-out rainbow option that incorporates two individual limit prices for the underlying assets, as also examined by Zvan et al. (2000). This option value becomes zero as soon as one of the two underlying assets breaches its respective barrier level. Consequently, even if the second asset remains within its barrier range and performs favourably, the option is rendered worthless if the first asset breaches its limit.

In quantum mechanics, a common problem involves determining the position of a particle moving freely within a box. Issues of this nature are commonly referred to as particle in a box problems. In these scenarios, a time-evolving particle cannot leave the box because it is subject to forces pushing it back once it has reached its boundary. In this case, the energy levels of the particles are quantified, a property arising from the satisfaction of both the Dirichlet and Neumann boundary conditions. These conditions ensure that the particle's wave function can be expressed as a superposition of the eigenstates (stationary states) of the Hamiltonian, as determined by the general solution of the Schrödinger equation (3.44). The solutions of different multi-particle Schrödinger equations are comprehensively presented by Griffiths and Schroeter (2018).

$$|\Psi(t)\rangle = \sum_{n} \langle \varphi_{E_n} | \Psi(0) \rangle e^{-\frac{i}{\hbar} E_n t} | \varphi_{E_n} \rangle \tag{3.44}$$

where $|\Psi(t)\rangle$ is the wave function of the particle at time t, $|\Psi(0)\rangle$ is the initial state of the particle, $\hbar$ is the Planck constant, E_n is the energy eigenvalue of the n-th quantum state, and $|\varphi_{E_n}\rangle$ is the n-th energy eigenstate of the Hamiltonian.

We first need to determine the eigenvalues and eigenvectors of the corresponding Hamiltonian operator. Specifically, we identify the appropriate solution for the Hermitian multi-asset BSM Hamiltonian operator (3.19). The final result is subsequently modified via the appropriate $\hat{\gamma}$ operators (3.16) to account for the necessary adjustments. Nevertheless, the Hermitian BSM Hamiltonian operator also contains partial cross-derivatives, which can be eliminated by diagonalisation to achieve greater ease of solution. We introduce y_1 and y_2 according to Eq. (3.45), where θ (3.46) ensures that the cross-derivative operator vanishes from the corresponding Hamilton operator. As a result of diagonalisation, the Hermitian BSM Hamiltonian operator can be expressed in the form (3.47), where λ_1 (3.48) and λ_2 (3.49) represent the eigenvalues of the coefficient matrix. Economically, this is a principal-component rotation of risk-neutral log-return risk; λ_1 and λ_2 are variance rates along orthogonal, uncorrelated diffusion directions.

$$\begin{aligned} y_1 &= x_1 \cos\theta + x_2 \sin\theta \\ y_2 &= -x_1 \sin\theta + x_2 \cos\theta \end{aligned} \tag{3.45}$$

$$\theta = \frac{1}{2} \arctan\left(-\frac{2\sigma_1\sigma_2\rho_{1,2}}{\sigma_2^2 - \sigma_1^2}\right) \tag{3.46}$$

$$\hat{h}^\dagger = -\frac{1}{2}\left(\lambda_1 \frac{\partial^2}{\partial y_1^2} + \lambda_2 \frac{\partial^2}{\partial y_2^2}\right) + \hat{V} \tag{3.47}$$

$$\lambda_1 = \frac{\sigma_1^2 + \sigma_2^2 + \sqrt{(\sigma_1^2 - \sigma_2^2)^2 + 4\sigma_1^2\sigma_2^2\rho_{1,2}^2}}{2} \tag{3.48}$$

$$\lambda_2 = \frac{\sigma_1^2 + \sigma_2^2 - \sqrt{(\sigma_1^2 - \sigma_2^2)^2 + 4\sigma_1^2\sigma_2^2\rho_{1,2}^2}}{2} \tag{3.49}$$

where $\hat{h}^\dagger$ is the Hermitian BSM Hamilton operator, σ_i is the constant volatility of the i-th underlying stock, ρ_{ij} is the constant Pearson correlation between the i-th and the j-th stocks, and $\hat{V}$ is a corresponding potential operator.

In quantum mechanics, the boundary conditions of a particle are inherently encoded in the potential energy operator $\hat{V}$, as it is directly influenced by the position operator. One specific realisation of this principle is the infinite potential well, a system in which the particle is confined and cannot escape the well. Consequently, the path dependence of various limit price options can be effectively addressed through their potential energy operator $\hat{V}$. The potential energy operator for the desired option, which is based on two underlying assets, can be determined by the boundary condition (3.50), where the potential energy becomes infinite outside of the logarithm of the lower limit prices L and the higher limit prices H. Here $+\infty$ is shorthand for the hard-wall limit $V_0 \to +\infty$, not a finite energy value or a point of the spectrum. Correspondingly, the informal notation $1/\infty = 0$ means only $\lim_{V_0 \to +\infty} V_0^{-1} = 0$. The Hamiltonian is realised on $L^2(\Omega)$ with homogeneous Dirichlet boundary conditions, where $\Omega = (L_1, H_1) \times (L_2, H_2)$ is open and the barrier points form its boundary $\partial\Omega$. This condition can also be expressed in terms of y by applying the transformations (3.45) to the corner points of the original rectangular domain. The corresponding coordinate bounds (3.51) are then given by the minimum and maximum values of the transformed coordinates at these points. In general, these bounds define an axis-aligned rectangle enclosing the transformed domain. The particle-in-a-box construction below is formulated on this enclosing domain and yields a separable semi-analytical approximation to the original barrier problem.

$$\hat{V} = \begin{cases} \frac{1}{2}\mu^{\top}\boldsymbol{\Sigma}^{-1}\mu + r, & \text{if } L_1 < x_1 < H_1 \text{ and } L_2 < x_2 < H_2 \\ \infty, & \text{otherwise} \end{cases} \tag{3.50}$$

$$\begin{aligned} \overline{L_1} &= \min\left\{L_1\cos\theta + L_2\sin\theta,\ L_1\cos\theta + H_2\sin\theta,\right. \\ &\qquad \left. H_1\cos\theta + L_2\sin\theta,\ H_1\cos\theta + H_2\sin\theta\right\} \\ \overline{H_1} &= \max\left\{L_1\cos\theta + L_2\sin\theta,\ L_1\cos\theta + H_2\sin\theta,\right. \\ &\qquad \left. H_1\cos\theta + L_2\sin\theta,\ H_1\cos\theta + H_2\sin\theta\right\} \\ \overline{L_2} &= \min\left\{-L_1\sin\theta + L_2\cos\theta,\ -L_1\sin\theta + H_2\cos\theta,\right. \\ &\qquad \left. -H_1\sin\theta + L_2\cos\theta,\ -H_1\sin\theta + H_2\cos\theta\right\} \\ \overline{H_2} &= \max\left\{-L_1\sin\theta + L_2\cos\theta,\ -L_1\sin\theta + H_2\cos\theta,\right. \\ &\qquad \left. -H_1\sin\theta + L_2\cos\theta,\ -H_1\sin\theta + H_2\cos\theta\right\} \end{aligned} \tag{3.51}$$

where $\hat{V}$ is the potential operator of the double knock-out Hermitian BSM Hamilton operator, μ represents the vector of the adjusted drift rate of the underlying asset's price process, r is a constant risk-free interest rate, $\boldsymbol{\Sigma}$ is the covariance matrix of the underlying assets, x_i is a logarithm of the i-th underlying asset's price, L_i is a logarithm of the i-th underlying asset's lower limit price, $\overline{L_i}$ is the transformed lower barrier of the L_i, H_i is a logarithm of the i-th underlying asset's higher limit price, and $\overline{H_i}$ is the transformed higher barrier of the H_i.

The stationary states of the two-asset Hermitian BSM Hamilton operator, can be represented as a product of independent stationary states (3.52). Individual stationary states are sought in the well-established form presented in Eq. (3.53). Moreover, by applying the appropriate boundary conditions, the quantised values of k, which represent the wavenumbers, can be determined via Eq. (3.54). The constants A and B are expressed via the normalisation condition (3.6). Consequently, the stationary state $|\varphi_{n,m}\rangle$ of the Hermitian BSM Hamilton operator is normalised when expressed in the form (3.55).

$$|\varphi\rangle = |\varphi_1\rangle|\varphi_2\rangle \tag{3.52}$$

$$\begin{aligned} |\varphi_1\rangle &= A\sin(k_1 y_1) + B\cos(k_3 y_1) \\ |\varphi_2\rangle &= C\sin(k_2 y_2) + D\cos(k_4 y_2) \end{aligned} \tag{3.53}$$

$$
\begin{aligned}
|\varphi_1\rangle &= A \sin\left(\frac{n\pi\left(y_1 - \overline{L_1}\right)}{\overline{H_1} - \overline{L_1}}\right) \\
|\varphi_2\rangle &= B \sin\left(\frac{m\pi\left(y_2 - \overline{L_2}\right)}{\overline{H_2} - \overline{L_2}}\right)
\end{aligned}
\tag{3.54}
$$

$$
|\varphi_{n,m}\rangle = \frac{\sqrt{2}}{\sqrt{\left(\overline{H_1} - \overline{L_1}\right)}} \frac{\sqrt{2}}{\sqrt{\left(\overline{H_2} - \overline{L_2}\right)}} \sin\left(\frac{n\pi\left(y_1 - \overline{L_1}\right)}{\overline{H_1} - \overline{L_1}}\right) \sin\left(\frac{m\pi\left(y_2 - \overline{L_2}\right)}{\overline{H_2} - \overline{L_2}}\right) \tag{3.55}
$$

where $|\varphi_{n,m}\rangle$ is the stationary state of the two-asset Hermitian BSM Hamilton operator, L_i is a logarithm of the i-th underlying asset's lower limit price, $\overline{L_i}$ is the transformed lower barrier of the L_i, H_i is a logarithm of the i-th underlying asset's higher limit price, $\overline{H_i}$ is the transformed higher barrier of the H_i, k is a wavenumber, A and B are normalisation constants.

The corresponding $E_{n,m}$ energies can be determined via a time-independent Schrödinger equation (3.56). This yields discrete spectra of energies because the boundary conditions allow only for specific quantised wavenumbers. Using the operators associated with the logarithm of the original price of the underlings, the $|\varphi_{n,m}\rangle$ stationary states (3.57) and their corresponding $E_{n,m}$ (3.58) energies satisfy the Schrödinger differential equation along with the boundary conditions for each underlying asset.

$$
-\frac{1}{2}\left(\lambda_1 \frac{\partial^2 |\varphi_{n,m}\rangle}{\partial y_1^2} + \lambda_2 \frac{\partial^2 |\varphi_{n,m}\rangle}{\partial y_2^2}\right) + \hat{V}|\varphi_{n,m}\rangle = E_{n,m}|\varphi_{n,m}\rangle \tag{3.56}
$$

$$
\begin{aligned}
|\varphi_{n,m}\rangle = &\frac{2}{\sqrt{(\overline{H_1} - \overline{L_1})(\overline{H_2} - \overline{L_2})}} \\
&\sin\left(\frac{n\pi(x_1\cos\theta + x_2\sin\theta - \overline{H_1})}{\overline{H_1} - \overline{L_1}}\right) \sin\left(\frac{m\pi(x_2\cos\theta - x_1\sin\theta - \overline{H_2})}{\overline{H_2} - \overline{L_2}}\right)
\end{aligned}
\tag{3.57}
$$

$$
E_{n,m} = \frac{1}{2}\left(\lambda_1 \frac{n^2\pi^2}{\left(\overline{H_1} - \overline{L_1}\right)^2} + \lambda_2 \frac{m^2\pi^2}{\left(\overline{H_2} - \overline{L_2}\right)^2}\right) + \hat{V} \tag{3.58}
$$

where $|\varphi_{n,m}\rangle$ is the stationary state of the two-asset Hermitian BSM Hamilton operator, $E_{n,m}$ is the energy eigenvalue of the quantum state, L_i is a logarithm of the i-th underlying asset's lower limit price, $\overline{L_i}$ is the transformed lower barrier of the L_i, H_i is a logarithm of

the i-th underlying asset's higher limit price, $\overline{H_i}$ is the transformed higher barrier of the H_i, and $\hat{V}$ is a corresponding potential operator.

To express this approximation in terms of the original two-asset Black-Scholes-Merton model, the stationary states must be transformed via the corresponding metric operator $\hat{\gamma}$ (3.59). The present value of the two-asset double knock-out rainbow option can be expressed by Formula (3.60), where we have substituted the payoff function (3.61) corresponding to the worst-of rainbow option.

$$\begin{aligned} |\varphi_{BSM}\rangle &= \hat{\gamma}|\varphi_{n,m}\rangle \\ \langle\varphi_{BSM}| &= \hat{\gamma}^{-1}\langle\varphi_{n,m}| \end{aligned} \tag{3.59}$$

$$|\Psi_{BSM}(0)\rangle = \sum_n \sum_m c_{n,m} e^{-\left(\frac{1}{2}\left(\lambda_1 \frac{n^2\pi^2}{(\overline{H_1}-\overline{L_1})^2} + \lambda_2 \frac{m^2\pi^2}{(\overline{H_2}-\overline{L_2})^2}\right) + \hat{V}\right)T} |\varphi_{BSM}\rangle \tag{3.60}$$

where

$$\begin{aligned} c_{n,m} &= \langle\varphi_{BSM}|\Psi_{BSM}(T)\rangle \\ |\Psi_{BSM}(T)\rangle &= \max\left(\min(e^{x_1(T)}, e^{x_2(T)}) - K, 0\right) \end{aligned} \tag{3.61}$$

$|\Psi_{BSM}(0)\rangle$ is the two-asset double knock-out worst-of rainbow option's present value, $|\varphi_{BSM}\rangle$ is the stationary state of the original non-Hermitian BSM Hamilton operator, $|\varphi_{n,m}\rangle$ is the stationary state of the Hermitian BSM Hamilton operator, $\hat{\gamma}$ is a Hermitian operator that transforms the given state into the corresponding solution of the non-Hermitian system, $c_{n,m}$ are the coefficients determining the final state, $|\Psi_{BSM}(T)\rangle$ is the corresponding payoff function, $\hat{V}$ is a corresponding potential operator, T represents the time remaining until maturity, K represents the strike price, L_i is a logarithm of the i-th underlying asset's lower limit price, $\overline{L_i}$ is the transformed lower barrier of the L_i, H_i is a logarithm of the i-th underlying asset's higher limit price, and $\overline{H_i}$ is the transformed higher barrier of the H_i.

The semi-analytical Schrödinger solution of the two-asset option (3.60) can be efficiently approximated numerically. Figures (3.1), (3.2), (3.3), (3.4), (3.5), illustrate different scenarios of the prices of a two-asset double knock-out worst-of rainbow option as a function of the two underlying assets' spot prices. In the examples, both assets have a maturity of 3 months and share a risk-free rate of 5%, while the volatilities differ between the two underlyings, the first asset S_1 exhibits volatility 40% and the second asset S_2 exhibits volatility 20%. The strike price is set at \$100 for the first three scenarios, while in the fourth case, it is adjusted to \$110. The limit price levels are also varied, with both underlying assets having identical limits of \$90 and \$120 in the first three examples, and differing limits in the final case, where the first underlying asset has a range of \$100 to \$140, while the second underlying asset spans from \$90 to \$130. In the first case, the correlation between the two underlying assets is assumed to be positive and relatively low,

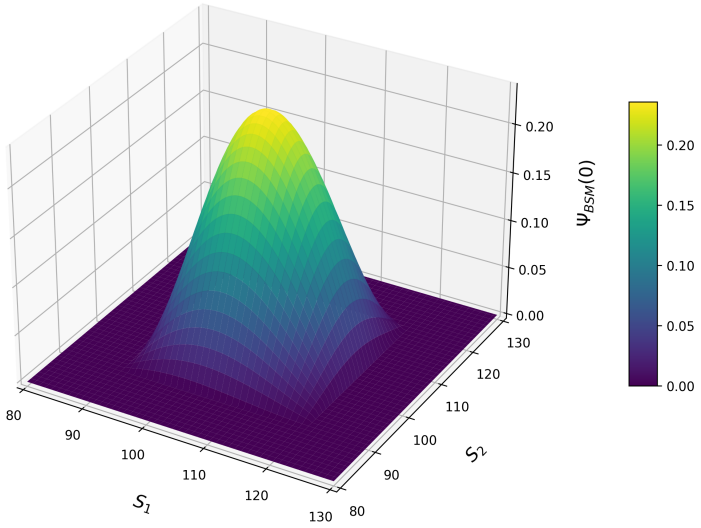

Figure 3.1: The surface plot of the two-asset double knock-out worst-of rainbow option price as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.05.

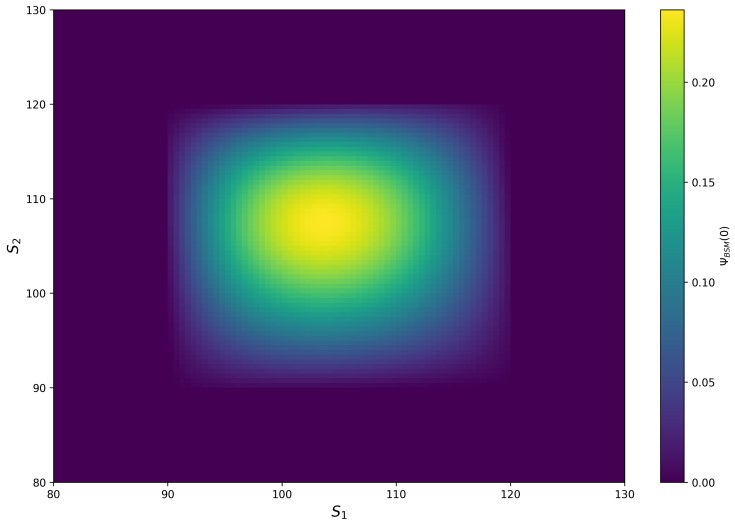

Figure 3.2: The heatmap of the two-asset double knock-out worst-of rainbow option price as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.05.

set at the level 0.05. In Figure (3.1), the pricing surface exhibits a bell-shaped distribution similar to a bivariate normal distribution. In the heatmap of Figure (3.2), we can observe that the high values of the option price lie above the strike price, corresponding to the region where the intrinsic value of the option becomes significantly higher, as both underlying assets move into the in the money range. The highest values occur when both asset prices are above the strike price and sufficiently distant from their respective limit prices, located approximately where both asset prices are near $S_1 = \$104$ and $S_2 = \$107$, achieving a maximum value of approximately $\Psi_{BSM}(0) \approx 0.24$. The high value region in

the heatmap is relatively symmetric, indicating that the low correlation between the assets ensures that changes in the spot price of one underlying asset have relatively independent effects on the option's price.

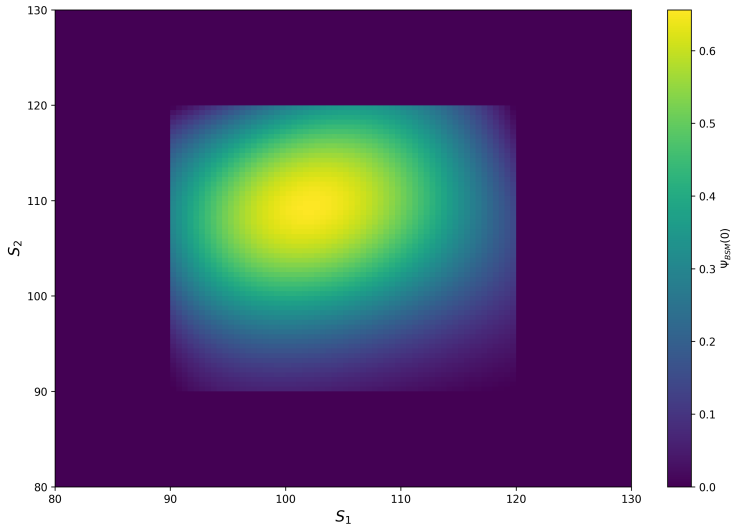

Figure 3.3: The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be 0.6.

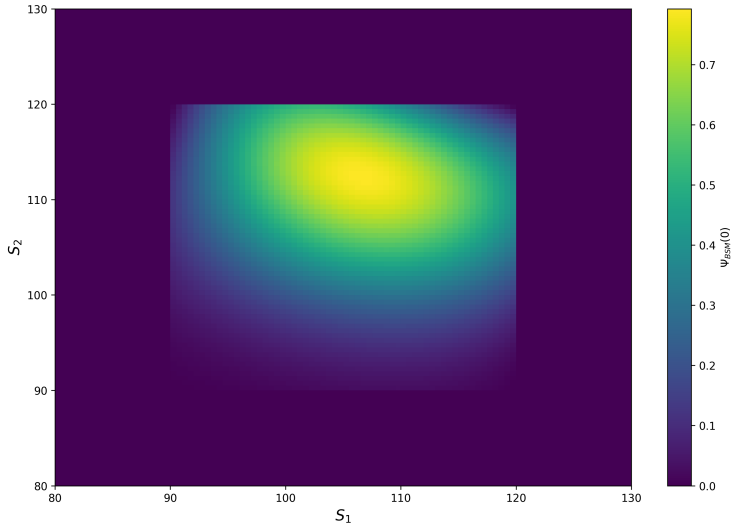

Figure 3.4: The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. The correlation between the two underlying assets is assumed to be -0.6.

Figure (3.3) corresponds to a scenario in which the correlation between the two underlying assets increases to 0.6 while all other parameters remain unchanged. In this case, the option prices exhibit a noticeable increase in magnitude compared with the lower correlation case, and the highest value is located approximately where both asset

prices are near S_1 = \$102 and S_2 = \$109, achieving a maximum value of approximately $\Psi_{BSM}(0) \approx 0.66$. The bright region, which represents high option values, becomes noticeably elongated. This elongation reflects the combined effects of the higher volatility of S_1 (40%) compared to S_2 (20%) and the increased correlation between the underlying assets. The bright region moves closer to the lower limit price of S_1, where a steep decline in the option value can be observed. Under lower correlation, the bell-shaped pricing surface is relatively symmetric, but with increased correlation, the high-value region becomes skewed and stretched.

In Figure (3.4), the correlation between the two underlying assets changes from a positive value to a negative value of -0.6, while all other parameters remain unchanged. The high value region shifts diagonally, with a stronger concentration in the upper right corner. This formalism reflects the inverse relationship between the two assets, where a negative correlation allows one asset's increase to counteract the other's decrease. Similarly to the high positive correlation case, the high value region is not symmetric and becomes more skewed. The bright region moves closer to the upper limit price of S_2, causing a sharper drop in the option value in that area. Overall, both high positive and high negative correlations result in skewed shapes of the pricing surface, shifting the high option value regions closer to the corresponding limit prices. The degree of the shift and elongation is influenced by the volatility structure of the underlying assets and leads to a steep decline in the option's value near the corresponding limit prices.

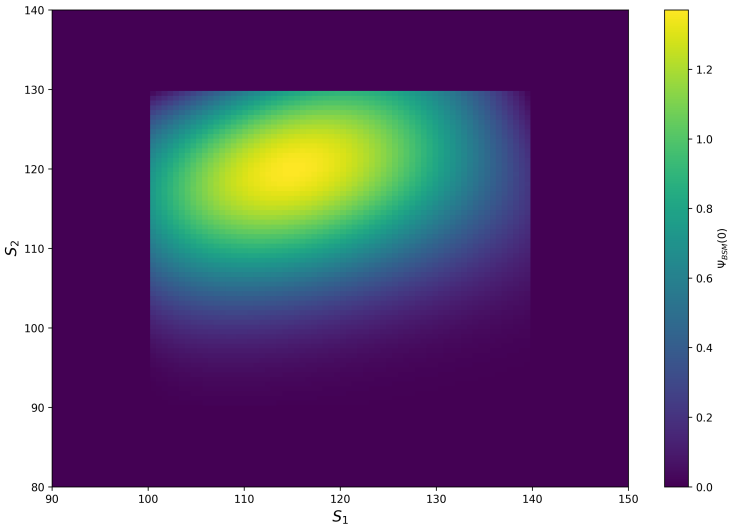

Figure 3.5: The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. S_1 has limits prices between \$100 and \$140, while S_2 has limit prices between \$90 and \$130. The correlation between the two underlying assets is assumed to be 0.6.

It is possible to apply different limit prices to the underlying assets, allowing for

increased flexibility of the conditions attached to the underlying assets. In Figure (3.5), compared with the positive case of the correlation 0.6, the limit prices are modified so that S_1 has limits between \$100 and \$140, while S_2 has limits between \$90 and \$130, with the strike price also adjusted to \$110. As a result of these changes, the region associated with high option values undergoes modifications in both shape and position. The high value region becomes more elongated along the S_1 axis due to the wider range of positions in the money, while simultaneously compressing along the S_2 due to the wider range of positions out of the money. Furthermore, the high value region shifts diagonally left upward, moving closer to the upper limit price of S_1 and the lower limit price of S_2. Furthermore, proximity of the high value region to the upper limit price of S_1 leads to a much steeper decline in the option value, as S_1 approaches its limit price. The modified limit prices result in a more distorted pricing surface, where the high value region is stretched and compressed in response to the adjusted structural parameters, demonstrating the sensitivity of multi-asset option pricing to changes in the barrier levels.

3.5 Conclusion

This study extends the pseudo-Hermitian Black–Scholes–Merton framework to accommodate multiple correlated underlying assets by identifying the appropriate $\hat{\eta}$-metric operator and constructing a Hermitian multi-asset BSM Hamiltonian operator. The transition to momentum-space offers a streamlined and elegant approach for solving the general multi-asset BSM differential equation. Furthermore, leveraging the well-established particles in a box framework from quantum mechanics enables efficient treatment of various financial instruments' path dependence, while the time-independent Schrödinger equation offers a clear semi-analytical solution for similar problems. The efficacy of this methodology is illustrated through the pricing of a two-asset double knock-out rainbow option, where even distinct barriers can be specified for each underlying asset. The presented methods can be easily generalised to options with more than two underlying assets, as well as to the pricing of other multi-asset options, such as basket options. Taken together, these findings emphasise the versatility of quantum-mechanical methodology for capturing the complexity of correlated multi-asset financial instruments and underscore the potential for future Schrödinger-based simulations to further advance theoretical understanding and practical applications in derivative pricing.

While the models discussed to this point assume a constant, time-invariant volatility parameter, they can be further developed by allowing volatility to vary over time or to depend on additional parameters, such as strike price or sentiment. The next chapter, therefore, investigates how volatility can be forecast more accurately using data-driven

approaches rather than traditional models, a capability crucial to both option pricing and risk management. In Chapter 4, we incorporate various sentiment indices into a proposed hybrid econometric and recurrent neural network model, demonstrating how different sentiments can enhance risk management, particularly Value-at-Risk forecasting.

Chapter 4

The Sentiment Augmented GARCH-LSTM Hybrid Model for Value-at-Risk Forecasting

Having established quantum-inspired pricing models for single-asset and multi-asset derivatives in Chapters 2 and 3, this final chapter turns to the empirical assessment of volatility and risk forecasts. We develop a sentiment-augmented GARCH–LSTM hybrid model that integrates news- and social-media-based sentiment indicators with traditional volatility models to improve Value-at-Risk forecasting accuracy. This work was carried out jointly with Balázs Egyed and builds on results published in Computational Economics (Léber and Egyed (2025)).

4.1 Introduction

Accurate volatility forecasting constitutes a cornerstone in the field of financial risk management and portfolio optimisation. The capacity to predict market turbulence with precision is imperative for the formulation of robust risk mitigation strategies and the optimisation of asset allocation decisions. In this context, Value-at-Risk has emerged as one of the most popular risk metrics. According to Duffie and Pan (1997) Value-at-Risk (VaR) is a statistical technique used to measure and quantify the level of financial risk of a portfolio over a specific time frame. It estimates the maximum loss that an investment portfolio is likely to suffer, given normal market conditions, over a period of time with a certain degree of confidence. VaR has gained popularity because it is easy to estimate and provides a single, summary measure of market risk for the entire range of financial instruments within a firm's portfolio.

Value-at-Risk has come under a lot of attack, Artzner et al. (1999) highlight that VaR is

not considered a coherent risk measure because of its lack of subadditivity. Subadditivity ensures that the risk of a combined portfolio should not exceed the sum of risks of individual portfolios. According to these authors violations of subadditivity with VaR can result in financial institutions underestimating risk, which can have implications both for internal risk management and for regulatory capital requirements. As a result of these concerns, Expected Shortfall (ES), which is an alternative risk measure that is coherent and inherently subadditive has been proposed. On the other hand, Daníelsson et al. (2013) showed that Value-at-Risk can indeed be used as a measure of risk in most practical applications, as it is generally subadditive for common types of assets. His paper finds that concerns regarding VaR's lack of subadditivity are not of significant importance, as violations are rare and mostly confined to assets with extremely fat tails or specific probability levels within return distributions. As such, for typical assets and applications, VaR remains a reliable and operationally simpler choice for risk measurement compared to other more complicated risk measures.

The Basel Committee on Banking Supervision provides guidelines and standards for banks to manage market risks, including requirements for capital reserves, risk assessment models, and specific reporting practices. As of this writing, the committee provides two possible measures of market risk, Value-at-Risk (VaR) and Expected Shortfall (ES). The move to Expected Shortfall from Value-at-Risk is significant because of the coherence of Expected Shortfall. Nonetheless, one salient challenge that arises with the adoption of Expected Shortfall is the intricacy of its backtesting. Regardless of the choice of the risk metric, VaR backtesting remains a key concept in the Basel framework due to its more straightforward application in contrast to backtesting of the Expected Shortfall. The criteria for backtesting necessitate the assessment of each trading desk's daily Value-at-Risk calculations at the 97.5% percent confidence level (Basel Committee on Banking Supervision, 2013, 2019).

There are different methods for estimating Value-at-Risk and considering stylised facts in these models is essential because they encapsulate the empirical properties observed in asset returns, which help predict market behaviour. In this context, Cont (2001) lists four important stylised facts, volatility clustering, conditional heavy tails, the leverage effect and intermittency. One of the most popular volatility forecast models is the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, which was introduced by Bollerslev (1986). This model can capture the phenomena of volatility clustering and conditional heavy tails in financial time series. Degiannakis et al. (2012) show that the GARCH model has the most satisfactory performance for the VaR forecasts before the financial crisis of 2008 for various indices, while during the financial crisis of 2008, the Asymmetric Power GARCH (APARCH) model was found to be more capable of

forecasting VaR efficiently because it is better at capturing extreme negative log returns and the asymmetry of the leverage effect.

On the other hand, for the models to meet these stylised facts, researchers have begun to develop artificial intelligence models, most of which are built on the predictions of GARCH models. Donaldson and Kamstra (1997) reported that the artificial neural network (ANN) model provides the best performance in terms of forecasting volatility because these models are capable of capturing complex, nonlinear relationships in the data that other models cannot. Hu et al. (2020) combined a GARCH model with Long Short-Term (LSTM) and Bidirectional Long Short-Term (BiLSTM) neural networks because these models have memory components and they are useful for capturing time-series characteristics, especially considering sequences of a past event. Kim and Won (2018) presented a comprehensive approach to forecasting stock price index volatility by integrating an LSTM neural network with multiple GARCH-type models such as the Exponential GARCH (EGARCH) model. Their findings also revealed that the hybrid LSTM models significantly outperform existing methodologies, achieving lower prediction errors across various metrics. Kakade et al. (2022) demonstrated that hybrid BiLSTM-GARCH models outperform every conventional model including hybrid GARCH-LSTM models also due to their ability to capture bidirectional dependencies in time series data, providing a more comprehensive understanding of the underlying volatility. García-Medina and Aguayo-Moreno (2024) show another particular advantage of GARCH-LSTM models, they present a powerful tool for predicting the volatility of cryptocurrency portfolios and optimising portfolio performance according to them. Applying a different approach, Buczynski and Chlebus (2023) introduce a neural network-based GARCH model and show its particular importance in the aspect of Value-at-Risk. Compared with GARCH, the GARCHNet models demonstrate better outcomes during the backtesting. The number of exceedances, which indicates when the actual losses surpass the predicted Value-at-Risk, is lower for GARCHNet models in several cases. GARCHNet models often have lower cost function values from the regulator's perspective, indicating potentially lower costs related to risk management activities. Wang et al. (2023) highlighted the advantages of incorporating multiple features of neural network-based models for forecasting cryptocurrency volatility. The authors identified key variables beyond lagged volatility, these include trading volume, blockchain metrics, Google Trends data, NASDAQ, S&P 500 adjusted close prices and policy uncertainty indices. By utilising these diverse internal and external determinants, multivariable neural network models outperform traditional models such as GARCH in terms of forecasting accuracy.

Our study extends the capabilities of conventional GARCH and GARCH-LSTM augmented neural network architectures by attaching sentiment indices drawn from both

conventional and social media platforms. Our objective is to enhance the accuracy of the existing models in forecasting market volatility and Value-at-Risk. We analyse the connection between market variances and the sentiment indices, with a particular focus on investigating the nonlinear connection between these variables. Drawing from the foundation established by prior research and our significant results of Granger causality tests, we develop two advanced models: a sentiment index augmented GARCH model and a sentiment index augmented GARCH-LSTM model. These enhanced models are designed to capitalise on the predictive signals embedded within sentiment data, which serve as proxies for investor mood and market perception, thereby providing a more nuanced understanding of market behaviours and future volatility. Grounded in our initial postulations, the integration of sentiment indices is posited not only to encapsulate the characteristically observed stylised facts of fat tails and volatility clustering within financial time series but also to adequately manage the phenomena of intermittency and leverage effects. Additionally, the inclusion of sentiment indices harnesses supplementary information that transcends conventional market data, potentially conferring an enhanced analytical edge in the prognostication of volatility patterns.

We evaluate conventional GARCH, hybrid GARCH-LSTM and our proposed models on empirical data. The forecasted volatilities and corresponding Value-at-Risk estimates derived from the various models under consideration are subjected to rigorous evaluation across the entirety of the S&P 500 index constituents. The comparative analysis of outcomes is evaluated with different test statistics in concert with diverse loss functions, thereby ensuring a comprehensive and multifaceted assessment of the predictive performance of each model. The empirical analysis consists of estimating Value-at-Risk forecasts at the 2.5% significance level one day ahead in a window of 357 test trading days (almost one and a half trading years). The analysis was written and conducted in Python and R.

The chapter is organised as follows. In subsection 4.2, we present a comprehensive literature review of the importance of sentiment indices and their possible risk management applications. In subsection 4.3, we present the data, computational resources and software used. In subsection 4.4, we describe the methodology and background for VaR backtesting. In subsection 4.5, we show the linear and nonlinear relationships between sentiment indices and social media indices. In subsection 4.6, we present the results of the different predictive models, and finally, in subsection 4.7 we conclude our results and make further proposals to improve the dissemination of forecasts.

4.2 Literature Review

One of the critical aspects influencing risk management decisions is the effect of volatility clustering and its consequences. According to Cont (2007) a possible explanation for the phenomenon of volatility clustering is the multi-agent-based model family in which the composition of agents with different trading behaviours changes over time. Lux and Marchesi (1998) differentiate agents according to whether they follow a chartist or a fundamentalist strategy. Chartists predict price trends based on past patterns, while fundamentalists trade based on intrinsic asset values, ignoring short-term fluctuations. Agents probabilistically switch strategies based on realised excess profits, leading to sudden volatility spikes when chartists dominate. This is subsequently followed by market stabilisation, a dynamic behaviour recognised as another stylised fact known as on-off intermittency. This behavioural switching accounts for not only volatility clustering, heavy tails, and intermittency but also the long-term memory effect observed in volatility. Kirman and Teyssiere (2002) proposed a different multiagent model emphasising interdependency among agents, leading to herding behaviour driven by social influences or rational responses to shared information. Agents adjust their forecasts and trading activities based on current market conditions and past outcomes, contributing to persistent volatility or long memory. The specific rules governing agents' decision updates, such as changing investment thresholds, significantly impact the overall volatility dynamics. According to Banerjee and Green (2015), noise traders, similar to chartists, are irrational agents who interpret noise as informative signals, leading to overreactions or underreactions in asset prices, and contributing to increased volatility and market mispricing. Noise traders are significantly influenced by information from the news and social media such as Twitter, while fundamentalists counterbalance this influence by pushing prices towards their fundamental values. Iqbal et al. (2023) emphasise the psychological biases of noise traders, whose sentiment-driven decisions differ from rational expectations and can create market inefficiencies. All these potential explanations contribute to the dynamics of volatility, which should be considered by risk management.

Many authors have explored the impact of various information sources, including social media and news, on market returns, volume, abnormal returns, and volatility. Antweiler and Frank (2004) investigated the influence of internet stock message boards on financial markets, by analysing over 1.5 million messages from Yahoo! Finance and Raging Bull concerning companies in the Dow Jones Industrial Average and the Dow Jones Internet Index. They discover that high message activity has a small but significant negative predictive effect on stock returns, with transaction costs eroding potential gains. These messages, often posted by noise traders, correlate with increased market volatility and trading volume, suggesting a predictive relationship between message activity and

future market volatility. Tetlock (2007) also examines the relationship between media content and stock market movements, focusing on the Wall Street Journal's "Abreast of the Market" column. By quantitatively measuring media pessimism, Tetlock found that high levels of pessimism predict downward pressure on stock prices, followed by a reversion to fundamentals, indicating that media content acts as a proxy for investor sentiment. Additionally, extreme media pessimism forecasts high market trading volumes and lower market returns lead to greater media pessimism. Tetlock et al. (2008) demonstrated that negative words significantly influence financial analysis more than positive words, as they indicate problems or risks associated with companies. They introduce quantitative measures of negative tone, such as the standardised fraction of negative words, which shows a strong correlation with stock returns and predictive power for various financial attributes. These measures efficiently predict low firm earnings in the upcoming quarter and low stock returns on the following trading day, highlighting the near-immediate impact of negative sentiment on market perception and short-term stock performance. This effect persists even when investors and analysts understand the implications of negative language, as negative words enhance predictions of future earnings beyond commonly accessible quantitative data. The research also reveals a delay in market reactions to negative information, indicating a lag in processing and responding to negative news.

The focus on the fraction of negative words in financial analysis is supported by Baumeister et al. (2001), who find that language contains more words for negative emotions than for positive emotions, emphasising the greater impact of negative emotions on cognitive processes and judgements of likeability. Negative behaviours are more diagnostic of a person's character, quickly categorising someone as bad, whereas good behaviours require consistent actions to achieve a positive categorisation. Loughran and McDonald (2011) highlight the distinct nature of financial vocabulary compared to conventional language, advocating for a specialised dictionary to accurately measure sentiment in corporate communications. Their research addresses the challenges of word categorisation, noting that positive words are often negated in financial contexts, which can cause misclassifications and errors in analysis. For example, phrases such as "did not benefit" can obscure negative information. Another interesting phenomenon that distorts the real meaning of texts is that readers evaluate a company's value based on documents. Considering this phenomenon journalists and authors usually tend to use cautious language, frequently avoiding negative words, instead of qualifying positive words in ways that may not be easily identifiable by different text analysis algorithms. The presence of negative language, however, tends to have a more significant and clear-cut impact. They recommend using their Fin-Neg list, which is better attuned to financial contexts and provides more accurate predictions of corporate performance and market distress than general-purpose negative word lists. This list significantly correlates with financial metrics such as stock

returns, volatility, and trading volume, offering a more reliable tool for sentiment analysis in finance and risk management.

Certain authors have already examined the impacts caused by social media, for example, Choi and Varian (2012) explored the use of Google search query data to estimate current economic activity in various industries. They hypothesise that the volume of industry-related Google searches correlates with current performance levels and can anticipate forthcoming data releases. Using Google Trends data, which provides normalised indices of search query volumes from January 1, 2004, they integrate this information into statistical models to estimate real-time economic activity, potentially predicting future events, including volatility. Hamid and Heiden (2015) also investigated the relationship between investor attention, measured by Google search volumes, and stock market volatility. They enhanced volatility forecasting models by incorporating Google search data and found that it can improve prediction accuracy. Their empirical similarity approach predicts future events by identifying historical market behaviour patterns and using similarity-weighted averaging. This method suggests that including investor attention data can refine volatility predictions beyond traditional models, offering innovative insights into economic and financial phenomena.

Deveikyte et al. (2022) explored sentiment analysis for predicting stock market volatility and returns, focusing on the FTSE100 index. Their study confirms a correlation between sentiment in financial news and stock market movements, with negative Twitter sentiment strongly predicting next-day market volatility (with a correlation coefficient of -0.7). By utilising topic recognition and Latent Dirichlet Allocation, they developed a classifier that achieved 63% accuracy in directional volatility prediction. However, Granger causality tests indicate that sentiment from news and social media does not causally predict changes in the FTSE100 index. Kranefuss and Johnson (2021) also investigated the impact of Twitter sentiment on stock market volatility forecasts using the ARIMA and ARFIMA models. They incorporate sentiment indicators from Loughran McDonald's lexicon and tweet volumes. Their results show that negative Twitter sentiment increases volatility, while positive sentiment decreases it, supporting the leverage effect theory. The volume of Tweets significantly correlates with market volatility, highlighting the growing relevance of social media as an indicator.

Some authors have already established the exploitation of nonlinear relationships between sentiments and financial attributes. Arslan (2024) investigated the challenge of forecasting Bitcoin prices for the following day and the impact of tweets on this price by applying the Long Short-Term Memory (LSTM) model. They utilised Empirical Mode Decomposition (EMD) as a data decomposition technique to discriminate between high

and low-frequency components, enabling a comprehensive examination of the attributes inherent in each component. Additionally, they generated a real-world dataset consisting of tweets related to Bitcoin and subsequently processed it for analysis. The proposed scheme demonstrates significant potential for contributing to the rationalisation and regulation of the cryptocurrency market, offering guidance to governments and assisting individuals in making profitable investment decisions. Akbiyik et al. (2023) studied Bitcoin volatility forecasting, incorporating data from Twitter as well. They constructed deep-learning models with more than 30 million Bitcoin-related tweets and 15-minute interval price data. Their findings reveal that metadata about Tweet authors, such as influence and follower count, better predict market behaviour than Tweet content or simple volume. Temporal convolutional networks outperform classical models, suggesting that identifying influential users is key to predicting Bitcoin volatility.

However, potential sentiment indicators can be constructed using not only Google and Twitter. For instance, Siganos et al. (2014) examine the relationship between investor sentiment on Facebook and stock market returns. Critiquing conventional sentiment measures such as the Michigan Consumer Sentiment survey, they propose using Facebook's Gross National Happiness Index (FGNHI) for daily sentiment analysis across twenty international markets. This method captures daily sentiment fluctuations and provides natural insights into user sentiment, covering a broader range of countries. Their findings support the hypothesis that positive sentiment biases stock market returns, aligning with behavioural finance theories. In general, a wide range of literature has already demonstrated the strong relationship between various sentiment indices and volatility. We aim to contribute to the literature by incorporating this relationship into the currently popular models and emphasising the nonlinear relationship between variance and sentiment indices.

4.3 Data Sources, Computational Resources and Software

The data used for the analysis were downloaded exclusively from the Bloomberg Terminal. Individual stocks in the S&P 500 stock index for March 2024 are used for the analysis. Individual stocks' adjusted close price, Twitter publication count, Twitter positive sentiment count, Twitter negative sentiment count, News publication count, News positive sentiment, and News negative sentiment count values at a daily frequency from Bloomberg's sentiment analysis tool are downloaded. The data used for analysis have both advantages and disadvantages. The algorithm behind the different sentiment data calculated by Bloomberg and the dictionary used is not public information. However, as mentioned in the literature review, the importance of finding an appropriate dictionary

for accurate sentiment analysis has been emphasised in the literature. On the other hand, presumably, a wide range of investors use this feature to set alerts and notifications and most financial institutions have Bloomberg subscriptions. Furthermore, after Twitter became X, the downloading and analysis of Tweets were prohibited and currently there is no other open-source alternative for downloading and analysing tweets. Fortunately, this feature has not been removed from Bloomberg after the company takeover, and subscribers are still provided with up-to-date sentiment information moreover they can download this information at different temporal resolutions for up to five years back.

For the computations in this study, we utilise a high-performance computing environment with an Intel(R) Core(TM) i7-9700K CPU operating at 3.60 GHz, featuring 8 cores and 8 threads and an NVIDIA GeForce RTX 2070 GPU with 8 GB of dedicated RAM. The software tools employed include R Studio for different GARCH model training, as well as Python via the Anaconda distribution and Jupyter Notebook for neural network model development and model backtesting. During the training of neural networks, we leverage the TensorFlow package along with the CUDA and cuDNN toolkits to enable GPU acceleration when available, significantly enhancing the computational efficiency and reducing the training time of our models. With this configuration, the evaluation of nonlinear Granger causality for all stocks in the S&P 500 took approximately five days, whereas the training of the prediction models on the entire dataset took approximately six days. The nonlinear Granger causality analysis was conducted using the Python package implemented by Rosol et al. (2024), while the expected loss functions implemented by Bayer (2019) were also utilised for backtesting the Value-at-Risk estimates obtained from various models.

4.4 Methodology

4.4.1 Realised Volatility

The realised volatility is employed as the target variable during the training of various neural network models, as recommended by Kim and Won (2018). The metric in Formula (4.1), considering n observations provides a precise measure of the historical volatility of an asset's returns over a specified period.

$$RV = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n-1}(r_i - \bar{r})^2}$$
$$r_i = \ln\left(\frac{p_{i+1}}{p_i}\right) \tag{4.1}$$

where n is the number of observations, p_i is the stock price at the end of the i-th

interval, r_i is the log return at the end of the i-th interval, RV is the realised volatility and $\bar{r}$ is the mean of r_i.

4.4.2 Sentiment Indices

One of the pivotal elements in our research is the methodology used to measure media sentiment. Our empirical analysis indicates that employing different sentiment indices results in varying outcomes. Among the various approaches we explore to quantify the negative tone of media, we ultimately adopt the negative sentiment ratio proposed by Tetlock et al. (2008) and subsequently utilised by Kranefuss and Johnson (2021). This decision is based on the compelling rationale presented in their arguments and the significant Granger causalities exhibited by most equities when applying this metric. Nevertheless, we also take into account the positive tone of media with the measure of the positive sentiment ratios, which means that we consider including the indices found in Formulas (4.2) as feature variables. In the analysis, the impact of the negative and positive sentiment indices is examined separately because although these indicators have a negative correlation, they cannot be perfectly expressed because of the inclusion of a nonnegligible amount of neutral sentiments.

$$\begin{aligned}
\text{Negative Twitter Sentiment Index}(t) &= -\frac{\text{Number of Negative Tweets}(t)}{\text{Total Number of Tweets}(t)} \\
\text{Positive Twitter Sentiment Index}(t) &= \frac{\text{Number of Positive Tweets}(t)}{\text{Total Number of Tweets}(t)} \\
\text{Negative News Sentiment Index}(t) &= -\frac{\text{Number of Negative News}(t)}{\text{Total Number of News}(t)} \\
\text{Positive News Sentiment Index}(t) &= \frac{\text{Number of Positive News}(t)}{\text{Total Number of News}(t)}
\end{aligned} \tag{4.2}$$

where, t denotes the index of the daily observation period for which the sentiment indices and the corresponding tweet and news counts are calculated.

4.4.3 Measures of Volatility Prediction Performance

To assess and compare the volatility prediction performance of different models, three standard measures of prediction errors are employed, the mean squared error (MSE), the root mean square percentage error (RMSPE) and the mean absolute error (MAE) which are represented in Formulas (4.3). In the case of RMSPE, a (10%) error indicates that the root-mean-square relative deviation of the forecasts from the observed values is (10%), which may be interpreted as a typical relative prediction error of approximately (10%). It does not imply that every forecast differs by exactly (10%), since the squaring of the errors gives greater weight to larger deviations.

$$\begin{aligned} \text{MSE} &= \frac{1}{n}\sum_{i=1}^{n}(\hat{y}_i - y_i)^2 \\ \text{RMSPE} &= \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i - \hat{y}_i}{y_i}\right)^2} \\ \text{MAE} &= \frac{1}{n}\sum_{i=1}^{n}|\hat{y}_i - y_i| \end{aligned} \tag{4.3}$$

where n is the number of observations, y_i is the actual value at observation i, $\hat{y}_i$ is the predicted value of a given model at observation i.

4.4.4 Granger Causality

Granger causality is used to measure linear causality between returns' variance and different sentiment indices. Granger causality determines whether one time series can help predict another time series. According to Hamilton (2020) in the context of two-time series Y_t and X_t, Granger causality can be tested using the regression model (4.4). The Granger causality test involves estimating two models, the first model with only lagged values of Y_t and the other with lagged values of both Y_t and X_t. The F-statistic is subsequently computed to evaluate the joint significance of the coefficients of X_{t-i} in predicting Y_t, grounded in the theoretical framework of Formula (4.5). The null hypothesis (H_0) is that $\beta_{2i} = 0$ for all i, while the alternative hypothesis (H_1) suggests that $\beta_{2i} \neq 0$ for at least one i.

$$Y_t = \sum_{i=1}^{p}\beta_{1i}Y_{t-i} + \sum_{i=1}^{q}\beta_{2i}X_{t-i} + \varepsilon_t \tag{4.4}$$

$$F = \frac{\left(\frac{SSR_r - SSR_{ur}}{m}\right)}{\left(\frac{SSR_{ur}}{n-k}\right)} \tag{4.5}$$

where Y_t is the dependent variable, X_t is the potential causal variable, ε_t is an error term, p and q are the lag orders, SSR_r is the sum of squared residuals from the restricted model (where the lagged values of X are omitted), SSR_{ur} is the sum of squared residuals from the unrestricted model (where the lagged values of X are included), m is the number of restrictions, which is the number of lagged X terms, n is the number of observations and k is the total number of parameters estimated in the unrestricted model (including both Y and X lags and a constant term).

In our study, we investigate the unidirectional Granger causality between conditional variance and sentiment indices based on Formula (4.6). We focused exclusively on

unidirectional Granger causality because our primary interest is to determine the predictive power of different sentiment indices. In this case, the restricted model includes only lagged values of $\hat{\varepsilon}_t$, while the unrestricted model incorporates lagged values of both $\hat{\varepsilon}_t$ and various lagged values of the specific Sentiment Index$_t$.

$$\text{Var}(r_t) = \text{E}(\hat{\varepsilon}_t^2 \mid \mathcal{H}_{t-1}) = \alpha_0 + \sum_{i=1}^{p} \alpha_i \hat{\varepsilon}_{t-i}^2 + \sum_{i=1}^{n} \gamma_i \text{Sentiment Index}_{t-i} + u_t \tag{4.6}$$

where r_t is the log return at time t, $\mathcal{H}_{t-1}$ denotes the history of the process until time t, Var(r_t) is the conditional variance of r_t, α_0 is a constant term, $\hat{\varepsilon}_{t-i}^2$ is the lagged squared residuals, α_i are the coefficients of the lagged squared residuals, Sentiment Index$_{t-i}$ is the specific lagged sentiment index with their respective γ_i coefficient and u_t denotes the normally distributed residuals term with zero expected value.

Furthermore, we tested the predictive power of sentiment indices using the nonlinear Granger causality approach proposed by Rosol et al. (2022). The authors introduce a Python package that measures nonlinear causality with the help of different neural networks. In this case, the unrestricted model is estimated based on a given neural network model represented by Formula (4.7). The restricted model is estimated as the same neural network without the lagged values of the specific Sentiment Index$_t$. To evaluate the significance of causality instead of the F-test, the researchers suggest the Wilcoxon signed-rank test in Formula (4.8). The null hypothesis (H_0) of the statistical test is that the median absolute error for a model based on past values of X is equal to or smaller than that for a model based on past values of X and Y. The alternative hypothesis (H_1) is that the model based on past values of both the X and Y time series has a smaller median absolute error than does the model based only on past values of X.

$$\begin{aligned} \text{Var}(r_t) = f(&\hat{\varepsilon}_{t-1}^2, \hat{\varepsilon}_{t-2}^2, \ldots, \hat{\varepsilon}_{t-p}^2, \\ &\text{Sentiment Index}_{t-1}, \text{Sentiment Index}_{t-2}, \ldots, \text{Sentiment Index}_{t-q}; \theta) \end{aligned} \tag{4.7}$$

$$\begin{aligned} d_i &= x_{i1} - x_{i2} \\ W^+ &= \sum_{d_i>0} R(|d_i|) \\ W^- &= \sum_{d_i<0} R(|d_i|) \\ W &= \min(W^+, W^-) \end{aligned} \tag{4.8}$$

where f is a specific neural network with a θ vector of hyperparameters, Var(r_t) is the conditional variance of the r_t log return, $\hat{\varepsilon}_{t-i}^2$ is the lagged squared residual,

Sentiment Index$_{t-i}$ is the lagged sentiment index, p and q are the number of lags of $\hat{\varepsilon}_t$ and Sentiment Index$_t$ respectively, x_i are the individual observations of the two samples, and $R(|d_i|)$ is the rank of the absolute difference $|d_i|$.

4.4.5 Generalized AutoRegressive Conditional Heteroskedasticity Model (GARCH)

The first pioneering work on volatility forecasting methods was the Autoregressive Conditional Heteroscedasticity (ARCH) model introduced by Engle (1982). The variance and constant mean formulas for the p-th-order ARCH(p) model are given by Formula (4.9). This model provides a framework for estimating the variance of the error term that is conditional on past information, which is an important aspect of accommodating heteroscedastic behaviour in time series data. The parameter p represents the number of lagged squared error terms included in the model.

$$\begin{aligned} r_t &= \mu + \sigma u_t \\ \sigma^2 &= \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \alpha_2 \epsilon_{t-2}^2 + \cdots + \alpha_p \epsilon_{t-p}^2 \end{aligned} \tag{4.9}$$

where r_t is the log return at time t, μ is a constant mean, σ^2 is the conditional variance, α_0 is a constant term (and is positive), $\alpha_1, \ldots, \alpha_p$ are coefficients to be estimated (which are nonnegative) and u_t and ϵ_t are the error terms that follow a normal distribution with mean zero.

The general extension of the ARCH model was introduced by Bollerslev (1986), who extended the model to incorporate not only past squared errors (which represent past variances) but also past conditional variances. This allows for a more flexible lag structure and can better account for volatility clustering in financial time series. The representation of the conditional variance and constant mean formulas of a GARCH(p, q) model is given by Formula (4.10). The parameters p and q determine the order of the model, which represents the number of lagged variance terms p and the number of lagged squared error terms q included in the model.

$$\begin{aligned} r_t &= \mu + \sigma u_t \\ \sigma^2 &= \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2 \end{aligned} \tag{4.10}$$

where r_t is the log return at time t, μ is a constant mean, σ^2 is the conditional variance at time t, α_0 is a constant term (and is positive), ε_{t-i}^2 are the past squared errors, α_i are the coefficients for the past squared errors (which are nonnegative), σ_{t-i}^2 are past conditional variances, and β_i are the coefficients for the past conditional variances (which

are nonnegative).

4.4.6 Sentiment Augmented GARCH Model

Engle and Ng (1993) provide a comprehensive examination of how new information is incorporated into volatility estimates, using daily return data from Tokyo Stock Price Index. They introduce the news impact curve to depict the relationship between unexpected returns and resulting predicted volatility, revealing the asymmetric effects of positive and negative news on volatility levels, this phenomenon is often referred to as the leverage effect. Although models such as the Exponential GARCH (EGARCH) model capture much of this asymmetry, they may overestimate the variability of conditional variance, indicating a potential overreaction to past return shocks. Extending GARCH models to include asymmetries caused by new information is crucial.

Based on the findings of previous literature and the asymmetric nature of new information's impact on volatility, we propose an extended GARCH model incorporating sentiment indices derived from both traditional and social media sources. The extended model is represented by Formula (4.11), in which the number of lags associated with the sentiment indices is selected by the Schwarz information criterion. We fit the proposed GARCH model using a hybrid optimisation technique, which enhances robustness and accuracy by sequentially employing multiple solvers.

$$
\begin{aligned}
r_t &= \mu + \sigma u_t \\
\sigma^2 &= \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2 + \sum_{i=1}^{r_1} \gamma_i \text{Twitter_neg}_{t-i} \\
&+ \sum_{i=1}^{r_2} \delta_i \text{Twitter_pos}_{t-i} + \sum_{i=1}^{r_3} \kappa_i \text{News_neg}_{t-i} + \sum_{i=1}^{r_4} \lambda_i \text{News_pos}_{t-i}
\end{aligned}
\tag{4.11}
$$

where r_t is a log return, μ is a constant mean, σ^2 is the conditional variance, α_0 is a constant term, ε_{t-i}^2 are the past squared errors, α_i are the coefficients for the past squared errors, σ_{t-i}^2 are past conditional variances and β_i are the coefficients for the past conditional variances. The additional terms are as follows: Twitter_neg_{t-i} (negative Twitter sentiment index at time $t-i$), Twitter_pos_{t-i} (positive Twitter sentiment index at time $t-i$), News_neg_{t-i} (negative News sentiment index at time $t-i$), and News_pos_{t-i} (positive News sentiment index at time $t-i$), whose coefficients are γ_i, δ_i, κ_i, and λ_i respectively.

4.4.7 Multilayer Perceptron

To evaluate nonlinear Granger causality as described in Formula (4.7), from among several neural network models, we use the multilayer perceptron neural network, which is preferred by Rosol et al. (2022). According to Rumelhart et al. (1986) a multilayer perceptron (MLP) is a class of feedforward artificial neural networks that consists of multiple layers of nodes, which are typically interconnected in a feedforward way. Each node in the network, often called a neuron, is a computational unit that takes input from the previous layer, applies a transformation, and provides output to the next layer. The key components of an MLP include the input layer, hidden layers, and the output layer. The input x_i to the i-th neuron is a weighted sum of outputs from neurons in the previous layer, as shown in Formula (4.12). The output y_i of the i-th neuron is determined by applying a linear or nonlinear activation function f to the input x_i, as shown in Formula (4.13). Common choices for f include the sigmoid function, hyperbolic tangent function (tanh) or rectified linear unit function (ReLU). The error (E) in the output is often computed using a cost function and a few examples are shown in Formula (4.3). To update the weights, backpropagation is used which involves computing the derivative of the error cost function with respect to each w weight using the chain rule shown in the Formula (4.14) and updating the weights according to Formula (4.15), where η is a learning rate parameter. Adaptive Moment Estimation (ADAM) is an optimisation algorithm that we use during neural network training, that provides an adaptive learning rate for each parameter. It leverages the first and second moments of the gradients to dynamically adjust the learning rate, resulting in faster convergence and improved performance on large datasets or models with many parameters (Kingma and Ba, 2015). Figure 4.1 shows an example of a schematic of a multilayer perceptron architecture with 4 x feature variables, one hidden layer with 3 n neurons and an output layer with one neuron. Every neuron's input in Figure 4.1 is the w_i-weighted sum of the previous units, and its output is the value of the sum evaluated by a specific activation function.

$$x_i = \sum_j w_{ij} y_j + b_i \tag{4.12}$$

$$y_i = f(x_i) = f\left(\sum_j w_{ij} y_j + b_i\right) \tag{4.13}$$

$$\frac{\partial E}{\partial w_{ij}} = \frac{\partial E}{\partial y_i} \cdot \frac{\partial y_i}{\partial x_i} \cdot \frac{\partial x_i}{\partial w_{ij}} = \frac{\partial E}{\partial y_i} \cdot \frac{\partial y_i}{\partial x_i} \cdot y_j \tag{4.14}$$

$$w_{ij}(t+1) = w_{ij}(t) - \eta \cdot \frac{\partial E}{\partial w_{ij}(t)} \tag{4.15}$$

where w_{ij} represents the weight from the j-th neuron in the previous layer to the i-th

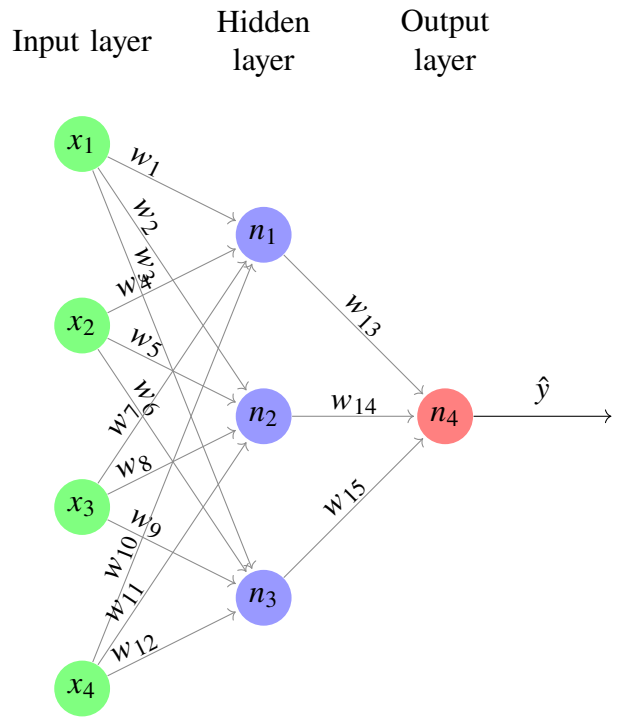

Figure 4.1: The schematic architecture of an arbitrary neural network with 4 inputs, one hidden layer with 3 neurons and an output layer with 1 neuron.

neuron in the current layer, y_j is the output from the j-th neuron in the previous layer, b_i is the bias term for the i-th neuron, f is a linear or nonlinear activation function, E is the cost function and η is a learning rate parameter.

When applying neural networks, one of the most important considerations is the issue of overfitting. Overfitting occurs when a model learns the training data too precisely, including noise and random fluctuations, mistaking them for genuine patterns. This typically leads to poor generalisation to new, unseen data which is a possible consequence of poor predictive power. One of the possible solutions is proposed by Srivastava et al. (2014), a technique known as dropout, which involves randomly removing units from the neural network during training. This approach prevents the coadaptation of features and compels the network to learn more robust and generalisable features. Another useful technique was presented by Morgan and Bourlard (1989), who proposed stopping the training process before the network has fully converged to the minimum training loss, based on the performance on a validation set. The training is halted when the network's error on the validation set starts to increase, indicating that the model is beginning to overfit the training data, rather than learning general patterns. The benefit of so-called early stopping is that it helps to maintain a good generalisation performance of the neural network on new, unseen data. By observing the network's performance on an independent test set and stopping the training at the right time, one can reduce the sensitivity to the choice of network size and the amount of training, thus helping to avoid the waste of

computational resources and potentially achieving more robust performance.

4.4.8 Long Short-Term Memory (LSTM)

Rather than combining GARCH models with a simple multilayer perceptron approach, it is advisable to use methods that account for the memory of conditional volatility. According to the literature, a more effective combination involves the Long Short-Term Memory (LSTM) recurrent network family (Hu et al., 2020, Kim and Won, 2018, Kakade et al., 2022, García-Medina and Aguayo-Moreno, 2024, Lu et al., 2016). Recurrent neural networks (RNN) are a type of neural network architecture that is particularly well-suited for processing sequences of data such as time series. The key feature of RNNs is that they have connections that loop back on themselves, allowing them to maintain information about previous inputs within their hidden state. This makes them powerful for tasks where context from earlier in the sequence is important for making predictions. LSTM is a specialised form of RNN proposed by Hochreiter and Schmidhuber (1997) and is designed to overcome the limitations of traditional RNNs, particularly the difficulty in learning dependencies between events that occur at vastly different points in time. It achieves this through the use of a complex cell structure that includes multiple gates controlling the flow of information. The formal form of one LSTM cell is represented in Formulas (4.16) and a visual representation of the formulas can be found in the schematic Figure 4.2. Each formula corresponds to one step in the operation of the LSTM cell at a given time step t.

$$
\begin{aligned}
\sigma(z) &= \frac{1}{1+\mathrm{e}^{-z}}, \qquad z \in \mathbb{R}, \\
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f), \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i), \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o), \\
\tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C), \\
C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t, \\
h_t &= o_t * \tanh(C_t).
\end{aligned}
\tag{4.16}
$$

where f_t is the forget gate activation vector, i_t is the input/update gate activation vector, o_t is the output gate activation vector, $\tilde{C}_t$ is the cell input activation vector, C_t is the cell state vector, h_t is the hidden state vector/output vector of the LSTM unit, σ denotes the sigmoid activation function, tanh is the hyperbolic tangent function, W are the weight matrices for each of the gates, b are the bias terms for each of the gates and $[h_{t-1}, x_t]$ indicates the concatenation of the previous hidden state and the current input. The $*$ operation represents the Hadamard product.

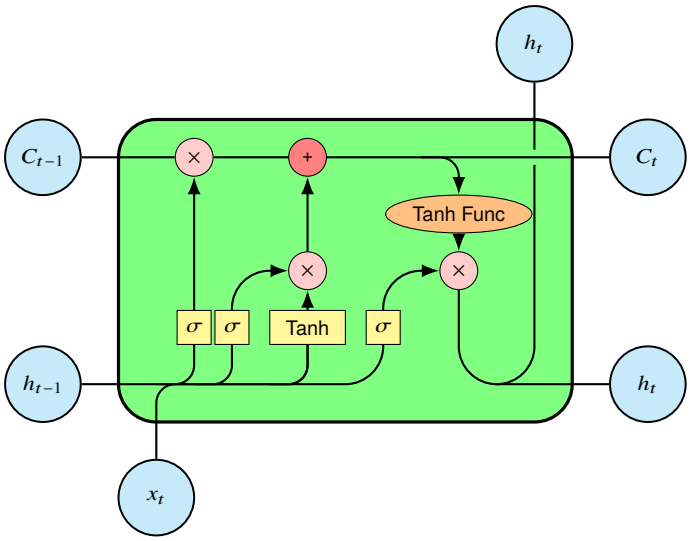

Figure 4.2: The schematic architecture of one LSTM cell.

Schuster and Paliwal (1997) advanced the LSTM cell-based neural networks, enhancing their performance on sequence prediction problems. In a standard LSTM network, information flows in a single direction through the sequence from past to future. However, in situations where the full context of the sequence is required to achieve the most accurate predictions, Bidirectional long short-term memory (BiLSTM) is especially advantageous. One BiLSTM layer consists of two LSTMs, one processes the sequence from start to end, while the other processes it from end to start. This setup allows the network to have both past and future context at each point in the sequence, potentially capturing patterns that may be missed by unidirectional LSTM models. Kakade et al. (2022) and Hu et al. (2020) demonstrated the superiority of GARCH-based BiLSTM models over traditional models, for this reason in our model, we also consider a BiLSTM layer.

An important limitation of the neural-network specifications considered in this chapter is their restricted structural interpretability and, consequently, their partly black-box character. Although LSTM and BiLSTM models can capture complex nonlinear dependencies and provide strong out-of-sample forecasts, their fitted parameters do not generally yield a transparent economic representation of how the forecast variable depends on individual inputs. Model-based sensitivities or feature-attribution measures may describe how a fitted predictor responds to its inputs, but without additional identifying assumptions they should not be interpreted as structural marginal effects or causal economic relationships. Accordingly, the methodology developed in this chapter is intended primarily for forecasting volatility and Value-at-Risk rather than for identifying or explaining the underlying economic mechanisms (Shmueli, 2010). Although explainable artificial intelligence (XAI) has emerged as a distinct research field aimed at improving the interpretability of machine-learning models, its application to financial-volatility forecasting remains an

important area for further investigation.

One promising direction in explainable artificial intelligence is represented by Kolmogorov - Arnold networks (KANs), whose learnable univariate functions can facilitate visual inspection and, in suitable cases, symbolic representation of fitted nonlinear relationships. Their application to interpretable financial-volatility forecasting has already been examined in the context of VIX prediction by Cho et al. (2025). Another, particularly novel direction is to adapt the cooperative-game-theoretic Harsányi performance-attribution framework introduced by Acerbi et al. (2026). In such an adaptation, economically meaningful input groups or modelling choices could serve as decision makers, their reference settings as benchmark actions, and an out-of-sample forecast score as the general return function. The resulting coalition attributions could decompose the model's excess predictive performance relative to the benchmark into individual contributions and interaction effects, thereby improving interpretability without constituting evidence of structural causality.

4.4.9 Value-at-Risk Estimation

Value-at-Risk (VaR) defines the potential loss in value of a portfolio over a period of time, given normal market conditions with a given confidence interval. According to the definition VaR can be expressed as Formula (4.17) (Duffie and Pan, 1997). If we assume that higher-order moments do not play a role and the distribution of returns can be considered normal, then Value-at-Risk can be estimated using the delta-normal method as shown in Formula (4.18).

$$\mathrm{VaR}_\alpha = F^{-1}(1-\alpha) \tag{4.17}$$

$$\mathrm{VaR}_\alpha = Z_\alpha \cdot \sigma \cdot \sqrt{\Delta t} \tag{4.18}$$

where VaR_α represents the Value-at-Risk at the α confidence level, F^{-1} is the inverse cumulative distribution function, Z_α is the quantile of the standard normal distribution corresponding to the α confidence level, σ is the volatility of the returns and Δt is the given time horizon.

4.4.10 Value-at-Risk Backtesting

One type of Value-at-Risk statistical test is the unconditional coverage test, we exclusively focus on the Kupiec test which is commonly used to assess the adequacy of a Value-at-Risk (VaR) model. It compares the actual number of exceedances of the VaR against the expected number under the assumed distribution. The likelihood ratio test is

used to compare the goodness-of-fit of two nested models for estimating the proportion of failure (POF) in a binary outcome scenario (0 for success, 1 for failure). The test statistic supposes that two models are available for estimating the proportion of failure, the null model M_0 and the alternative model M_1. In the null model M_0, where we assume a fixed proportion of failure, represented by π_0, the corresponding likelihood function is given by Formula (4.19). In the case of the alternative model M_1, we estimate the proportion of failure as a parameter, denoted by π_1, whose likelihood function is given by Formula (4.20). The likelihood ratio test statistic compares the two models according to Formula (4.21). Under certain regularity conditions, the likelihood ratio test statistic LR_{POF} asymptotically follows a chi-squared distribution with one degree of freedom. According to the null hypothesis (H_0) the proportion of failure is equal to π_0, while the alternative hypothesis (H_1) suggests that the proportion of failure is different from π_0.

$$\mathcal{L}(M_0) = \prod_{i=1}^{n} (\pi_0)^{y_i} (1 - \pi_0)^{1-y_i} \tag{4.19}$$

$$\mathcal{L}(M_1) = \prod_{i=1}^{n} (\pi_1)^{y_i} (1 - \pi_1)^{1-y_i} \tag{4.20}$$

$$LR_{POF} = -2 \ln \left(\frac{\mathcal{L}(M_0)}{\mathcal{L}(M_1)} \right) \tag{4.21}$$

where y_i is the binary outcome for observation i (0 for success, 1 for failure), π_0 represents the proportion of failures in the M_0 model, π_1 represents the proportion of failures in the M_1 model and n is the total number of observations.

Another type of Value-at-Risk test is the independence test, in this case, we focus on the Christoffersen and Pelletier (2004) duration test which examines the time between the i-th and $(i-1)$-th violations. The correctly specified Value-at-Risk model would imply independently and identically distributed durations. The durations under the null hypothesis should follow an exponential distribution with the probability density function expressed by Formula (4.22) and a constant hazard function. To test the independence of these durations, we can compare the observed distribution against a Weibull distribution as described by Formula (4.23). The null hypothesis (H_0) of the statistical test is that $b = 1$, therefore the durations are memoryless. The alternative hypothesis (H_1) is that $(b \neq 1)$, implies that durations are not memoryless.

$$f_{\text{exp}}(D) = pe^{-pD} \tag{4.22}$$

$$f_W(D) = ab^b D^{b-1} \exp\left(-bD^b\right) \tag{4.23}$$

where $D_i = t_i - t_{i-1}$ is the duration, in which t_i is the time of the i-th violation, p is the rate parameter of the exponential distribution and a, b are the scale and shape parameters of the Weibull distribution respectively.

According to Campbell (2005) an accurate VaR measure must satisfy both the independence and unconditional coverage properties. Tests that simultaneously evaluate these properties are capable of identifying deficiencies in Value-at-Risk measures. Kupiec's unconditional coverage test and Christoffersen and Pelletier's duration test can be combined into one statistical test expressed by Formula (4.24). The combined likelihood ratio statistic follows a chi-squared distribution with 2 degrees of freedom.

$$LR_{\text{Joint}} = LR_{\text{POF}} + LR_{\text{Duration}} \tag{4.24}$$

where LR_{POF} is the likelihood ratio statistic for Kupiec's Proportion of Failures test and LR_{duration} is the likelihood ratio statistic for Christoffersen and Pelletier's duration-based test.

Value-at-Risk is usually also tested by different loss functions. These functions inherently adopt a negative orientation, attributing higher values to instances of model failure, thus penalising inaccurate predictions more severely. The VaR model that yields the minimal loss value is deemed superior, as it demonstrates a greater efficacy in managing and mitigating risk. This evaluative process ensures that the chosen model aligns closely with the risk preferences and operational objectives of the risk managers. Lopez et al. (1999) formalised loss functions and proposed a quadratic loss function known as the magnitude loss function. Lopez's binomial loss is presented in the Formula (4.25) and the quadratic loss function can be expressed as an expected value of the $C_{m,t}$ cost function.

$$C_{m,t+1} = \begin{cases} 1 + \left(r_{t+1} - \text{VaR}_{m,t+1}\right)^2 & \text{if } r_{t+1} < \text{VaR}_{m,t+1} \\ 0 & \text{if } r_{t+1} \geq \text{VaR}_{m,t+1} \end{cases} \tag{4.25}$$

where $C_{m,t+1}$ is the cost function for model m at time $t+1$, r_{t+1} is the actual return at time $t+1$ and $\text{VaR}_{m,t+1}$ is the predicted value of Value-at-Risk for model m at time $t+1$.

According to Sarma et al. (2003), the loss function of a firm is designed to capture the economic costs associated with the failures of a Value-at-Risk (VaR) model from a firm's perspective. This includes not only direct losses when the actual portfolio returns fall below the VaR threshold but also the opportunity cost of holding excessive capital. The function emphasises the trade-off between safety and capital efficiency, penalising both model failures and excessive capital allocation. The binomial loss of the firms can be expressed by Formula (4.26) and the firm loss is an expected value of the l_t loss functions.

$$l_{t+1} = \begin{cases} (r_{t+1} - \text{VaR}_{t+1})^2 & \text{if } r_{t+1} < \text{VaR}_{t+1} \\ -\alpha \text{VaR}_{t+1} & \text{otherwise} \end{cases} \tag{4.26}$$

where l_{t+1} is the loss at time $t+1$, r_{t+1} is the actual return at time $t+1$, VaR_{t+1} is the predicted Value-at-Risk at time $t+1$ and α is the opportunity cost of capital.

According to González-Rivera et al. (2004), the smoothed Value-at-Risk (VaR) loss function provides an advanced method for evaluating VaR models. Unlike traditional nondifferentiable loss functions, it offers a smooth approximation that facilitates optimisation and statistical inference. In Formula (4.27) the smooth approximation function $m_\delta(a, b)$ makes the loss function differentiable, which is crucial for the application of gradient-based optimisation techniques.

$$\begin{aligned} \widetilde{Q} &= P^{-1} \sum_{t=R}^{T} \left(\alpha - m_\delta(r_{t+1}, \text{VaR}^\alpha_{t+1})\right) (r_{t+1} - \text{VaR}^\alpha_{t+1}) \\ m_\delta(a, b) &= [1 + \exp\{\delta(a-b)\}]^{-1} \end{aligned} \tag{4.27}$$

where α is the quantile level for the Value-at-Risk (VaR) calculation, r_{t+1} is the actual return or value at time $t+1$, VaR^α_{t+1} is the predicted Value-at-Risk at time $t+1$ for the given quantile level α, P is the prediction period, R is the estimation period, T is the total number of periods and δ is the smoothness parameter for the function $m_\delta(a, b)$, which controls the approximation of the indicator function.

4.5 Predictive Power of Sentiment Indices

In this section, we utilise both linear and nonlinear Granger causality statistical tests at different significance levels to assess the predictive power of various sentiment indices on the variance of firms listed on the S&P 500. It is essential to note that the focus here is on predictive power rather than strict causality, highlighting the capability of sentiment indices to forecast future stock variances based on historical observations. To ensure the robustness of our findings, we leverage the entire available period of historical data (5 years) for each firm in the S&P 500.

4.5.1 Linear Granger Causality

In the context of assessing the predictive power of different sentiment indices on stocks, we begin by ensuring that both the firm's variance and the related sentiment index are stationary. This is a critical step as nonstationary data can lead to spurious results in Granger causality tests. To determine stationarity, we utilise the Augmented Dickey-Fuller

(ADF) test (Hamilton, 2020). In Table 4.1 we present the results of the ADF tests, which indicate the proportion of companies in the S&P 500 whose variance and standardised sentiment indices are together stationary at the 5% significance level.

Table 4.1: The proportion of companies in the S&P 500 whose variance and standardised sentiment indices are together stationary at the 5% significance level

Sentiment Index	Propotion of Stocks
Negative Twitter Sentiment Index	91%
Positive Twitter Sentiment Index	90%
Negative News Sentiment Index	66%
Positive News Sentiment Index	92%

The linear Granger causality test is conducted only in cases where both the variance and the given standardised sentiment index indicators are stationary at the 5% significance level. The statistical tests are performed between the variance and the different standardised sentiment indices with lags ranging from 1 to 10. The proportions of companies for which the linear Granger causality between different sentiment indices and variance is significant at different levels are presented in Tables 4.2 4.3 4.4. The linear Granger causality tests reveal significantly better results in terms of the relationship between the variance of stock returns and the negative Twitter sentiment index than other sentiment indices. These results for negative tone align with those of Tetlock et al. (2008), highlighting the predictive power of negative sentiments for volatility and variance. The analysis reveals that the variance of more than one-third of the firms in the S&P 500 can be effectively predicted at the 10% significance level using the standardised negative Twitter sentiment index from the preceding period. However, following the first lag, this proportion begins to rapidly decrease and the proportion of the preceding period at 1% significance is only 10%. Nonetheless, the proportion of significant results at the 5% and 10% significance levels caused by negative Twitter sentiment remains surprisingly high, even at the 10th lag. The proportion of significant results observed with other sentiment indices is considered notable only at the 10% significance level.

Table 4.2: Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 1% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	9%	3%	2%	1%
2	10%	3%	2%	2%
3	10%	3%	2%	1%
4	10%	4%	3%	1%
5	7%	4%	2%	1%
6	7%	4%	3%	2%
7	7%	4%	4%	3%
8	8%	3%	3%	3%
9	7%	3%	4%	3%
10	6%	3%	4%	3%

Table 4.3: Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 5% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	27%	9%	6%	6%
2	21%	10%	5%	5%
3	20%	8%	6%	3%
4	18%	9%	5%	3%
5	17%	8%	6%	3%
6	16%	6%	7%	3%
7	15%	6%	7%	4%
8	16%	5%	6%	4%
9	15%	5%	5%	3%
10	14%	6%	5%	4%

Table 4.4: Proportions of observed companies for which the linear Granger causality between sentiment indices and variance are significant at the 10% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	35%	16%	11%	13%
2	30%	14%	8%	9%
3	26%	13%	8%	7%
4	24%	12%	8%	6%
5	24%	11%	8%	5%
6	21%	10%	8%	5%
7	20%	8%	7%	5%
8	20%	8%	7%	5%
9	20%	8%	8%	5%
10	21%	8%	8%	5%

4.5.2 Nonlinear Granger Causality

To determine whether there is a nonlinear relationship between the variance and different sentiment indices, we evaluate the nonlinear Granger causality approach proposed by Rosol et al. (2022). The restricted and unrestricted models of the statistical test are specified as described in Formula (4.7), where following the authors' recommendation, we utilise a multilayer perceptron to approximate the nonlinear function. The specifications of the multilayer perceptron and the training process are summarised in Table 4.5. To achieve the best results, the training is conducted using multiple learning rates. Furthermore, the training process is repeated five times with different initial weights in each case. The model that achieves the smallest Residual Sum of Squares (RSS) on the validation set, utilising the corresponding trained weights, is employed to make predictions on the test set.

Table 4.5: Multilayer Perceptron's Hyperparameters of Non-linear Granger Causality Test

Parameters	**Value**
Number of Hidden Layers	2
Number of Neurons at Each Layer	100
Activation Functions at Each Layer	Rectified Linear Unit (ReLU)
Epochs	30
Learning Rates	0.0001, 0.00001
Batch Size	32
Dropout Rate after Every Hidden Layer	5%
Validation Ratio	20%
Test Ratio	30%
Loss Function	MSE

To investigate the nonlinear relationship between variance and various sentiment indices, we consider the lags of these sentiment indices from 1 to 10. The proportions of companies for which the nonlinear Granger causality between various sentiment indices and variance is significant at the 1%, 5%, and 10% significance levels are presented in Tables 4.6, 4.7, and 4.8. The results indicate that, at all conventional significance levels (1%, 5%, and 10%), the variance of a larger proportion of the examined companies can be predicted by the given sentiment index compared to the result of linear Granger causality. The hierarchical structure of the significant causality proportions between different sentiment indices remains consistent. Specifically, at every significance level the highest proportion of significant Granger causality is observed between the negative Twitter sentiment index and the variance. This is followed by the proportions of the positive Twitter sentiment index, negative News sentiment index, and finally the positive News sentiment index. Lowering the significance level only slightly reduces the proportion of significant results and at a given significance level increasing the number of lags noticeably reduces the proportion of significant results. The first lag of the given Twitter sentiment indices offers predictive insights of variances for more than three-quarters of the observed companies. In contrast, the first lag of the News sentiment indices provides useful information for predicting the variance in less than half of the companies examined. However, this proportion remains significant, indicating that for nearly half of the companies, news sentiment indices can offer valuable predictive insights. Therefore, despite being less effective than Twitter sentiment indices, the predictive power of news sentiment indices should not be neglected. A similar phenomenon is observed when considering the last examined lag. In this case, the Twitter sentiment indices exhibit significant predictive power for more than half of the companies studied. In contrast, the corresponding predictive capability of the News sentiment indices becomes negligible. These results suggest that in a nonlinear model, indices calculated from Twitter sentiment contain significantly more information for predicting stock variance than indices derived from news sentiment.

4.6 Performance Evaluation and Backtesting of Predictive Models

In this chapter, we present the hyperparameters of our predictive models and the configuration used for training and backtesting. Additionally, we present the key observations on randomly selected stocks, followed by a comprehensive presentation of the results for all stocks in the S&P 500 index.

Table 4.6: Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 1% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	81%	79%	32%	28%
2	78%	76%	32%	24%
3	74%	72%	26%	26%
4	69%	68%	20%	18%
5	69%	64%	19%	18%
6	65%	61%	14%	15%
7	64%	59%	12%	13%
8	60%	58%	9%	9%
9	58%	55%	9%	12%
10	54%	54%	9%	8%

Table 4.7: Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 5% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	83%	81%	39%	37%
2	79%	79%	37%	31%
3	78%	75%	32%	31%
4	72%	74%	26%	26%
5	73%	68%	23%	23%
6	70%	65%	19%	20%
7	68%	64%	18%	18%
8	67%	62%	15%	15%
9	64%	58%	14%	16%
10	62%	59%	13%	11%

4.6.1 Hyperparameters of Neural Networks

We use the one-day ahead expanding rolling window technique to forecast 1-day Value-at-Risk values at the 2.5% significance level using four models. The first model is the traditional GARCH(1,1) model presented in Formula (4.10), referred to as $GARCH$. The second model is an extended GARCH model presented in Formula (4.11), referred to as $GARCH_X$. The third and fourth models are LSTM neural network models, both of which are based on GARCH(1,1) forecasts. These models differ solely in their input variables. We refer to the first neural network model as the NN model which consists solely of contemporaneous and lagged volatility forecasts from the GARCH(1,1) model. The second neural network model is referred to as the $NN_Sentiment_Index$ model, whose input is augmented with different lagged sentiment index values. The look-back period of

Table 4.8: Proportions of observed companies for which the nonlinear Granger causality between sentiment indices and variance are significant at the 10% significance level across different lags.

Lag	Negative Twitter Sentiment Index	Positive Twitter Sentiment Index	Negative News Sentiment Index	Positive News Sentiment Index
1	84%	83%	42%	41%
2	81%	79%	40%	35%
3	78%	77%	36%	33%
4	74%	76%	30%	31%
5	75%	72%	26%	27%
6	73%	67%	23%	24%
7	69%	66%	22%	23%
8	67%	64%	19%	18%
9	67%	61%	18%	20%
10	65%	63%	16%	15%

both models is one month (21 trading days). The target variable in both cases is the realised volatility computed from 10 days. First, we train the GARCH(1,1) model exclusively on the first three months (63 trading days) of the entire 5-year time series. Second, the LSTM neural networks are trained on 70% of the remaining data. The training is repeated three times with different initial weights in each case. The best model is selected based on the lowest RMSPE on the validation set, which comprises 30% of the training set. Finally, we backtest the predictive capabilities of the different models on the test dataset, which is the commonly applied 30% of the observed dataset (almost one and a half trading years, 357 trading days). According to Campbell (2005), the minimum period examined during the backtest is one year because the possibility that no violations occur during a short period such as one year is not trivial. Both of the neural networks' first hidden layers are BiLSTM layers, followed by three LSTM hidden layers. During hyperparameter tuning, we observe that the application of multiple BiLSTM layers negatively impacts the performance of both models. The selected hyperparameters of the two models are presented in Table 4.9. To avoid overfitting, we employ minibatch training, dropout between the hidden layers, and early stopping.

4.6.2 Results on Randomly Selected Stocks

The following analysis illustrates the performance differences among various models using three randomly selected stocks characterised by high volatility and intermittency. The selected stocks are Alphabet Inc. (GOOG), International Business Machines Corporation (IBM) and Tesla Inc. (TSLA). The results of the Value-at-Risk (VaR) at the 2.5% significance level forecasted by different models over the test period are visually presented in Figures 4.3 4.4 4.5. The forecasts generated by the $GARCH$ and $GARCH_X$

Table 4.9: The hyperparameters of neural networks for volatility forecasting

Parameter	Value
Hidden layers	BiLSTM, LSTM, LSTM, LSTM
Number of LSTM cells in hidden layers	128, 64, 32, 16
Activation functions of neurons	Hyperbolic Tangent
Dropout ratio between hidden layers	0.1
Learning rate	0.001
Epochs	200
Early stopping patience	30 Epochs
Batch Size	32
Loss function	RMSPE

models exhibit significant differences, attributed to the additional information considered in the $GARCH_X$ model. Compared to the $GARCH$ model, the $GARCH_X$ model is more responsive to unforeseen changes in returns, and its absolute magnitudes are occasionally greater. Furthermore, in many instances, the Value-at-Risk levels forecasted by the $GARCH_X$ model are generally lower than those forecasted by the $GARCH$ model. This endeavours to address the underestimation of Value-at-Risk observed in the GARCH model family as noted by Lee and Su (2012).

The neural network models, $NN_Sentiment_Index$ and NN exhibit heightened sensitivity to immediate changes, indicating their ability to swiftly respond to abrupt market fluctuations. The NN model, which relies exclusively on lagged conditional variances predicted by the $GARCH$ model, attempts to react more promptly than the simple $GARCH$ model. However, the NN model is unable to capture all changes, for instance, there are certain changes to which only the $GARCH_X$ model can respond promptly, due to the additional information it incorporates. The sentiment-augmented neural network $NN_Sentiment_Index$ leverages the additional information utilised by the $GARCH_X$ model to better capture and react to sudden and intense market changes. The absolute magnitudes of the VaR forecasts produced by the different neural network models are substantially greater than those of the other models' forecasts. However, the instance of the NN model shows that the flexibility provided by the neural network alone is not sufficient. The additional information provided by the different sentiment indices also increases the accuracy of the VaR forecasts which ensures the immediate response to most sudden changes with appropriate absolute magnitudes. The $NN_Sentiment_Index$ and $GARCH_X$ models establish lower VaR levels in high volatility periods. However, similar to the NN model, during less volatile periods, the $NN_Sentiment_Index$ model predicts higher VaR levels, which prevents the model from overestimating the VaR during these times, as is often the case with $GARCH$ models (Lee and Su, 2012). This suggests that the adequate integration of sentiment data in different models contributes to a more con-

servative risk assessment in high-volatility periods, which prevents the underestimation of Value-at-Risk and the flexibility provided by the neural network ensures that the VaR level is not overestimated during less volatile periods. Based on the selected visual examples, it is evident that the *NN_Sentiment_Index* model performs the best among all the models, because this model combines the indispensable complexity offered by the neural networks with the advantages of providing additional information on sentiment indices.

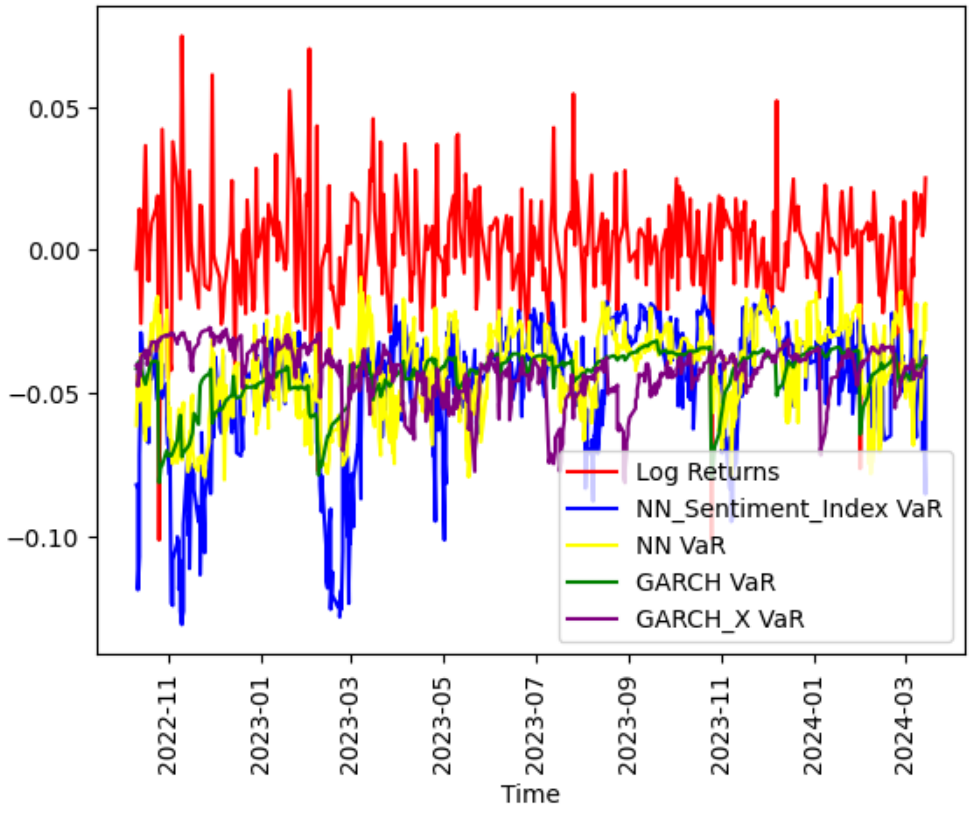

Figure 4.3: Forecasts of the 2.5% Value-at-Risk for Alphabet Inc. stock in the test dataset as predicted by multiple models.

In addition to evaluating the visual results, we also conduct various statistical tests on the different VaR forecasts. The discrepancies from realised volatility for these stocks, the statistical backtests of VaR, and the expected losses determined by various loss functions are summarised in Tables 4.10 4.11 4.12. Table 4.10 reveals that the *NN_Sentiment_Index* model consistently outperformed the other models in forecasting realised volatility for all randomly selected stocks. A pairwise comparison of the models revealed that the *NN_Sentiment_Index* model performed better in most cases across all three metrics and outperformed in at least two metrics. For the randomly selected stocks under investigation, distinguishing among the remaining three models is not straightforward. Consequently, the comparison among these models is presented in the aggregated results subsection, which encompasses a larger sample size to provide a comprehensive analysis.

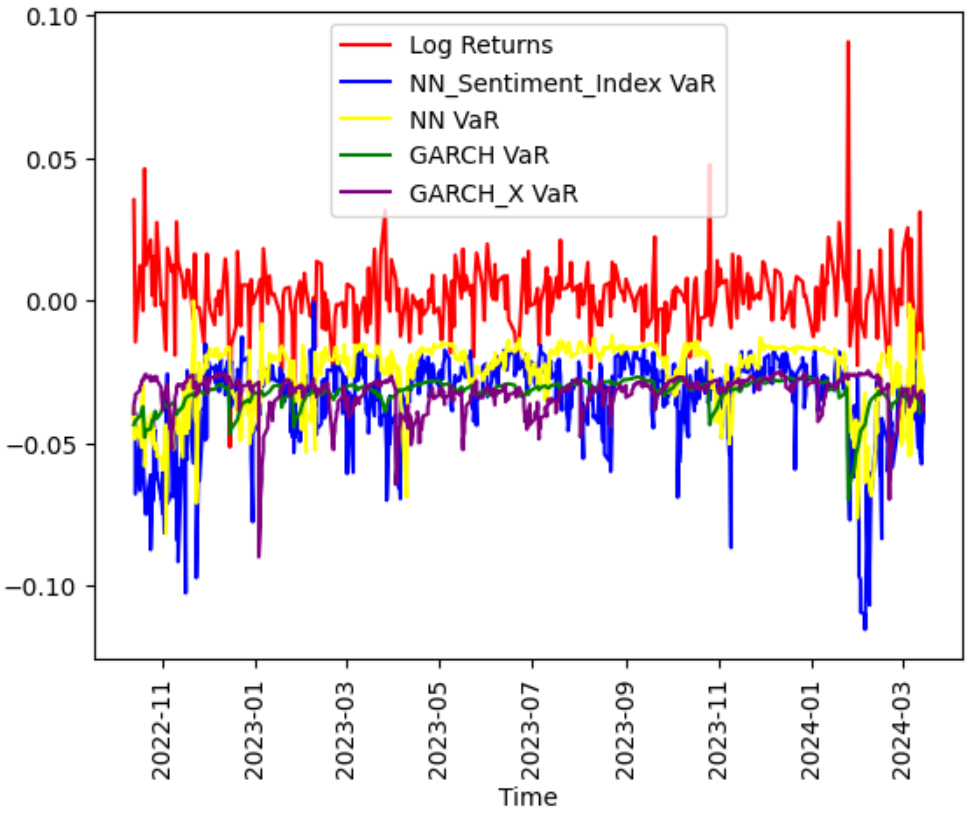

Figure 4.4: Forecasts of the 2.5% Value-at-Risk for International Business Machines Corporation stock in the test dataset as predicted by multiple models.

One of the most critical aspects of Value-at-Risk (VaR) backtesting is compliance with statistical tests. In the case of randomly selected stocks Table 4.11 shows that the GARCH model fails to meet any of the observed statistical tests at the 5% significance level, which confirms the general finding that GARCH models tend to underestimate or overestimate the actual Value-at-Risk (Lee and Su, 2012). In contrast, nearly all the other models satisfy the Kupiec test at the 5% significance level. However, only the VaR predictions generated by the $NN_Sentiment_Index$ model satisfy both the Christoffersen Duration and Joint tests. The observed results for the selected stocks are also reflected in the aggregated outcomes, demonstrating that only the $NN_Sentiment_Index$ model neither underestimates nor overestimates the Value-at-Risk, while the time intervals between violations exhibit no memory effect.

Finally, we assess the expected loss values estimated using different loss functions, as presented in Table 4.12. The table indicates that, in almost all cases, most of the different losses calculated with the VaR predictions from the $GARCH$ model are the smallest. This is likely because although the $GARCH$ model either underestimates or overestimates the VaR in certain periods, it generally tends to overestimate the VaR for the stocks examined during the whole period. The second smallest magnitude of the loss is achieved by the

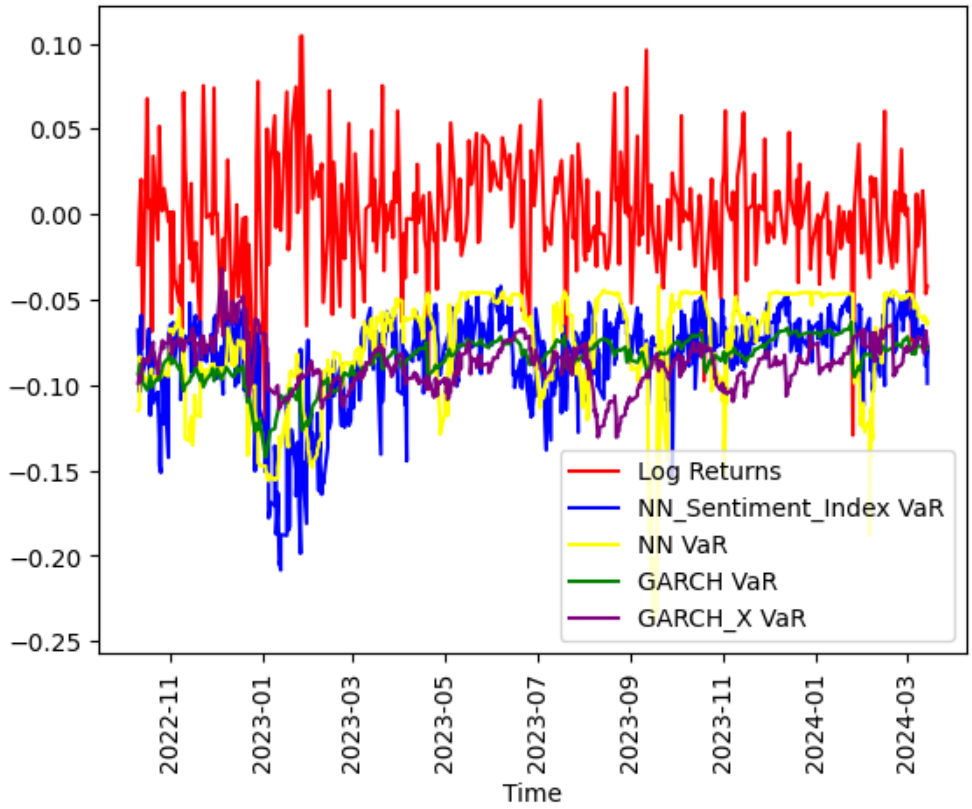

Figure 4.5: Forecasts of the 2.5% Value-at-Risk for Tesla Inc. stock in the test dataset as predicted by multiple models.

models enhanced with sentiment indices. While the traditional GARCH models appear to be the best in terms of minimising expected losses, a broader perspective reveals a different conclusion. When considering which models achieve the lowest losses while also complying with various statistical tests, the $NN_Sentiment_Index$ model stands out as the best. For the examined stocks, the $GARCH$ model does not satisfy any statistical tests, while the $NN_Sentiment_Index$ model not only adheres to the stringent requirements of statistical validation but also provides competitive loss values, making it the preferred choice for accurate and reliable Value-at-Risk forecasting.

Table 4.10: Metrics measuring the forecasting ability of volatility predicted by different models

	NN_Sentiment			NN			GARCH			GARCH_X		
	MSE	RMSPE	MAE	MSE	RMSPE	MAE	MSE	RMSPE	MAE	MSE	RMSPE	MAE
GOOG	4.504e-04	0.9768	1.966e-02	4.534e-04	0.9810	1.973e-02	4.547e-04	0.9786	1.973e-02	4.598e-04	0.9768	1.976e-02
IBM	1.663e-04	0.9674	1.174e-02	1.698e-04	0.9876	1.193e-02	1.685e-04	0.9797	1.186e-02	1.693e-04	0.9771	1.185e-02
TSLA	1.279e-03	0.9410	3.384e-02	1.297e-03	0.9602	3.431e-02	1.305e-03	0.9586	3.435e-02	1.310e-03	0.9524	3.343e-02

MSE denotes the Mean Square Error, RMSPE denotes the Root Mean Squared Percentage Error and MAE denotes the Mean Absolute Error. The different expressions of the error functions can be found in Formulas (4.3)

Table 4.11: Results of different statistical tests at the 5% significance level

	NN_Sentiment			NN			GARCH			GARCH_X		
	Kupiec	Christoffersen	Joint	Kupiec	Christoffersen	Joint	Kupiec	Christoffersen	Joint	Kupiec	Christoffersen	Joint
GOOG	FR H0	FR H0	FR H0	FR H0	R H0	R H0	R H0	R H0	R H0	FR H0	R H0	R H0
IBM	FR H0	FR H0	FR H0	FR H0	R H0	R H0	R H0	R H0	R H0	R H0	R H0	R H0
TSLA	FR H0	R H0	FR H0	FR H0	R H0	R H0	R H0	R H0	R H0	FR H0	R H0	R H0

We denote FR 0 to indicate that we failed to reject the null hypothesis of a given statistical test at the 5% significance level, while R 0 denotes that we rejected the null hypothesis of a given statistical test at the 5% significance level.

Table 4.12: Metrics measuring the expected loss of Value-at-Risk predictions predicted by different models

	NN_Sentiment			NN			GARCH			GARCH_X		
	QL	FL	SL	QL	FL	SL	QL	FL	SL	QL	FL	SL
GOOG	1.960e-02	6.142e-02	-8.006e-03	2.521e-02	6.748e-02	-7.607e-03	8.394e-03	5.197e-02	-8.629e-03	2.521e-02	6.689e-02	-7.933e-03
IBM	7.026e-03	4.443e-02	-8.681e-03	1.967e-02	4.714e-02	-7.567e-03	5.618e-03	3.774e-02	-8.599e-03	8.429e-03	4.062e-02	-8.554e-03
TSLA	1.959e-02	1.090e-01	-5.600e-03	3.359e-02	1.107e-01	-5.488e-03	1.119e-02	9.555e-02	-6.095e-03	2.800e-02	1.137e-01	-4.987e-03

QL denotes the Quadratic Loss found in Formula (4.25), FL denotes the Firm Loss found in Formula (4.26), and SL denotes the Smooth Loss found in Formula (4.27).

4.6.3 Aggregated Results of Stocks in the S&P 500 Index

The analysis conducted on the randomly selected stocks is extended to include all stocks in the S&P 500 index. The aggregated results of this comprehensive analysis are presented in this subsection. First, we compare the different models pairwise to determine the proportion of stocks where each model more accurately predicts realised volatility relative to the others based on the specified metric. Second, we assess the proportion of stocks where one model realised less expected loss compared to the other model based on the specified metric. Finally, we evaluate the proportion of stocks for which one model more or less frequently rejected the null hypothesis than did the other model in the given statistical tests.

The results are presented in Tables 4.13, 4.14 4.15, where the first two tables show the proportions of observed stocks determined by comparing how often the model in the given row outperforms the model in the corresponding column based on the specified metric. For instance in Table 4.13, the proportion of 69.57% demonstrates that the $NN_Sentiment_Index$ model has a lower MSE than does to the NN model. The $NN_Sentiment_Index$ model outperforms the NN model in terms of the MSE for approximately 69.57% of the observed stocks evaluated. Naturally, if we examine the NN row, it becomes evident that in $1 - 69.57\% = 30.43\%$ of the cases, the NN model achieves a lower MSE value compared to the $NN_Sentiment_Index$ model. In contrast, the proportions presented in Table 4.15 should be interpreted differently. These proportions indicate how often the model in the given row performs better or worse than the model in the corresponding column. Hence, these values represent relative differences in proportions. For instance, in Table 4.15 the proportion of -6.38% at the intersection of the $NN_Sentiment_Index$ (row) and NN (column) for Kupiec means that in 6.38% of the stocks, the $NN_Sentiment_Index$ model rejects the null hypothesis less frequently than does the NN model. A negative value indicates that the NN model outperforms the $NN_Sentiment_Index$ model in terms of the Kupiec statistical test. The average magnitude of the volatility prediction errors of the different models compared to that of the GARCH model is presented in Figure 4.6.

The $GARCH_X$ model exhibits much better performance in terms of error metrics compared to the $GARCH$ model in the case of the prediction of realised volatility. Figure 4.6 shows that the $GARCH_X$ model exhibits significantly smaller average errors from realised volatility compared to the $GARCH$ model based on all average metrics. Additionally, in nearly three-quarters of the examined stocks, the realised volatility is predicted more accurately by incorporating additional sentiment information compared to the $GARCH$ model. This ratio is reversed when examining expected losses based on various metrics. In this case, the $GARCH$ model predicts smaller losses in three-quarters of the

Table 4.13: The proportion of stocks for which the realised volatility is better predicted by the given model than by the other models based on the specified metric.

	NN_Sentiment_Index	NN	GARCH	GARCH_X
NN_Sentiment_Index				
MSE		69.57%	63.19%	62.34%
RMSPE		79.36%	66.17%	54.47%
MAE		76.38%	65.53%	56.81%
NN				
MSE	30.43%		49.79%	43.19%
RMSPE	20.64%		37.66%	22.34%
MAE	23.62%		44.04%	27.87%
GARCH				
MSE	36.81%	50.21%		38.09%
RMSPE	33.83%	62.34%		24.04%
MAE	34.47%	55.96%		26.81%
GARCH_X				
MSE	37.66%	56.81%	61.91%	
RMSPE	45.53%	77.66%	75.96%	
MAE	43.19%	72.13%	73.19%	

instances compared to the $GARCH_X$ model. However, according to the statistical tests, the $GARCH_X$ model demonstrates less frequent underestimation or overestimation of the predicted VaR values. The $GARCH_X$ model rejects the null hypothesis of Kupiec test less frequently than the $GARCH$ model in approximately 15% of the stocks (75 stocks) and performs better in other statistical tests as well. Overall, by incorporating additional information derived from different sentiments, we can achieve better predictive capability for realised volatility in significantly more stocks and increase compliance with statistical tests, though this comes at the cost of potentially higher expected losses for certain stocks.

Comparing the NN model to the $GARCH$ model, the GARCH(1,1) based NN model can also predict smaller deviations from realised volatility than the simple $GARCH$ model in nearly half of the examined stocks, although the average error values of the NN's predictions are greater than those of the $GARCH$ model. This is because while the NN model provides better predictions in the majority of cases, the improvement in these instances is marginal. These marginal improvements are offset by the cases where the $GARCH$ model performs better. Consequently, the average error values for the NN model are higher. One possible explanation for why the hybrid model does not perform better than the simple model is that, despite the use of multiple regularisation techniques, the neural network model may have overfitted the remaining stocks. The NN model does not show significant improvement in reducing expected losses compared to the standard $GARCH$ model, achieving lower expected losses in only approximately 10% of the examined stocks.

Table 4.14: The proportion of stocks for which the loss is smaller than the given model compared to the other model based on the specified metric.

	NN_Sentiment_Index	**NN**	**GARCH**	**GARCH_X**
NN_Sentiment_Index				
Quadratic Loss		88.30%	15.96%	34.26%
Firm Loss		60.00%	17.87%	30.21%
Smooth Loss		68.51%	23.83%	35.74%
NN				
Quadratic Loss	10.85%		7.45%	13.19%
Firm Loss	40.00%		14.04%	26.17%
Smooth Loss	31.49%		11.91%	24.68%
GARCH				
Quadratic Loss	80.00%	92.13%		69.79%
Firm Loss	82.13%	85.96%		78.30%
Smooth Loss	76.17%	88.09%		73.83%
GARCH_X				
Quadratic Loss	63.19%	86.17%	22.13%	
Firm Loss	69.79%	73.83%	21.70%	
Smooth Loss	64.26%	75.32%	26.17%	

Nevertheless, the results of the statistical tests improved significantly, for instance, the Kupiec test was satisfied more frequently in 35% of the examined cases (175 stocks), and Christoffersen's duration and Joint tests were also satisfied more frequently in one-quarter of the stocks. It is also observed that the NN model accepts the various statistical tests in significantly more instances than does the $GARCH_X$ model. Overall, the neural network-based hybrid model demonstrates better volatility forecasting capabilities in many cases and satisfies statistical tests significantly more often than do the GARCH-based models; however, this comes at the cost of higher expected losses in many instances. The NN model's potential is evident, but additional enhancements are required to fully realise its capabilities.

The $NN_Sentiment_Index$ model combines all the advantages contained in the previously presented $GARCH_X$ and NN models. While the neural network ensures the appropriate complexity of the model, the additional information derived from various sentiments provides the opportunity to identify unexpected changes. Consequently, the $NN_Sentiment_Index$ model can respond appropriately to both expected and unexpected changes not only in terms of timing but also in absolute magnitude, in contrast to the other examined models. The $NN_Sentiment_Index$ model volatility forecasting capability surpasses the other models' ability. When comparing pairwise with the other models, it achieves better results in approximately three-quarters of the examined stocks compared to the NN model and demonstrates superior predictive capability for approximately three-

Table 4.15: The proportion of stocks for which the null hypothesis of the given test statistics was rejected more or less frequently compared to the other model.

	NN_Sentiment_Index	**NN**	**GARCH**	**GARCH_X**
NN_Sentiment_Index				
Kupiec		-6.38%	28.72%	14.04%
Christoffersen		20.00%	44.04%	37.44%
Joint		11.06%	33.19%	27.87%
NN				
Kupiec	6.38%		35.11%	20.43%
Christoffersen	-20.00%		24.04%	17.45%
Joint	-11.06%		22.13%	16.81%
GARCH				
Kupiec	-28.72%	-35.11%		-14.68%
Christoffersen	-44.04%	-24.04%		-6.60%
Joint	-33.19%	-22.13%		-5.32%
GARCH_X				
Kupiec	-14.04%	-20.43%	14.68%	
Christoffersen	-37.45%	-17.45%	6.60%	
Joint	-27.87%	-16.81%	5.32%	

fifths of the examined stocks compared to the different GARCH models. Furthermore, the average error value relative to the $GARCH$ model is notably significant, being more than twice as large as the improvement achieved by the $GARCH_X$ model. In terms of expected losses, the model positions itself between the NN and $GARCH_X$ models. It significantly predicts lower expected losses in more cases, based on various metrics, than does the NN model. However, compared to the $GARCH_X$ model, it forecasts lower expected losses in only approximately one-third of the cases. Statistically, there is no evidence of a significant discrepancy in the proportion of stocks for which the NN model meets the criteria of the Kupiec test. However, substantial improvement is observed in the duration and joint statistical tests. The model satisfies the duration test in approximately an additional 20% of the examined stocks (almost 100 stocks) and the joint test in an additional 10% (more than 50 stocks) compared to the NN model. Ultimately, the $NN_Sentiment_Index$ model is the one that most accurately predicts volatility across the largest number of observed stocks among all the proposed models. Furthermore, the Value-at-Risk predicted by the $NN_Sentiment_Index$ model ensures the smallest expected loss for the majority of the observed stocks while satisfying the different statistical tests.

4.7 Conclusion and Further Suggestions

In this chapter, we propose incorporating the additional information provided by sentiments of news and social media into well-known conditional volatility forecasting

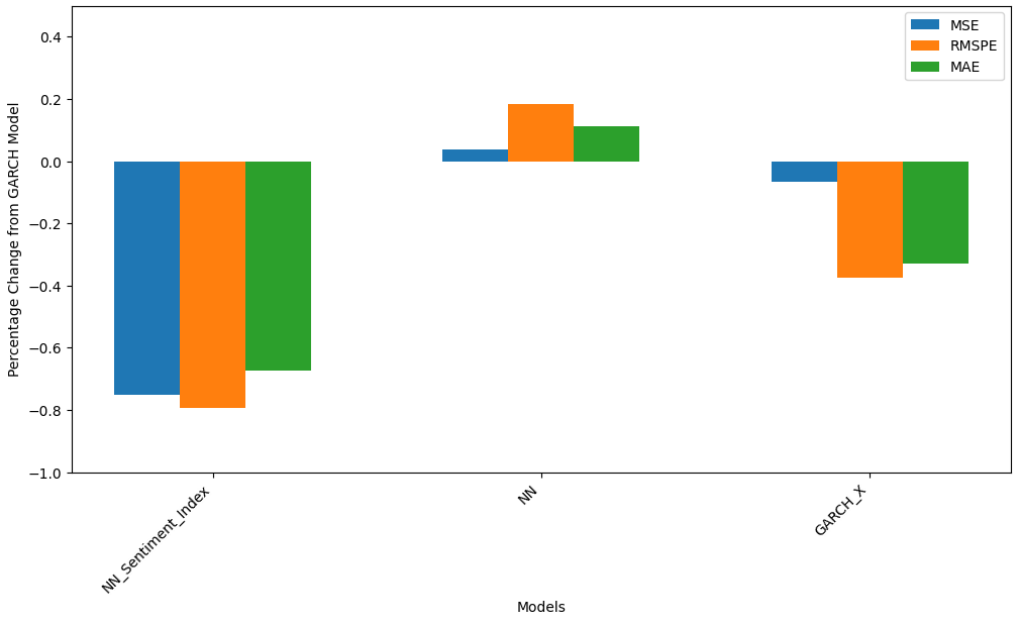

Figure 4.6: The volatility prediction capability of the given models compared to that of the GARCH model

models. The underlying idea is that sentiment indices can predict unexpected changes that cannot be anticipated based solely on information on previous returns. First, we explore the nonlinear functional relationship between various sentiment indices and conditional volatility. Following this, we propose two models enhanced with sentiment indices to achieve more accurate conditional volatility predictions, and thereby more precise Value-at-Risk estimations.

We conducted an empirical analysis to compare different models. During the analysis, we estimate Value-at-Risk (VaR) values at a 2.5% significance level for all stocks listed in the S&P 500. The models are trained and backtested over a five-year period. Based on the predictions the sentiment augmented GARCH model is more effective at predicting realised volatility across various metrics than the standard GARCH model. The additional information from sentiment indices reduces instances of underestimation or overestimation of the actual Value-at-Risk. GARCH prediction-based neural networks are capable of responding promptly to rapid returns changes. However, we demonstrate that these reactions are sometimes missing or often insufficiently large for accurate VaR estimation. Our proposed sentiment augmented neural network model combines the necessary complexity of neural networks and the ability to detect unexpected changes provided by the additional information of different sentiment indices. The model, therefore, is capable of responding not only promptly to unforeseen changes but also to changes of an appropriate magnitude.

Based on our results, the *NN_Sentiment_Index* model most accurately predicts con-

ditional volatility among all the models examined. Furthermore, for most of the examined stocks, the Value-at-Risk estimated by the model is the one that realises the smallest expected loss while satisfying the conventional statistical tests of Value-at-Risk. These features of the sentiment augmented neural network model demonstrate its superior performance and reliability in financial risk prediction. Therefore, we recommend further development and integration of this model into risk management practices to enhance predictive accuracy and minimise potential financial losses.

During the training process for individual stocks, we refrained from performing hyperparameter optimisation. We apply the same sufficiently complex architecture and hyperparameters in all cases to ensure that the estimated results from the models are easily comparable. Therefore, we believe that the results could be further improved if hyperparameter optimisation is conducted separately for each stock. Unfortunately, we were only able to conduct the training and backtesting of the models over the past five years. However, if one has access to older sentiment data, it would be also possible to train the proposed models in a stressed environment, which we believe would result in even more conservative VaR predictions. In this study, we exclusively forecast 1-day ahead Value-at-Risk, but the architecture of the neural network with LSTM cells could enable longer-term forecasts, which, in our opinion, could achieve even greater improvement compared to conventional models.

Chapter 5

Conclusion and Future Research

Overall, the dissertation examines derivative pricing and risk management from three closely connected perspectives. Chapter 2 presents a numerical simulation of a pseudo-Hermitian Hamiltonian for single-asset path-dependent options. By formulating the problem in terms of Arrow-Debreu state prices within a Hilbert space framework, the chapter implements Hamiltonian spectral simulation and a shift-invert Lánczos variant for derivative valuation and risk-sensitivity estimation. The numerical results show that the proposed algorithms produce highly accurate option values, state-price densities, digital option values, Delta estimates, and knock-out probabilities. They also show that, on finer grids, the proposed shift-invert Lánczos Hamiltonian simulation attains essentially the same accuracy as the full spectral solution with much more favourable runtime scaling. The benchmark comparisons further indicate that the Hamiltonian-based simulations remain fully competitive with standard finite-difference, tree lattice, and Monte Carlo procedures, especially for risk measures that depend sensitively on the state-price density. More broadly, the numerical Hamiltonian framework is not limited to the contracts studied in the chapter and can be extended to other path-dependent derivatives for which no closed-form analytical solution is available.

Chapter 3 generalises the single-asset framework to multiple correlated underlying assets. Its main contribution is the derivation of the pseudo-Hermitian Hamilton operator of the multi-asset Black-Scholes-Merton model. Building on this result, the chapter develops semi-analytical valuation formulas for multi-asset plain vanilla European claims and double-knockout rainbow options, including cases with asset-specific barrier levels. In this way, the chapter shows that the Hamiltonian simulation framework developed in the single-asset setting can be extended to higher-dimensional derivative-pricing problems.

Moving beyond derivative valuation, Chapter 4 advances the more precise measurement of financial risk by introducing data-driven methods for risk forecasting. It investigates whether sentiment indicators extracted from news and social media contain useful

information for volatility and Value-at-Risk forecasting. The chapter extends the standard GARCH framework with sentiment variables and develops sentiment-augmented recurrent neural network specifications that combine volatility forecasts by the GARCH model with sentiment information, particularly via LSTM architectures. The empirical results indicate that sentiment contains predictive information about future volatility, and a nonlinear model of this predictive content yields further gains in forecast accuracy. The empirical analysis is conducted on a broad scale across the full cross-section of stocks included in the S&P 500 index. Overall, the results indicate that sentiment augmentation improves realised-volatility forecasts and, in many cases, yields more reliable Value-at-Risk estimates than benchmark models, while the proposed sentiment-augmented neural-network specification provides the strongest overall performance across the examined stocks.

The findings presented above also permit a direct evaluation of the four hypotheses formulated in Chapter 1.

H1. H1 is supported within the examined numerical setting. The results of Chapter 2 show that the pseudo-Hermitian Hamiltonian framework is fully competitive with the Crank–Nicolson, binomial-tree, and Monte Carlo benchmarks, with particularly strong performance for quantities that depend sensitively on the Arrow–Debreu state-price density. On sufficiently fine grids, the shift-invert Lánczos method retains essentially the accuracy of the full-spectral solution while exhibiting more favourable empirical runtime scaling.

H2. Chapter 3 demonstrates that the pseudo-Hermitian similarity transformation can be generalised to correlated multi-asset Black–Scholes–Merton pricing problems and that it maps the corresponding non-Hermitian pricing Hamiltonian to a Hermitian representation. For European terminal payoffs, the resulting free-particle formulation recovers the conventional multi-asset Black–Scholes–Merton pricing kernel, thereby enabling the valuation of standard multi-asset European options. For the examined two-asset double knock-out worst-of rainbow option with asset-specific barriers, the particle-in-a-box construction approximates the transformed barrier problem and yields a computationally tractable semi-analytical approximation method. Taken together, these results support H2.

H3. H3 is partially supported because statistically significant predictive ability was not found for all examined stocks. At the first lag and the 5% significance level, the nonlinear tests detected significant predictive relationships for 83%, 81%, 39%, and 37% of the examined stocks for negative Twitter/X, positive Twitter/X, negative news, and positive news sentiment, respectively. The corresponding linear detection rates were 27%, 9%, 6%, and 6%. Thus, the results indicate that nonlinear

specifications capture a substantial part of the predictive information contained in sentiment that remains undetected by linear tests.

H4. H4 is also partially supported because sentiment augmentation improved forecast performance for a majority, but not all of the examined stocks. Relative to GARCH, the sentiment-augmented GARCH model produced lower errors for 61.91%, 75.96%, and 73.19% of the stocks under MSE, RMSPE, and MAE, respectively. Relative to the corresponding GARCH–LSTM-based benchmark, the sentiment-augmented model improved these measures for 69.57%, 79.36%, and 76.38% of the stocks. The Value-at-Risk results also improved overall, as sentiment augmentation generally increased the non-rejection rates across the applied backtests. For the GARCH specifications, improvements occurred under the coverage, duration, and joint tests, whereas for the GARCH–LSTM-based specifications the gains were concentrated in the duration and joint tests. The improvement was therefore substantial overall, although it was not uniform across all model specifications and backtests.

On the basis of these numerical and empirical findings, the principal results of the dissertation can be formulated as the following thesis statements:

T1. Within the examined single-asset Black–Scholes–Merton settings, the pseudo-Hermitian Hamiltonian spectral method reconstructs the Arrow–Debreu state-price density with high numerical accuracy and is fully competitive with the Crank–Nicolson, binomial-tree, and Monte Carlo benchmark methods for option valuation, with particularly strong performance for density-sensitive quantities. On sufficiently fine grids, the shift-invert Lánczos approximation reproduces the full-spectral estimates with essentially unchanged accuracy while exhibiting a more favourable empirical runtime.

T2. The pseudo-Hermitian similarity transformation applied in the single-asset framework of T1 is generalisable to correlated multi-asset Black–Scholes–Merton pricing problems and maps the corresponding non-Hermitian pricing Hamiltonian to a Hermitian representation. For European terminal payoffs, the resulting free-particle formulation recovers the conventional multi-asset Black–Scholes–Merton pricing kernel and thereby enables the valuation of standard multi-asset European options. For the examined two-asset double knock-out worst-of rainbow option with asset-specific barriers, the particle-in-a-box construction yields a computationally tractable semi-analytical approximation.

T3. Lagged positive and negative sentiment indices derived from news and Twitter/X contain statistically significant predictive information for the conditional variance

of daily returns among substantial proportions of the examined S&P 500 stocks. The considerably higher detection rates of the nonlinear tests show that nonlinear specifications capture a substantial part of the predictive information contained in sentiment that remains undetected by the standard linear test.

T4. Augmenting the GARCH and GARCH–LSTM-based models with sentiment information reduces out-of-sample realised-volatility forecast errors for a majority of the examined stocks under each of MSE, RMSPE, and MAE. These forecasting gains are accompanied by stronger overall Value-at-Risk backtesting results across the coverage, duration, and joint tests, with the sentiment-augmented specifications attaining higher non-rejection rates in five of the six model-by-test comparisons.

These results also point to several directions for future research. A natural next step is to connect Chapters 2 and 3 more closely by developing a numerical Hamiltonian simulation framework for multi-asset path-dependent derivatives, ideally with computational techniques that remain tractable in higher dimensions. A further natural direction follows from our observation that volatility can be forecast with high accuracy using data-driven approaches. This suggests that, rather than relying exclusively on parametrically specified dynamics, it may be advantageous to estimate the Hamiltonian operator used for option pricing using data-driven methods and then employ the resulting operator directly for valuation and risk sensitivity estimation. Finally, since the pricing methodology is formulated in operator form, the development of efficient quantum algorithms for these valuation problems remains a promising longer-term research direction. Progress in this area could open the way to new computational methods for pricing and risk analysis of increasingly complex derivative structures.

References

Acerbi, C., Csoka, P., Herings, P., 2026. Performance Attribution: the Harsanyi Method. WorkingPaper. CentER, Center for Economic Research. CentER Discussion Paper Nr. 2026-004.

Aït-Sahalia, Y., Lo, A.W., 1998. Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. The Journal of Finance 53, 499–547. doi:https://doi.org/10.1111/0022-1082.215228.

Akbiyik, M.E., Erkul, M., Kämpf, K., Vasiliauskaite, V., Antulov-Fantulin, N., 2023. Ask "Who", not "What": Bitcoin Volatility Forecasting with Twitter Data, in: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, pp. 688–696.

An, X., Liu, F., Zheng, M., Anh, V.V., Turner, I.W., 2021. A space-time spectral method for time-fractional Black-Scholes equation. Applied Numerical Mathematics 165, 152–166. doi:https://doi.org/10.1016/j.apnum.2021.02.009.

Antweiler, W., Frank, M.Z., 2004. Is all that talk just noise? The information content of internet stock message boards. The Journal of Finance 59, 1259–1294. doi:https://doi.org/10.1111/j.1540-6261.2004.00662.x.

Araneda, A.A., Villena, M.J., 2021. Computing the CEV option pricing formula using the semiclassical approximation of path integral. Journal of Computational and Applied Mathematics 388, 113244. doi:https://doi.org/10.1016/j.cam.2020.113244.

Arrow, K.J., 1964. The Role of Securities in the Optimal Allocation of Risk-bearing. The Review of Economic Studies 31, 91–96. doi:https://doi.org/10.1007/978-1-349-15492-0_19.

Arslan, S., 2024. Bitcoin Price Prediction Using Sentiment Analysis and Empirical Mode Decomposition. Computational Economics , 1–22 doi:https://doi.org/10.1007/s10614-024-10588-3.

Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1999. Coherent measures of risk. Mathematical Finance 9, 203–228. doi:https://doi.org/10.1111/1467-9965.00068.

Aurentz, J.L., Kalantzis, V., Saad, Y., 2017. Cucheb: A GPU implementation of the filtered Lanczos procedure. Computer Physics Communications 220, 332–340. doi:https://doi.org/10.1016/j.cpc.2017.06.016.

Baaquie, B.E., 2007. Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates. Cambridge University Press.

Baaquie, B.E., Corianó, C., Srikant, M., 2004. Hamiltonian and potentials in derivative pricing models: exact results and lattice simulations. Physica A: Statistical Mechanics and its Applications 334, 531–557. doi:https://doi.org/10.1016/j.physa.2003.10.080.

Banerjee, S., Green, B., 2015. Signal or noise? Uncertainty and learning about whether other traders are informed. Journal of Financial Economics 117, 398–423. doi:https://doi.org/10.1016/j.jfineco.2015.05.003.

Basel Committee on Banking Supervision, 2013. Fundamental review of the trading book: A revised market risk framework. 11.04.2024 URL: https://www.bis.org/publ/bcbs265.pdf.

Basel Committee on Banking Supervision, 2019. Minimum capital requirements for market risk. 11.04.2024 URL: https://www.bis.org/bcbs/publ/d352.pdf.

Baumeister, R.F., Bratslavsky, E., Finkenauer, C., Vohs, K.D., 2001. Bad is stronger than good. Review of General Psychology 5, 323–370. doi:https://doi.org/10.1037/1089-2680.5.4.323.

Bayer, S., 2019. VaR-Backtesting. URL: https://github.com/BayerSe/VaR-Backtesting.git.

Bender, C.M., Boettcher, S., 1998. Real Spectra in Non-Hermitian Hamiltonians Having $\mathcal{PT}$ Symmetry. Physical Review Letters 80, 5243–5246. doi:https://doi.org/10.1103/PhysRevLett.80.5243.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327. doi:https://doi.org/10.1016/0304-4076(86)90063-1.

Boyle, P.P., Lau, S.H., 1994. Bumping up against the barrier with the binomial method. The Journal of Derivatives 1, 6–14. doi:https://doi.org/10.3905/jod.1994.407891.

Brown, H., Hobson, D., Rogers, L., 2001. Robust Hedging of Barrier Options. Mathematical Finance 11, 285–314. doi:https://doi.org/10.1111/1467-9965.00116.

Buczynski, M., Chlebus, M., 2023. GARCHNet: Value-at-Risk Forecasting with GARCH Models Based on Neural Networks. Computational Economics , 1–31 doi:https://doi.org/10.1007/s10614-023-10390-7.

Campbell, S.D., 2005. A Review of Backtesting and Backtesting Procedures. Finance and Economics Discussion Series 2005-21. Board of Governors of the Federal Reserve System (U.S.). doi:https://doi.org/10.17016/FEDS.2005.21.

Carr, P., Itkin, A., 2021. Semi-Analytical Solutions for Barrier and American Options Written on a Time-Dependent Ornstein–Uhlenbeck Process. The Journal of Derivatives 29, jod.2021.1.133. doi:https://doi.org/10.3905/jod.2021.1.133.

Chamberlain, G., Rothschild, M., 1983. Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets. Econometrica 51, 1281–1304. doi:https://doi.org/10.2307/1912275.

Cho, S.Y., Lee, S., Kim, H.G., 2025. Forecasting VIX using interpretable Kolmogorov-Arnold networks. Expert Systems with Applications 294, 128781. doi:https://doi.org/10.1016/j.eswa.2025.128781.

Choi, H., Varian, H., 2012. Predicting the present with Google Trends. Economic record 88, 2–9. doi:https://doi.org/10.1111/j.1475-4932.2012.00809.x.

Christoffersen, P., Pelletier, D., 2004. Backtesting value-at-risk: A duration-based approach. Journal of Financial Econometrics 2, 84–108. doi:https://doi.org/10.1093/jjfinec/nbh004.

Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative finance 1, 223. doi:https://doi.org/10.1080/713665670.

Cont, R., 2007. Volatility clustering in financial markets: empirical facts and agent-based models, in: Long memory in economics. Springer, pp. 289–309. doi:https://doi.org/10.1007/978-3-540-34625-8_10.

Contreras, M., Pellicer, R., Villena, M., Ruiz, A., 2010. A quantum model of option pricing: When Black–Scholes meets Schrödinger and its semi-classical limit. Physica A: Statistical Mechanics and its Applications 389, 5447–5459. doi:https://doi.org/10.1016/j.physa.2010.08.018.

Contreras G., M., Llanquihuén, A., Villena, M., 2016. On the Solution of the Multi-Asset Black-Scholes Model: Correlations, Eigenvalues and Geometry. Journal of Mathematical Finance 06, 562–579. doi:https://doi.org/10.4236/jmf.2016.64043.

Cortes, C.L., Gray, S.K., 2022. Quantum Krylov subspace algorithms for ground- and excited-state energy estimation. Phys. Rev. A 105, 022417. doi:https://doi.org/10.1103/PhysRevA.105.022417.

Daníelsson, J., rn N. Jorgensen, B., Samorodnitsky, G., Sarma, M., de Vries, C.G., 2013. Fat tails, VaR and subadditivity. Journal of Econometrics 172, 283–291. doi:https://doi.org/10.1016/j.jeconom.2012.08.011.

Davydov, D., Linetsky, V., 2003. Pricing Options on Scalar Diffusions: An Eigenfunction Expansion Approach. Operations Research 51, 185–209. doi:https://doi.org/10.1287/opre.51.2.185.12782.

Degiannakis, S., Floros, C., Livada, A., 2012. Evaluating value-at-risk models before and after the financial crisis of 2008: International evidence. Managerial Finance 38, 436–452. doi:https://doi.org/10.1108/03074351211207563.

Deveikyte, J., Geman, H., Piccari, C., Provetti, A., 2022. A sentiment analysis approach to the prediction of market volatility. Frontiers in Artificial Intelligence 5, 836809. doi:https://doi.org/10.3389/frai.2022.836809.

Donaldson, R., Kamstra, M., 1997. An artificial neural network-GARCH model for international stock return volatility. Journal of Empirical Finance 4, 17–46. doi:https://doi.org/10.1016/S0927-5398(96)00011-4.

Duffie, D., Pan, J., 1997. An overview of value at risk. The Journal of Derivatives 4, 7–49. doi:http://dx.doi.org/10.3905/jod.1997.407971.

Engle, R.F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica , 987–1007 doi:https://doi.org/10.2307/1912773.

Engle, R.F., Ng, V.K., 1993. Measuring and testing the impact of news on volatility. The Journal of Finance 48, 1749–1778. doi:https://doi.org/10.1111/j.1540-6261.1993.tb05127.x.

Epperly, E.N., Lin, L., Nakatsukasa, Y., 2022. A Theory of Quantum Subspace Diagonalization. SIAM Journal on Matrix Analysis and Applications 43, 1263–1290. doi:https://doi.org/10.1137/21M145954X.

Fernández, F.M., 2020. Algebraic treatment of non-Hermitian quadratic Hamiltonians. Journal of Mathematical Chemistry 58, 2094–2107. doi:https://doi.org/10.1007/s10910-020-01165-8.

Fisher, T., 2007. Variations on the Vanna-Volga Adjustment. Bloomberg Research Paper .

García-Medina, A., Aguayo-Moreno, E., 2024. LSTM–GARCH hybrid model for the prediction of volatility in cryptocurrency portfolios. Computational Economics 63, 1511–1542. doi:https://doi.org/10.1007/s10614-023-10373-8.

Gibbs, A.L., Su, F.E., 2002. On Choosing and Bounding Probability Metrics. International Statistical Review 70, 419–435. doi:https://doi.org/10.2307/1403865.

Gonzalez-Conde, J., Rodríguez-Rozas, Á., Solano, E., Sanz, M., 2023. Efficient Hamiltonian simulation for solving option price dynamics. Physical Review Research 5, 043220. doi:https://doi.org/10.1103/PhysRevResearch.5.043220.

González-Rivera, G., Lee, T.H., Mishra, S., 2004. Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood. International Journal of Forecasting 20, 629–645. doi:https://doi.org/10.1016/j.ijforecast.2003.10.003.

Griffiths, D.J., Schroeter, D.F., 2018. Introduction to Quantum Mechanics. Cambridge University Press.

Hamid, A., Heiden, M., 2015. Forecasting volatility with empirical similarity and Google Trends. Journal of Economic Behavior & Organization 117, 62–81. doi:https://doi.org/10.1016/j.jebo.2015.06.005.

Hamilton, J.D., 2020. Time series analysis. Princeton University Press.

Hansen, L.P., Jagannathan, R., 1991. Implications of Security Market Data for Models of Dynamic Economies. Journal of Political Economy 99, 225–262. doi:http://dx.doi.org/10.1086/261749.

Hansen, L.P., Jagannathan, R., 1997. Assessing Specification Errors in Stochastic Discount Factor Models. The Journal of Finance 52, 557–590. doi:https://doi.org/10.2307/2329490.

Hansen, L.P., Richard, S.F., 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. Econometrica 55, 587–613. doi:https://doi.org/10.2307/1913601.

Harrison, J., Kreps, D.M., 1979. Martingales and arbitrage in multiperiod securities markets. Journal of Economic Theory 20, 381–408. doi:https://doi.org/10.1016/0022-0531(79)90043-7.

Haven, E.E., 2002. A discussion on embedding the Black–Scholes option pricing model in a quantum physics setting. Physica A: Statistical Mechanics and its Applications 304, 507–524. doi:https://doi.org/10.1016/S0378-4371(01)00568-4.

Herman, D., Googin, C., Liu, X., Sun, Y., Galda, A., Safro, I., Pistoia, M., Alexeev, Y., 2023. Quantum computing for finance. Nature Reviews Physics 5, 450–465. doi:https://doi.org/10.1038/s42254-023-00603-1.

Hicks, W., 2020. Closed Quantum Black-Scholes: Quantum Drift and the Heisenberg Equation of Motion. Journal of Stochastic Analysis 1. doi:http://dx.doi.org/10.31390/josa.1.1.06.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Computation 9, 1735–1780. doi:https://doi.org/10.1162/neco.1997.9.8.1735.

Hu, Y., Ni, J., Wen, L., 2020. A hybrid deep learning approach by integrating LSTM-ANN networks with GARCH model for copper price volatility prediction. Physica A: Statistical Mechanics and its Applications 557, 124907. doi:https://doi.org/10.1016/j.physa.2020.124907.

Hull, J.C., 2023. Risk management and financial institutions. John Wiley & Sons.

Hull, J.C., Basu, S., 2016. Options, Futures, and Other Derivatives. Pearson Education India.

Iqbal, N., Gul, F., Mubarik, F., 2023. Investor Sentiments and Stock Returns: A Study on Noise Traders. Journal of Positive School Psychology 7, 53–64. doi:https://doi.org/10.1016/S1874-8651(10)60010-5.

Jana, T., Roy, P., 2011. Supersymmetry in option pricing. Physica A: Statistical Mechanics and its Applications 390, 2350–2355. doi:https://doi.org/10.1016/j.physa.2011.02.027.

Jana, T., Roy, P., 2012. Pseudo Hermitian formulation of the quantum Black–Scholes Hamiltonian. Physica A: Statistical Mechanics and its Applications 391, 2636–2640. doi:https://doi.org/10.1016/j.physa.2011.12.012.

Johnson, H., 1987. Options on the Maximum or the Minimum of Several Assets. The Journal of Financial and Quantitative Analysis 22, 277–283. doi:https://doi.org/10.2307/2330963.

Kakade, K., Jain, I., Mishra, A.K., 2022. Value-at-Risk forecasting: A hybrid ensemble learning GARCH-LSTM based approach. Resources Policy 78, 102903. doi:https://doi.org/10.1016/j.resourpol.2022.102903.

Khodayari, L., Ranjbar, M., 2019. A computationally efficient numerical approach for multi-asset option pricing. International Journal of Computer Mathematics 96, 1158–1168. doi:https://doi.org/10.1080/00207160.2018.1458096.

Kim, H.Y., Won, C.H., 2018. Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. Expert Systems with Applications 103, 25–37. doi:https://doi.org/10.1016/j.eswa.2018.03.002.

Kingma, D., Ba, J., 2015. Adam: A Method for Stochastic Optimization, in: International Conference on Learning Representations (ICLR), San Diega, CA, USA.

Kirby, W., Motta, M., Mezzacapo, A., 2023. Exact and efficient Lanczos method on a quantum computer. Quantum 7, 1018. doi:https://doi.org/10.22331/q-2023-05-23-1018.

Kirman, A., Teyssiere, G., 2002. Microeconomic models for long memory in the volatility of financial time series. Studies in Nonlinear Dynamics & Econometrics 5. doi:https://doi.org/10.2202/1558-3708.1083.

Kranefuss, E., Johnson, D.K., 2021. Does Twitter Strengthen Volatility Forecasts? Evidence from the S&P 500, DJIA and Twitter Sentiment Analysis. Evidence from the S&P 500. doi:https://dx.doi.org/10.2139/ssrn.3786251.

Kumar, S., Wilmott, C.M., 2025. Simulating the non-Hermitian dynamics of financial option pricing with quantum computers. Scientific Reports 15, 13268. doi:https://doi.org/10.1038/s41598-025-97245-3.

Lánczos, C., 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. J. Res. Natl. Bur. Stand. B 45, 255–282. doi:https://doi.org/10.6028/jres.045.026.

Léber, D., 2021. Útvonalfüggő opciók kvantummechanikai árazása: Az időfüggetlen Schrödinger egyenlet és a Feynman-féle pályaintegrál segítségével. Budapesti Corvinus Egyetem, Gazdálkodástudományi Kar. URL: https://szd.lib.uni-corvinus.hu/13333/.

Léber, D., 2025. Quantum mechanical approach to pricing multi-asset path-dependent options. Physica A: Statistical Mechanics and its Applications 675, 130644. doi:https://doi.org/10.1016/j.physa.2025.130644.

Léber, D., Egyed, B., 2025. The sentiment augmented GARCH-LSTM hybrid model for value-at-risk forecasting. Computational Economics , 1–41 doi:https://doi.org/10.1007/s10614-025-11042-8.

Lee, C.F., Su, J.B., 2012. Alternative statistical distributions for estimating value-at-risk: theory and evidence. Review of Quantitative Finance and Accounting 39, 309–331. doi:https://doi.org/10.1007/s11156-011-0256-x.

Lehoucq, R., Sorensen, D., Yang, C., 1998. ARPACK Users' Guide: Solution of Large-scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods. Software, Environments, and Tools, Society for Industrial and Applied Mathematics.

Lin, J., Palmer, K., 2013. Convergence Of Barrier Option Prices In The Binomial Model. Mathematical Finance 23, 318–338. doi:https://doi.org/10.1111/j.1467-9965.2011.00501.x.

Lin, J., Ritchken, P., 2006. On pricing derivatives in the presence of auxiliary state variables. Journal of Derivatives 14, 29. doi:https://doi.org/10.3905/jod.2006.667549.

Linetsky, V., 1997. The Path Integral Approach to Financial Modeling and Options Pricing. Computational Economics 11, 129–163. doi:https://doi.org/10.1023/A:1008658226761.

Linetsky, V., 2004a. Lookback options and diffusion hitting times: A spectral expansion approach. Finance and Stochastics 8, 373–398. doi:https://doi.org/10.1007/s00780-003-0120-5.

Linetsky, V., 2004b. Spectral Expansions for Asian (Average Price) Options. Operations Research 52, 856–867. doi:https://doi.org/10.1287/opre.1040.0113.

Linetsky, V., 2007. Spectral Methods in Derivatives Pricing, in: Birge, J.R., Linetsky, V. (Eds.), Financial Engineering. Elsevier. volume 15 of *Handbooks in Operations Research and Management Science*, pp. 223–299. doi:https://doi.org/10.1016/S0927-0507(07)15006-4.

Lopez, J.A., et al., 1999. Methods for evaluating value-at-risk estimates. Economic Review 2, 3–17. doi:https://doi.org/10.21314/JOR.1999.012.

Loughran, T., McDonald, B., 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. The Journal of Finance 66, 35–65. doi:https://doi.org/10.1111/j.1540-6261.2010.01625.x.

Lu, X., Que, D., Cao, G., 2016. Volatility forecast based on the hybrid artificial neural network and GARCH-type models. Procedia Computer Science 91, 1044–1049. doi:https://doi.org/10.1016/j.procs.2016.07.145.

Lux, T., Marchesi, M., 1998. Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents. IFAC Proceedings Volumes 31, 7–10. doi:https://doi.org/10.1016/S1474-6670(17)40450-2.

Mollahasani, N., 2023. A Hybrid Spectral-Finite Difference Method for Numerical Pricing of Time-Fractional Black–Scholes Equation. Computational Economics 64, 841–869. doi:https://doi.org/10.1007/s10614-023-10441-z.

Morgan, N., Bourlard, H., 1989. Generalization and parameter estimation in feedforward nets: Some experiments. Advances in Neural Information Processing Systems 2. doi:https://dl.acm.org/doi/10.5555/2969830.2969907.

Mostafazadeh, A., 2002. Pseudo-Hermiticity versus PT symmetry: The necessary condition for the reality of the spectrum of a non-Hermitian Hamiltonian. Journal of Mathematical Physics 43, 205–214. doi:https://doi.org/10.1063/1.1418246.

Peherstorfer, B., Gómez, P., Bungartz, H.J., 2015. Reduced models for sparse grid discretizations of the multi-asset Black-Scholes equation. Advances in Computational Mathematics 41, 1365–1389. doi:https://doi.org/10.1007/s10444-015-9421-4.

Rasulov, A., Rakhmatov, R., Nafasov, A., 2016. Evaluation the Price of Multi-Asset Rainbow Options Using Monte Carlo Method. Journal of Applied Mathematics and Physics 4, 178–182. doi:https://doi.org/10.4236/jamp.2016.41021.

Riley, K.F., Hobson, M.P., Bence, S.J., 2006. Mathematical Methods for Physics and Engineering: A Comprehensive Guide. 3 ed., Cambridge University Press.

Ritchken, P., Sankarasubramanian, L., Vijh, A.M., 1993. The Valuation of Path Dependent Contracts on the Average. Management Science 39, 1202–1213. doi:https://doi.org/10.1287/mnsc.39.10.1202.

Rogers, L.C.G., Shi, Z., 1995. The Value of an Asian Option. Journal of Applied Probability 32, 1077–1088. doi:https://doi.org/10.2307/3215221.

Rosol, M., Mlyńczak, M., Cybulski, G., 2022. Granger causality test with nonlinear neural-network-based methods: Python package and simulation study. Computer Methods and Programs in Biomedicine 216, 106669. doi:https://doi.org/10.1016/j.cmpb.2022.106669.

Rosol, M., Mlyńczak, M., Cybulski, G., 2024. Nonlincausality. https://github.com/mrosol/Nonlincausality.git.

Rubinstein, M., 1991. Somewhere over the rainbow. Risk 4, 63–66.

Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning representations by back-propagating errors. Nature 323, 533–536. doi:https://doi.org/10.1038/323533a0.

Saad, Y., 2011. Numerical Methods for Large Eigenvalue Problems: Revised Edition. Classics in Applied Mathematics, Society for Industrial and Applied Mathematics.

Sarma, M., Thomas, S., Shah, A., 2003. Selection of Value-at-Risk models. Journal of Forecasting 22, 337–358. doi:https://doi.org/10.1002/for.868.

Schuster, M., Paliwal, K.K., 1997. Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing 45, 2673–2681. doi:https://doi.org/10.1109/78.650093.

Shmueli, G., 2010. To Explain or to Predict? Statistical Science 25. doi:https://doi.org/10.1214/10-STS330.

Siganos, A., Vagenas-Nanos, E., Verwijmeren, P., 2014. Facebook's daily sentiment and international stock markets. Journal of Economic Behavior & Organization 107, 730–743. doi:https://doi.org/10.1016/j.jebo.2014.06.004.

Sorensen, D.C., 1997. Implicitly Restarted Arnoldi/Lanczos Methods for Large Scale Eigenvalue Calculations, in: Keyes, D.E., Sameh, A., Venkatakrishnan, V. (Eds.), Parallel Numerical Algorithms. Springer Netherlands, Dordrecht, pp. 119–165. doi:https://doi.org/10.1007/978-94-011-5412-3_5.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. The Journal of Machine Learning Research 15, 1929–1958. doi:https://dl.acm.org/doi/10.5555/2627435.2670313.

Stamatopoulos, N., Egger, D.J., Sun, Y., Zoufal, C., Iten, R., Shen, N., Woerner, S., 2020. Option Pricing using Quantum Computers. Quantum 4, 291. doi:https://doi.org/10.22331/q-2020-07-06-291.

Stulz, R., 1982. Options on the minimum or the maximum of two risky assets: Analysis and applications. Journal of Financial Economics 10, 161–185. doi:https://doi.org/10.1016/0304-405X(82)90011-3.

Tetlock, P.C., 2007. Giving content to investor sentiment: The role of media in the stock market. The Journal of Finance 62, 1139–1168. doi:https://doi.org/10.1111/j.1540-6261.2007.01232.x.

Tetlock, P.C., Saar-Tsechansky, M., Macskassy, S., 2008. More than words: Quantifying language to measure firms' fundamentals. The Journal of Finance 63, 1437–1467. doi:https://doi.org/10.1111/j.1540-6261.2008.01362.x.

Vukovic, O., 2015. On the Interconnectedness of Schrodinger and Black-Scholes Equation. Journal of Applied Mathematics and Physics 3, 1108–1113. doi:http://dx.doi.org/10.4236/jamp.2015.39137.

Wade, B., Khaliq, A., Yousuf, M., Vigo-Aguiar, J., Deininger, R., 2007. On smoothing of the Crank–Nicolson scheme and higher order schemes for pricing barrier options. Journal of Computational and Applied Mathematics 204, 144–158. doi:https://doi.org/10.1016/j.cam.2006.04.034.

Wang, Y., Andreeva, G., Martin-Barragan, B., 2023. Machine learning approaches to forecasting cryptocurrency volatility: Considering internal and external determinants. International Review of Financial Analysis 90, 102914. doi:https://doi.org/10.1016/j.irfa.2023.102914.

Wilmott, P., 2013. Paul Wilmott Introduces Quantitative Finance. The Wiley Finance Series, Wiley.

Woo, J., Youn Kim, W., Choi, S., 2024. Efficient Shift-and-Invert Preconditioning for Multi-GPU Accelerated Density Functional Calculations. Journal of Chemical Theory and Computation 20, 7443–7452. doi:https://doi.org/10.1021/acs.jctc.4c00721.

Wróblewski, M., 2014. Quantum physics methods in share option valuation. Technical Transactions , 23–40 doi:https://doi.org/10.4467/2353737XCT.14.021.2604.

Yeşiltaş, Ö., 2023. The Black–Scholes equation in finance: Quantum mechanical approaches. Physica A: Statistical Mechanics and its Applications 623, 128909. doi:https://doi.org/10.1016/j.physa.2023.128909.

Zvan, R., Vetzal, K., Forsyth, P., 2000. PDE methods for pricing barrier options. Journal of Economic Dynamics and Control 24, 1563–1590. doi:https://doi.org/10.1016/S0165-1889(00)00002-6.

Publications Underlying the Dissertation

Léber, D., 2025. Quantum mechanical approach to pricing multi-asset path-dependent options. Physica A: Statistical Mechanics and its Applications 675, 130644. doi:https://doi.org/10.1016/j.physa.2025.130644.

Léber, D., Egyed, B., 2025. The sentiment augmented GARCH-LSTM hybrid model for value-at-risk forecasting. Computational Economics , 1–41 doi:https://doi.org/10.1007/s10614-025-11042-8.