



Politecnico
di Torino



Politecnico di Torino

Master of Science in Quantum Engineering

A Quantum-Enhanced Regime-Switching Model for Financial Time Series Forecasting

Academic Supervisor:

Prof. Giovanna Turvani

Candidate:

Giorgia Mazzaro

Company Supervisors:

Dr. Davide Caputo

Dr. Davide Tezza

Academic Year 2024–2025

Abstract

The ability to model and forecast financial time series is essential for modern economic strategy, guiding critical decisions in risk management, asset allocation, and derivative pricing. While traditional econometric models, such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH), perform well in stable conditions, they often fail to capture the abrupt regime shifts that characterise real-world markets. Detecting these structural breaks as they emerge, rather than *a posteriori*, remains a major challenge for effective forecasting.

Developed in collaboration with Data Reply, this thesis investigates how quantum mechanics can offer both a novel mathematical framework and a computational advantage to address this challenge. The financial market, with its vast number of interacting agents, inherent uncertainty, and rapid shifts in collective sentiment, provides a fertile domain for applying quantum-inspired methods. Where classical models struggle with structural non-linearities and high-dimensional data, quantum formalisms offer a new lens to model this complexity.

Specifically, this research investigates two complementary quantum approaches. First, it explores the quantum probability formalism, wherein a Schrödinger-like trading equation gives rise to discrete energy levels in market dynamics, providing a theoretical basis for the multimodal distributions observed in asset returns. This approach offers a richer interpretation of market evolution by modelling phenomena such as the superposition of investor beliefs and interference effects. Second, this work provides experimental validation for the correspondence between these quantum-like states and financial market regimes by using a quantum two-sample test to detect structural breaks. By comparing return distributions from rolling time windows, this test identifies market transitions with greater sensitivity than classical methods, and its output is used to select between distinct econometric models (ARIMA-GARCH), which have been pre-calibrated on the uni-modal and multi-modal regimes, enhancing risk management.

Keywords: Financial Time Series Forecasting, Quantum Machine Learning, Quantum Finance, Regime Switching, Quantum Two-Sample Test, Quantum ARIMA-GARCH

Contents

1	Introduction	4
1.1	Research Objectives	5
1.2	Thesis Structure	5
2	Financial Background and Classical Modelling	7
2.1	The Nature of Financial Time Series	8
2.2	Stylized Facts: The Anomalies	8
2.2.1	Volatility Clustering	9
2.2.2	Heavy Tails and Leptokurtosis	9
2.3	Classical Forecasting Models	10
2.3.1	Autocorrelation	10
2.3.2	ARIMA: Modelling the Linear Trend	10
2.3.3	GARCH: Modelling the Volatility	11
2.4	The Failure Point: Regime Switching	12
3	Quantum Fundamentals	13
3.1	The Postulates of Quantum Mechanics	14
3.2	The Qubit and its Geometric Representation	22
3.3	Quantum Gates and Circuits	25
3.3.1	Single-Qubit Gates	26
3.3.2	Multi-Qubit Gates and Universality	28
3.4	Entanglement: A Uniquely Quantum Resource	29
3.5	Elements of Quantum Machine Learning	31
3.5.1	The Hybrid Paradigm	32
3.5.2	Data Encoding Strategies	32
3.5.3	Parameterised Quantum Circuits	33
3.5.4	The Quantum Kernel Trick	33
3.5.5	State of the Art: Applications in Finance	34
4	Quantum Probability Formulation	36
4.1	Quantum Re-interpretation of Classical Probability	37

4.2	Alternative Quantum Approaches: The Harmonic Oscillator	38
4.3	The Asset Return Model: Active Trading Intention	39
4.4	The Schrödinger-like Trading Equation	41
4.4.1	Fourier Decomposition: The ω -market	41
4.4.2	The Governing Equation and the Potential	41
5	Methodology: The Experimental Pipeline	44
5.1	Phase 1: Regime Identification	45
5.1.1	Data Pre-processing	45
5.1.2	The Statistical Engine: Hall-York Test	45
5.1.3	The Search for the Threshold E_0	46
5.1.4	Regime Labelling and Persistence	47
5.2	Phase 2: Quantum Statistical Validation	47
5.2.1	Synthetic Validation Model	48
5.2.2	Feature Engineering: Capturing Temporal Dynamics	49
5.2.3	The Classical Two-Sample Test (MMD)	50
5.2.4	The Quantum Two-Sample Test	51
5.2.5	Parametric Quantum Circuit Architectures	52
5.2.6	Circuit Selection: Synthetic Benchmarking	52
5.2.7	Data Sampling and Validation Strategy	53
5.2.8	Bootstrap Resampling	53
5.3	Phase 3: Hybrid Forecasting Strategy	54
5.3.1	Model Specification	54
5.3.2	Validation Protocol: Walk-Forward vs Cross-Validation	55
5.3.3	Validation Metric	55
5.3.4	Generalisation and Model Robustness	55
5.4	Implementation and Scalability Notes	56
5.4.1	Simulation Strategy	56
5.4.2	Scalability Considerations	56
6	Results and Discussion	57
6.1	Phase 1: Regime Identification (SZSE 2011-2020)	57
6.1.1	Macro-Analysis and Historical Validation	58
6.1.2	Regime Dynamics and Persistence	59
6.2	Phase 2: Quantum Statistical Validation	60
6.2.1	Baseline Stability Analysis	60
6.2.2	Robustness to Time and Sample Size	61
6.2.3	Regime Discrimination Test (Detection)	62
6.3	Phase 3: Forecasting Application	63

6.3.1	Methodology: The Merged Crisis Model	64
6.3.2	Results: Informational Efficiency	64
6.3.3	Concept Drift and Dynamic Adaptation	65
7	Conclusion and Future Directions	67
7.1	Summary of Findings	67
7.2	Implications: Concept Drift and Dynamic Adaptation	68
7.3	Future Research Directions	68
A	Code and Data Availability	70
B	Additional Quantum Circuit Architectures	71
B.1	Evaluated Architectures	71
B.2	Benchmark Comparison	73
B.3	Detailed Benchmark Results	73
C	Identified Market Regimes	75
D	Supplementary Context Plots	76
D.1	Baseline Stability (2011–2013)	76
D.2	Robustness Check (2016–2018)	76
Appendices		70
Bibliography		78

Chapter 1

Introduction

Financial markets represent one of the most intricate complex systems in the modern world. Unlike physical systems governed by immutable laws, markets are driven by the interactions of millions of heterogeneous agents acting on imperfect information. Consequently, the central challenge in quantitative finance is not merely predicting the future price of an asset but quantifying and managing the profound uncertainty that characterises these systems.

For decades, the standard approach to this challenge has been to seek equilibrium. Traditional econometric models such as the Autoregressive Integrated Moving Average (ARIMA) and Generalised Autoregressive Conditional Heteroskedasticity (GARCH) rely on the assumption that the underlying stochastic process governing returns is fundamentally stable. These tools have proved exceptionally powerful for modelling markets in their stable states. However, they share a critical vulnerability: they are often static. They struggle to account for the abrupt and non-linear structural breaks that define financial crises. When a market transitions from a period of stability to one of turbulence, the statistical properties of the system change so drastically that a model trained on the former regime becomes obsolete when applied to the latter.

This thesis addresses this limitation by proposing a change in perspective. Instead of forcing a single classical model to fit the entire history of a time series, we adopt a physics-inspired framework that views the market as a system capable of existing in distinct energy levels. Drawing upon the recent formalism of *Quantum Probability* proposed by Li [1], we model the market transition between unimodal and multimodal states as a natural phase transition driven by trading volume rather than a failure of statistics.

It is important to clarify that we do not assert that the stock market is a microscopic quantum object. Rather, we employ the mathematical machinery of quantum mechanics because it offers a superior language for describing systems with interfer-

ence effects and state superposition. To validate this theoretical stance, we adapt the *Quantum Two-Sample Test* methodology recently proposed by Garvin et al. [2]. By leveraging the high-dimensional feature space of a quantum kernel, we aim to demonstrate that these market regimes are statistically distinct and that identifying them is the prerequisite for any robust forecasting strategy.

The development of this work was carried out in collaboration with *Data Reply*. This partnership provided a valuable opportunity to bridge academic research with industrial application standards, allowing for the deepening of both theoretical and practical expertise in Quantum Computing through the exploration of novel concepts, libraries, and real-world financial applications.

1.1 Research Objectives

The primary goal of this work is to bridge the gap between the theoretical potential of Quantum Finance and rigorous empirical validation. Specifically, this thesis aims to:

1. Analyse the statistical properties of the SZSE Component Index to identify the limitations of classical Gaussian assumptions and confirm the presence of multimodality.
2. Apply a volume-based energy threshold to segment the time series into distinct calm and turbulent regimes.
3. Implement and execute a *Quantum Two-Sample Test* to provide statistical validation that these regimes represent fundamentally different probability distributions.
4. Demonstrate how this regime-aware approach can inform a hybrid forecasting pipeline.

1.2 Thesis Structure

The remainder of this thesis is organised as follows.

Chapter 2 establishes the classical financial context. It defines the fundamental concepts of returns and stationarity, reviews the standard ARIMA-GARCH modelling framework, and highlights the specific stylized facts of the market that motivate our search for a new approach.

Chapter 3 provides the necessary primer on quantum computing. It reviews the mathematical formalism of Hilbert spaces, qubits, and quantum circuits, laying the groundwork for the more advanced concepts applied in later chapters.

Chapter 4 introduces the theoretical core of the work. It presents the *Quantum Probability* framework for asset return modelling and details how trading volume can be interpreted as an energy parameter that drives the system from a ground state to excited, multimodal states.

Chapter 5 describes the computational methodology. It explains the Quantum Machine Learning techniques employed, specifically focusing on Quantum Kernels and the implementation of the *Quantum Two-Sample Test* used to measure the distance between market regimes.

Chapter 6 presents the empirical analysis of the SZSE dataset. It details the identification of the regime-switching threshold, the results of the quantum statistical tests, and demonstrates the performance of a regime-aware hybrid forecasting pipeline.

Chapter 7 summarises the findings, discusses the implications for financial risk management, and outlines avenues for future research in quantum-enhanced forecasting.

Chapter 2

Financial Background and Classical Modelling

The quest to find order in the apparent chaos of financial markets has long fascinated researchers. The foundation of quantitative finance is formally traced to 1900, when Louis Bachelier proposed in his thesis *Théorie de la Spéculation* that stock price fluctuations could be modelled using the mathematics of Brownian motion [3]. Notably, his mathematical derivation of diffusion processes predated Einstein's physical explanation of the same phenomenon by five years.

This paradigm evolved throughout the 20th century, culminating in two pillars of modern finance: the Black-Scholes model [4] and the Efficient Market Hypothesis [5]. The Black-Scholes model provided the first closed-form solution for pricing options by assuming returns follow a log-normal distribution. Concurrently, Fama's Efficient Market Hypothesis asserted that prices instantly reflect all available information, implying that future price movements are fundamentally unpredictable and follow a random walk. While mathematically elegant, this Gaussian framework often clashes with empirical reality, failing to account for the crashes and extreme events observed in actual markets.

This chapter establishes the theoretical toolkit required to understand the econometric analysis performed in this thesis. The chapter is organised as follows. In Section 2.1, we define the financial time series and the specific return metrics used. In Section 2.2, we examine the stylized facts of asset returns that contradict classical Gaussian assumptions. In Section 2.3, we review the standard forecasting tools, ARIMA and GARCH. Finally, in Section 2.4, we identify the inability of these static models to handle regime switching, motivating the quantum approach.

2.1 The Nature of Financial Time Series

Financial econometrics applies statistical methods to financial data [6]. The primary data source is the *time series*, a sequence of data points indexed in chronological order. In this study, we analyse the Shenzhen Stock Exchange (SZSE) Component Index. This index tracks the 500 largest companies on the Shenzhen exchange and serves as a key barometer for the Chinese technology and manufacturing sectors.

Asset Returns: The Velocity of the Market

A common error in early analysis is to model the price of an asset P_t directly. Prices typically wander arbitrarily and have no fixed mean, a property known as non-stationarity. To perform rigorous statistical analysis, we must transform prices into *returns*. We distinguish between two types of returns used in this thesis.

Interday Log-Returns

For time series forecasting and statistical testing, we use the *logarithmic return*. It is defined as the natural logarithm of the ratio of consecutive closing prices:

$$r_t = \ln \left(\frac{P_t^{\text{Close}}}{P_{t-1}^{\text{Close}}} \right) \quad (2.1)$$

Log-returns are preferred because they are time-additive and typically display statistical properties closer to stationarity than raw prices.

Intraday Returns

Later in this work, specifically for the Lin Li framework discussed in Chapter 4, we will require a measure of the daily trading energy. For this purpose, we define the *intraday return* which captures the price excursion during a single trading session:

$$r_t^{\text{intra}} = \ln \left(\frac{P_t^{\text{Close}}}{P_t^{\text{Open}}} \right) \quad (2.2)$$

Unless specified otherwise, the term return in the remainder of this text refers to the Interday Log-Return r_t .

2.2 Stylized Facts: The Anomalies

Decades of empirical research have revealed that financial returns do not behave like the idealised random walks of Black-Scholes. They exhibit a set of persistent

statistical properties known as *stylized facts* [7].

2.2.1 Volatility Clustering

The most pervasive fact is that volatility is not constant. In finance, volatility is the standard proxy for risk. Mandelbrot noted that large changes tend to be followed by large changes, of either sign, while small changes tend to be followed by small changes [8]. This phenomenon is known as *volatility clustering*. It implies that while the direction of the market may be unpredictable, the magnitude of the move is correlated with the recent past. Figure 2.1 illustrates this for the SZSE index, where periods of relative calm are punctuated by clusters of extreme turbulence.

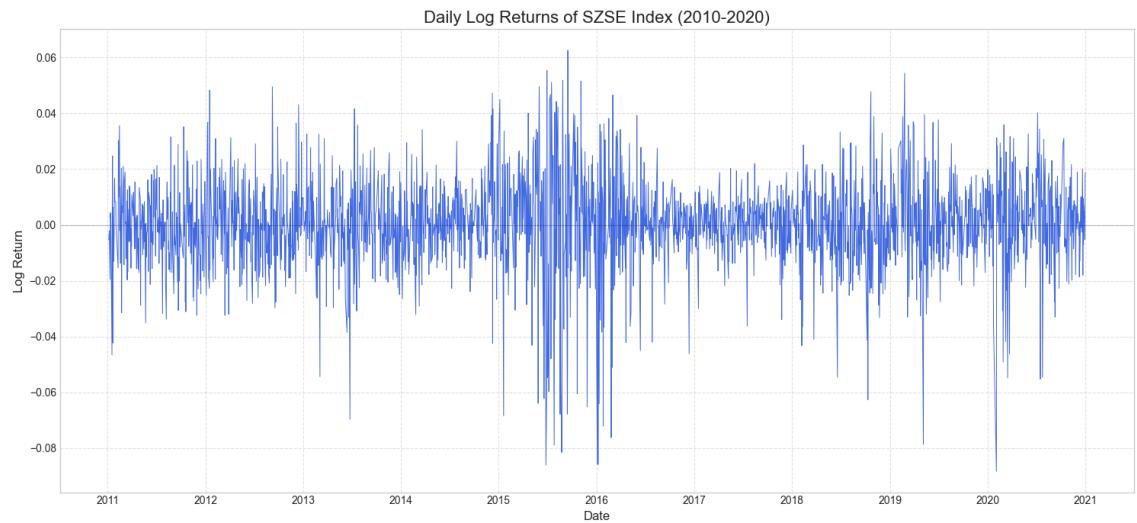


Figure 2.1: Daily log-returns of the SZSE Component Index (2011-2020). The data exhibits clear volatility clustering, where high-variance days are grouped together.

2.2.2 Heavy Tails and Leptokurtosis

In a Gaussian world, events exceeding 3 standard deviations (3σ) are extremely rare. In financial markets, 5σ or even 10σ events happen with alarming frequency. This is known as the heavy tail problem. We quantify this deviation from normality using two metrics calculated on our SZSE dataset:

- *Skewness* (g_1) measures asymmetry. A normal distribution has $g_1 = 0$. Our analysis of the SZSE index yields a skewness of -0.72. This negative value indicates a left-skewed distribution, meaning that large negative returns or crashes are more common and more violent than large positive returns.
- *Excess Kurtosis* (k_4) measures the thickness of the tails. A normal distribution has an excess kurtosis of 0. The SZSE returns exhibit an excess kurtosis of

3.9. This property, known as leptokurtosis, confirms that the probability of extreme events is significantly higher than classical models assume.

Figure 2.2 visualises this reality. The empirical distribution in blue has a sharper peak and fatter tails than the theoretical Gaussian shown by the black dashed line.

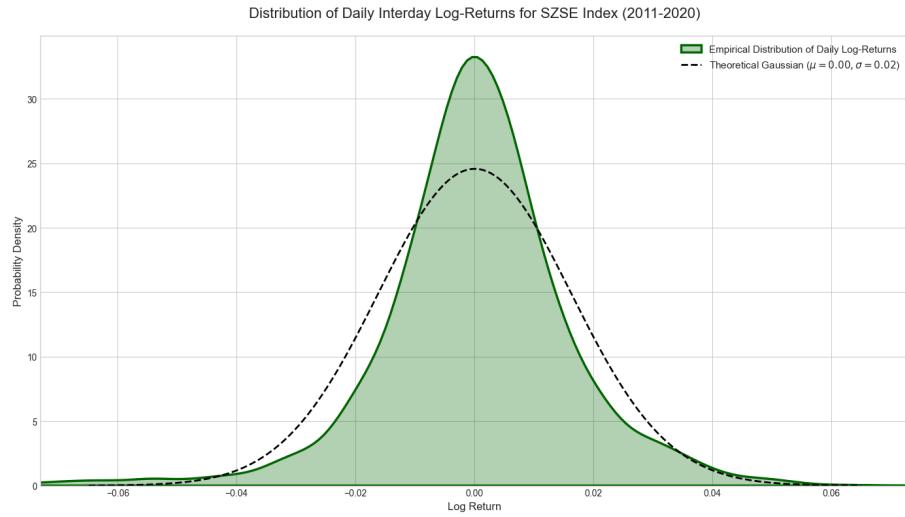


Figure 2.2: Kernel Density Estimate of SZSE returns compared to a Normal distribution. The higher peak and fatter tails indicate a deviation from Gaussian assumptions.

2.3 Classical Forecasting Models

To navigate this complexity, econometrics has developed models that attempt to capture the memory of the market. Understanding these models requires defining a core concept: *autocorrelation*.

2.3.1 Autocorrelation

Autocorrelation measures the relationship between a variable's current value and its past values. If the price of a stock today is strongly influenced by its price yesterday, the series is said to be autocorrelated. Positive autocorrelation indicates a trend, while negative autocorrelation indicates mean reversion. We measure this using the Autocorrelation Function (ACF), which plots the correlation at different time lags. This tool is essential for selecting the parameters of the models described below.

2.3.2 ARIMA: Modelling the Linear Trend

The *Autoregressive Integrated Moving Average* (ARIMA) model [9] is the standard tool for forecasting the mean of a time series. An ARIMA(p, d, q) model is defined

by three parameters:

1. *Autoregression (p)*: This parameter dictates how many past days influence the current price. An $AR(1)$ model implies that today's return depends directly on yesterday's return. This captures the momentum of the market.
2. *Integration (d)*: This parameter specifies how many times the raw data must be differenced to become stationary. For financial prices which wander arbitrarily, we typically use $d = 1$ which is equivalent to using returns.
3. *Moving Average (q)*: This parameter dictates how the model reacts to past forecast errors, often termed shocks. An $MA(1)$ model implies that if the model made a surprise error yesterday, that error will continue to ripple into today's prediction.

These parameters are typically chosen by minimising the Akaike Information Criterion (AIC), a metric that balances the model's accuracy against its complexity to prevent overfitting.

2.3.3 GARCH: Modelling the Volatility

While ARIMA predicts the direction of the price, it assumes the risk or variance is constant. As we saw in Section 2.2, this is false since volatility comes in clusters. To model this, Engle and Bollerslev developed the GARCH(p, q) framework [10].

The standard GARCH(1,1) model predicts the variance σ_t^2 based on three terms:

$$\sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (2.3)$$

Here, the coefficients have precise physical meanings:

- *Reaction to Shocks (α)* determines how jumpy the volatility is. A high α means the market reacts violently to a sudden crash.
- *Persistence (β)* determines the memory of the volatility. A high β means that once the market enters a crisis, it stays turbulent for a long time.
- *Stability* requires that $\alpha + \beta < 1$.

In our SZSE analysis, we typically observe high β values, confirming the sticky nature of financial crises.

2.4 The Failure Point: Regime Switching

The combination of ARIMA for the mean and GARCH for the volatility constitutes the standard industry baseline. However, these models share a fundamental limitation in that they are global. They estimate a single set of parameters $\phi, \theta, \alpha, \beta$ that attempts to fit the entire history of the market.

But markets are not static. As we will demonstrate in Chapter 4, the market undergoes phase transitions, shifting between a Ground State characterised by calm unimodal distributions and Excited States characterised by turbulent multimodal distributions. A static GARCH model trained on a decade of data behaves like a thermostat set to the average annual temperature. It will be too cold in winter and too hot in summer. This limitation motivates our research into whether Quantum Probability can be used to detect these regimes dynamically and switch between models adaptively.

Chapter 3

Quantum Fundamentals

Financial markets represent complex adaptive systems driven by the intricate interactions of countless heterogeneous agents. Chapter 2 established that traditional modelling frameworks often struggle to capture the stylized facts of these systems, such as heavy tails and regime shifts. This limitation motivates the search for a more expressive mathematical language. This thesis explores whether quantum theory can offer this richer framework, positioning our work within the active research field of *Quantum Finance*.

The potential of applying quantum principles to finance is twofold. First, it provides a new mathematical syntax where concepts like superposition and complex probability amplitudes allow uncertainty and agent interaction to be treated naturally. Second, it offers a new computational tool. Quantum computing achieves its power by leveraging unique resources such as interference which, when combined in multi-qubit systems through entanglement, enable computational possibilities intractable for classical machines.

Clarifying the terminology surrounding quantum computational power is crucial at this stage. *Quantum supremacy* refers to the demonstration that a programmable device can solve a problem that no classical computer can solve in a feasible amount of time, regardless of the practical utility of that problem. A more practical goal is *quantum advantage*, which denotes the point where a quantum computer solves a relevant real-world problem significantly faster or more accurately than the best-known classical algorithms [11]. The research presented in this thesis operates in the current Noisy Intermediate-Scale Quantum (NISQ) era, where the primary focus is the pursuit of quantum advantage on specific, carefully chosen tasks rather than demonstrating universal supremacy.

This chapter constructs the theoretical foundation required to understand the proposed methodology. Section 3.1 defines the axiomatic postulates of quantum mechanics. Section 3.2 derives the fundamental unit of quantum information, the

qubit, and its geometric representation. Section 3.3 introduces the operational tools of quantum computation, namely gates and circuits, which are necessary to manipulate quantum information. Finally, Section 3.5 details the core concepts of Quantum Machine Learning, laying the groundwork for the parameterised circuits and quantum kernels used in the empirical analysis.

3.1 The Postulates of Quantum Mechanics

Quantum mechanics is the theory that describes the dynamics of the microscopic world. Its framework is built upon a concise set of fundamental axioms, or postulates, that connect an abstract mathematical description to physically observable phenomena. It is standard practice in foundational texts to introduce the theory through these formal rules [12–14]. We will follow this canonical approach. As the language of this theory is fundamentally that of linear algebra, our approach will be to pause within the discussion of each postulate to formally define and clarify the key mathematical concepts it employs.

Postulate 1 (State Space). *Associated to any isolated physical system is a complex vector space with inner product (that is, a Hilbert space) known as the state space. The system is completely described by its state vector, which is a unit vector in the systems state space.*

This first postulate is the cornerstone of the theory. It makes a profound claim: the state of a physical system is a *vector*. These vectors live in a mathematical structure known as a *Hilbert space*. Conceptually, a Hilbert space is a complex vector space equipped with an inner product, which allows us to define geometric notions such as length and angle, essential for normalization and orthogonality.

A final, more technical requirement is that the space must be *complete*¹. While this is a critical distinction in infinite-dimensional settings, it is a standard theorem of linear analysis that any finite-dimensional inner product space is automatically complete [15]. As this thesis is concerned exclusively with finite-dimensional systems, this condition is always satisfied.

To work with these vectors, we use the elegant formalism introduced by Paul Dirac [14].

Definition 3.1.1 (Bra-Ket Notation). *In Dirac notation, the elements of the Hilbert space are represented as follows:*

¹In mathematics, a space is *complete* if every Cauchy sequence of elements in the space converges to a limit that is also within the space. This property is what formally distinguishes a Hilbert space from more general inner product spaces.

- A state vector is denoted by a ket, $|\psi\rangle$. It is mathematically equivalent to a column vector.
- For every ket $|\psi\rangle$, there exists a corresponding bra, $\langle\psi|$, which is its conjugate transpose (a row vector of complex conjugates).
- The inner product of a state $|\psi\rangle$ with a state $|\phi\rangle$ is denoted by $\langle\phi|\psi\rangle$. Two states are orthogonal if their inner product is zero.

The direct consequence of the vector space structure is the *principle of superposition*. Since states are vectors, we can form linear combinations: if $|a\rangle$ and $|b\rangle$ are two valid states, then their sum (properly normalised) is also a valid state.

The perfect illustration of this postulate is the *qubit*, the fundamental unit of quantum information. While a classical bit is confined to the states 0 or 1, a qubit's state is a vector in a two-dimensional Hilbert space. The notation for this space is \mathbb{C}^2 , where the ' \mathbb{C} ' signifies that the space is defined over the complex numbers, and the superscript '2' indicates its dimension. A dimension of two means that any vector in the space can be written as a linear combination of just two basis vectors that span the space [12].

By convention, we choose a particular orthonormal basis called the *computational basis*, whose vectors are denoted by $|0\rangle$ and $|1\rangle$. We can give these abstract kets an explicit representation as column vectors:

$$|0\rangle \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (3.1)$$

A general state $|\psi\rangle$ is therefore a superposition of these basis vectors:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad (3.2)$$

where the complex coefficients $\alpha, \beta \in \mathbb{C}$ are known as *probability amplitudes*, and must satisfy the normalization condition

$$|\alpha|^2 + |\beta|^2 = 1. \quad (3.3)$$

A full, detailed exploration of the qubit's properties, including its geometric representation on the Bloch sphere, is provided in Section 3.2.

It is precisely this ability to inhabit a continuum of superposition states that makes the qubit a far richer information carrier than its classical counterpart. This richness, however, presents a paradox: although the state is described by continuous complex amplitudes, a single measurement will always yield just one classical bit.

The resolution to this puzzle lies in the nature of quantum measurement: a process that does not merely reveal pre-existing information, but actively transforms the state itself.

To describe such transformations, whether induced by measurement or by time evolution, we require a new mathematical object: the *operator*. Operators are the tools that allow us to model how quantum states change and how physical quantities are extracted from them.

Operators and Observables The state vector $|\psi\rangle$ provides a complete, yet abstract, description of a quantum system. To connect this mathematical object to the physical world, we must represent the tangible properties we can actually measure in an experiment, such as position, momentum, and energy. In the language of quantum mechanics, these measurable quantities are called *observables*, and they are represented by a special class of mathematical objects known as *operators*.

Formally, an operator \hat{A} is a linear map that acts on a ket to produce another ket within the same Hilbert space. In finite-dimensional systems, every linear operator admits a matrix representation, which we denote by A (without the hat). The action of \hat{A} on a state vector corresponds to matrix multiplication in this representation. To construct and interpret these matrices, it is useful to first define the *outer product*.

Definition 3.1.2 (Outer Product). *Given a ket $|\psi\rangle$ and a bra $\langle\phi|$, their outer product is the operator $|\psi\rangle\langle\phi|$. This operator maps a vector $|\chi\rangle$ to the scaled vector $|\psi\rangle\langle\phi|\chi\rangle$. In matrix terms, if $|\psi\rangle$ and $|\phi\rangle$ are column vectors, the outer product corresponds to the multiplication of a column vector by a row vector, resulting in a square matrix.*

For example, the outer product of the basis state $|0\rangle$ with itself yields:

$$|0\rangle\langle 0| = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}. \quad (3.4)$$

Any operator can be defined by its action on the basis vectors. For a single qubit, a general operator \hat{A} can be expressed as a linear combination of outer products:

$$\hat{A} = a|0\rangle\langle 0| + b|0\rangle\langle 1| + c|1\rangle\langle 0| + d|1\rangle\langle 1|, \quad (3.5)$$

where the complex coefficients a, b, c, d determine the operators action.

A key insight is that the columns of an operator's matrix are determined by its

action on the basis vectors. If an operator \hat{A} acts on the computational basis as:

$$\hat{A} |0\rangle = a |0\rangle + c |1\rangle \quad (3.6)$$

$$\hat{A} |1\rangle = b |0\rangle + d |1\rangle \quad (3.7)$$

then the matrix A representing \hat{A} is:

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}. \quad (3.8)$$

The first column corresponds to $\hat{A} |0\rangle$, and the second to $\hat{A} |1\rangle$. This dual perspective (abstract operator and concrete matrix) will be particularly useful when we introduce quantum gates in Section 3.3, where operators correspond to unitary matrices that manipulate qubit states.

While an operator can be defined by its action on any basis, a more powerful way to interpret its action comes from analysing its *eigenvectors* and *eigenvalues*.

Definition 3.1.3 (Eigenvectors and Eigenvalues). *Let \hat{A} be a linear operator acting on a Hilbert space. A non-zero vector $|\psi\rangle$ is called an eigenvector of \hat{A} if the action of \hat{A} on $|\psi\rangle$ simply scales it by a complex number λ , called the eigenvalue:*

$$\hat{A} |\psi\rangle = \lambda |\psi\rangle. \quad (3.9)$$

This equation tells us that $|\psi\rangle$ is unchanged in direction by the operator: it is only stretched or rotated in phase. Geometrically, eigenvectors are those special directions that remain invariant under the transformation, up to a scaling factor.

Example: Pauli-Z Operator. Consider the Pauli-Z operator:

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (3.10)$$

Its action on the computational basis states is:

$$Z |0\rangle = |0\rangle, \quad Z |1\rangle = -|1\rangle. \quad (3.11)$$

Here, $|0\rangle$ and $|1\rangle$ are eigenvectors of Z , with eigenvalues $+1$ and -1 , respectively. This means that Z leaves $|0\rangle$ unchanged, while it flips the sign of $|1\rangle$.

The concept of eigenvectors is particularly powerful when applied to the operators central to quantum mechanics. For the operators that describe physical pro-

cesses, a crucial property holds: their eigenvectors can be chosen to form a complete orthonormal basis for the state space.

This basis, known as the *eigenbasis*, is extremely useful. When an operator is expressed in its own eigenbasis, its matrix representation simplifies to a diagonal form, where the diagonal entries are precisely its eigenvalues. The physical role of an operator is directly encoded in the mathematical properties of its eigenvalues. Based on these properties, we can identify two fundamental classes of operators that govern all of quantum dynamics: the unitary operators for evolution, and the Hermitian operators for measurement.

Definition 3.1.4 (Unitary Operator). *An operator \hat{U} is unitary if its conjugate transpose is also its inverse ($\hat{U}^\dagger \hat{U} = I$). Its eigenvalues are always complex numbers with a modulus of one, i.e., of the form $e^{i\theta}$, where $\theta \in \mathbb{R}$.*

This mathematical condition has a profound physical meaning: it guarantees that the length of the state vector is preserved during the transformation, thus conserving the total probability of the system. All *quantum gates*, which we will detail in Section 3.3, are by definition unitary operators.

Definition 3.1.5 (Hermitian Operator). *An operator \hat{A} is Hermitian if it is equal to its own conjugate transpose ($\hat{A} = \hat{A}^\dagger$). Its eigenvalues are always real numbers.*

Since the outcome of a physical measurement must be a real number, this property makes Hermitian operators the unique mathematical candidates for representing physical *observables*.

An observable is any measurable property of a system; prominent examples include position, momentum, and energy (represented by the Hamiltonian operator).

Equipped with this mathematical language of Hilbert spaces, state vectors, and operators, we can now state the second postulate, which governs the dynamics of an unobserved quantum system.

Postulate 2 (Evolution). *The evolution of a closed quantum system is described by a unitary transformation. That is, the state $|\psi_1\rangle$ of the system at time t_1 is related to the state $|\psi_2\rangle$ at a later time t_2 by a unitary operator $U(t_1, t_2)$ such that:*

$$|\psi_2\rangle = U |\psi_1\rangle \quad (3.12)$$

This is the *law of motion* for an unobserved system. A more refined version

describes the evolution in continuous time through the *Schrödinger equation*:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle \quad (3.13)$$

where \hat{H} is a Hermitian operator known as the *Hamiltonian*, which corresponds to the total energy of the system, and \hbar is the reduced Planck constant ($\hbar = h/2\pi$). The two formulations are equivalent, as the solution to the Schrödinger equation gives the unitary operator:

$$U(t_2, t_1) = e^{-i\hat{H}(t_2-t_1)/\hbar}. \quad (3.14)$$

The linearity of this evolution is a defining feature of quantum mechanics and has significant implications. As we will see in Section 3.5, it means that the activation functions used in Quantum Machine Learning must be fundamentally different from the non-linear functions, such as ReLU, that are common in classical deep learning.

The unitary evolution described in Postulate 2 governs the dynamics of a quantum system as long as it remains isolated. However, to extract information, an observer must interact with it. This act, known as *measurement*, marks a profound and conceptually challenging departure from unitary evolution. While the mathematical rules for predicting its outcomes are extraordinarily successful, their physical interpretation is the subject of intense debate, a foundational issue known as the *measurement problem* [13]. The formalism posits that measurement is inherently probabilistic and irreversible, a stark contrast to the deterministic and reversible nature of unitary evolution. We now formalize these operational rules.

Postulate 3 (Measurement). *Quantum measurements are described by a collection of measurement operators, $\{M_m\}$, acting on the state space. The index m refers to the possible measurement outcomes. If the state is $|\psi\rangle$ immediately before measurement, the probability of result m is given by the Born rule:*

$$p(m) = \langle \psi | M_m^\dagger M_m | \psi \rangle \quad (3.15)$$

and the state after measurement is:

$$\frac{M_m |\psi\rangle}{\sqrt{p(m)}}. \quad (3.16)$$

The measurement operators satisfy the completeness relation:

$$\sum_m M_m^\dagger M_m = I. \quad (3.17)$$

This is the most general formulation of a measurement. A vast and intuitive subclass are the *projective measurements*, which are directly associated with observables. In this case, the measurement operators have the special structure of projection operators.

Definition 3.1.6 (Projection Operator). *A projection operator (or projector) is an operator P that is both Hermitian ($P = P^\dagger$) and idempotent ($P^2 = P$). The outer product of a normalised state vector with its corresponding bra, $P_{|\psi\rangle} = |\psi\rangle\langle\psi|$, is a projector onto the one-dimensional subspace spanned by $|\psi\rangle$.*

In a projective measurement, the operators $\{M_m\}$ are a set of orthogonal projectors $\{P_m\}$ that correspond to the eigenspaces of the observable being measured, and the outcome m is the corresponding eigenvalue. To make this concrete, consider the measurement of a qubit in the computational basis, which corresponds to measuring the Pauli-Z observable. The projectors onto its eigenspaces are:

$$P_0 = |0\rangle\langle 0| = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad P_1 = |1\rangle\langle 1| = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \quad (3.18)$$

For a general state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the probability of obtaining the outcome associated with state $|0\rangle$ is then:

$$p(0) = \langle\psi|P_0|\psi\rangle = |\alpha|^2. \quad (3.19)$$

The third postulate is profoundly counter-intuitive. It first establishes that outcomes are fundamentally *probabilistic*. Second, it describes the *collapse of the wave function*: measurement is an invasive process that projects the state onto a new one. This disturbance implies a fundamental limit on which properties can be known simultaneously. If two observables do not commute, the *Heisenberg Uncertainty Principle* imposes a strict trade-off [16]. For position (x) and momentum (p), this is famously expressed by the inequality:

The third postulate is profoundly counter-intuitive. It first establishes that outcomes are fundamentally *probabilistic*. Second, it describes the *collapse of the wave function*: measurement is an invasive process that projects the state onto a new one. This disturbance implies a fundamental limit on which properties can be known simultaneously. If two observables do not commute, the *Heisenberg Uncertainty Principle* imposes a strict trade-off [16]. For position (x) and momentum (p), this is famously expressed by the inequality:

$$\sigma_x\sigma_p \geq \frac{\hbar}{2}, \quad (3.20)$$

where σ_x and σ_p are the standard deviations of the position and momentum measurements, respectively.

Having established the rules for single systems, we must now consider how to describe systems composed of multiple parts. This requires a new mathematical tool for combining their respective state spaces.

Definition 3.1.7 (Tensor Product). *The tensor product, denoted by the symbol \otimes , is an operation that combines vector spaces to create a larger composite space whose dimension is the product of the individual dimensions.*

Let us consider two qubits in the states:

$$|a\rangle = \begin{pmatrix} a_0 \\ a_1 \end{pmatrix}, \quad |b\rangle = \begin{pmatrix} b_0 \\ b_1 \end{pmatrix}.$$

Their tensor product is the four-dimensional vector:

$$|a\rangle \otimes |b\rangle = \begin{pmatrix} a_0 b_0 \\ a_0 b_1 \\ a_1 b_0 \\ a_1 b_1 \end{pmatrix}. \quad (3.21)$$

With this operation defined, we can now state the final postulate.

Postulate 4 (Composite Systems). *The state space of a composite physical system is the tensor product of the state spaces of the component physical systems. Moreover, if we have systems numbered 1 through n , and system number i is prepared in the state $|\psi_i\rangle$, then the joint state of the total system is $|\psi_1\rangle \otimes |\psi_2\rangle \otimes \cdots \otimes |\psi_n\rangle$.*

This mathematical structure allows for the existence of *entangled states*: global states of a composite system that cannot be described by assigning a definite state to each of its individual parts. This property of non-separability is the fundamental resource enabling quantum information processing. Its implications are groundbreaking, particularly in cybersecurity, where it is the basis for provably secure communication, and in finance, where it offers potential speed-ups for complex optimisation problems. The nature of this resource will be discussed in detail in Section 3.4.

These postulates form the complete foundation of quantum mechanics. They provide the rules for describing the state of any quantum system, how that state evolves in time, and what happens when we attempt to measure it. With these abstract laws now established, we are ready to explore their concrete implications by applying them to the central objects of quantum computation: the qubit.

3.2 The Qubit and its Geometric Representation

The postulates provide the abstract rules of quantum theory. We now apply these rules to the fundamental object of quantum computation: the *qubit*. A classical bit can exist in one of two definite states, 0 or 1. By contrast, a qubit leverages the principle of superposition to occupy a much richer state space.

Definition 3.2.1 (Qubit). *A qubit is a two-level quantum system whose state space is the two-dimensional complex Hilbert space \mathbb{C}^2 . An orthonormal basis for this space is chosen by convention and denoted as $\{|0\rangle, |1\rangle\}$, representing the classical bit values.*

An arbitrary pure state of a qubit, $|\psi\rangle$, is a unit vector in this space, written as a linear combination of the basis states:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (3.22)$$

where the coefficients $\alpha, \beta \in \mathbb{C}$ are complex numbers known as *probability amplitudes*.

Although this algebraic description is complete, it can feel abstract. While the state appears to require four real parameters (the real and imaginary parts of α and β), we can show that two are redundant.

First, let us express the complex amplitudes in their polar form: $\alpha = r_\alpha e^{i\delta_\alpha}$ and $\beta = r_\beta e^{i\delta_\beta}$. The state can be written as:

$$|\psi\rangle = r_\alpha e^{i\delta_\alpha} |0\rangle + r_\beta e^{i\delta_\beta} |1\rangle. \quad (3.23)$$

The normalization condition, $|\alpha|^2 + |\beta|^2 = 1$, implies that $r_\alpha^2 + r_\beta^2 = 1$. Since r_α and r_β are real and non-negative, this constraint allows us to represent them using a single real parameter, an angle θ , such that:

$$r_\alpha = \cos(\theta/2) \quad \text{and} \quad r_\beta = \sin(\theta/2), \quad \text{for } \theta \in [0, \pi]. \quad (3.24)$$

Substituting these into the state vector yields:

$$|\psi\rangle = \cos(\theta/2) e^{i\delta_\alpha} |0\rangle + \sin(\theta/2) e^{i\delta_\beta} |1\rangle. \quad (3.25)$$

The second key consideration is the irrelevance of the *global phase factor*. Two states that differ only by a global phase, such as $|\psi\rangle$ and $e^{i\gamma}|\psi\rangle$, are physically indistinguishable. We can therefore multiply the entire state by a factor of $e^{-i\delta_\alpha}$ without changing its physical properties. This effectively removes the phase from

the first amplitude:

$$|\psi'\rangle = e^{-i\delta_\alpha} |\psi\rangle = \cos(\theta/2) |0\rangle + \sin(\theta/2) e^{i(\delta_\beta - \delta_\alpha)} |1\rangle. \quad (3.26)$$

The only phase that remains is the *relative phase*, $\phi = \delta_\beta - \delta_\alpha$. These two steps have reduced the four initial real parameters to just two: the polar angle θ and the azimuthal (relative phase) angle ϕ .

This reduction to two real parameters motivates a powerful geometric visualisation: the *Bloch sphere* (Figure 3.1). This construction maps the state space of a single qubit onto the surface of a unit sphere in \mathbb{R}^3 , where the state is uniquely defined by the angles θ and ϕ we have just derived:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\phi} \sin\left(\frac{\theta}{2}\right) |1\rangle, \quad \text{where } \theta \in [0, \pi] \text{ and } \phi \in [0, 2\pi). \quad (3.27)$$

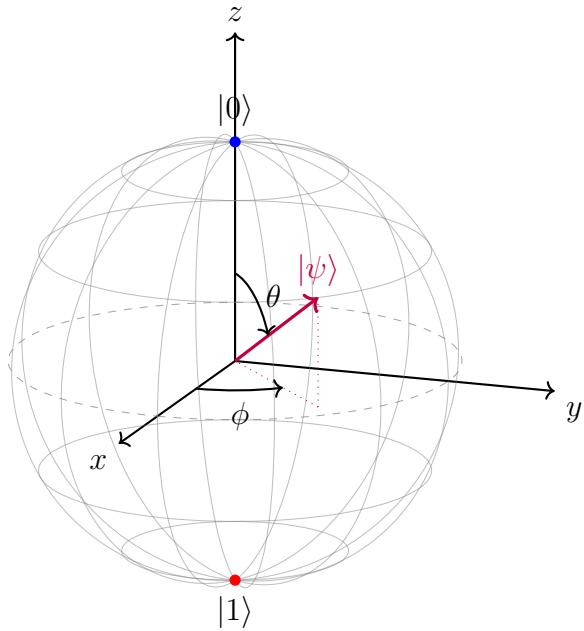


Figure 3.1: The Bloch sphere as a geometric representation of a single-qubit pure state. Each vector $|\psi\rangle$ corresponds to a point on the sphere defined by the angles θ and ϕ . The states $|0\rangle$ and $|1\rangle$ are located at the north and south poles, respectively.

In this mapping, the state $|0\rangle$ corresponds to the *north pole* ($\theta = 0$), while $|1\rangle$ is found at the *south pole* ($\theta = \pi$). Orthogonal states occupy antipodal points.

This representation is more than illustrative; it helps build intuition about the behaviour of quantum gates, which, as we will explore in Section 3.3, act as rotations of the state vector on the surface of the Bloch sphere.

Pure vs. Mixed States. The Bloch sphere, however, tells a richer story, as illustrated in Figure 3.2. The points on the *surface* of the sphere represent *pure states*, for which our knowledge of the system is maximal (corresponding to a Bloch vector of length $|\vec{r}| = 1$). More generally, a quantum system can be in a *mixed state*, which corresponds to a vector in the *interior* of the sphere ($|\vec{r}| < 1$). This can arise either from classical uncertainty over an ensemble of states or from decoherence.

To handle both scenarios, we use the *density operator*, ρ . A generic, or *partially mixed state*, is represented by a vector of length $0 < |\vec{r}| < 1$. The extreme case is the *maximally mixed state*, which represents a complete lack of information about the qubit's orientation. This corresponds to the very center of the sphere, a vector of length $|\vec{r}| = 0$.²

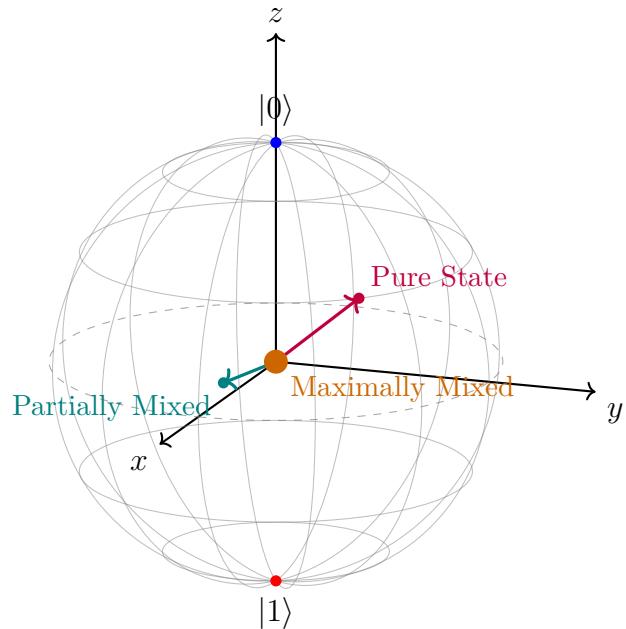


Figure 3.2: Representation of pure and mixed states on the Bloch sphere. Pure states lie on the surface ($r = 1$), partially mixed states lie in the interior ($0 < r < 1$), and the maximally mixed state lies at the origin ($r = 0$).

Definition 3.2.2 (Density Operator). *The density operator ρ for a quantum system is a positive semi-definite operator with unit trace ($\text{tr}(\rho) = 1$).*

- For a pure state $|\psi\rangle$, the density operator is the projector onto that state:

$$\rho = |\psi\rangle\langle\psi|. \quad (3.28)$$

²The most general form for a single-qubit density operator is $\rho = \frac{1}{2}(I + \vec{r} \cdot \vec{\sigma})$, where $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is the vector of Pauli matrices and \vec{r} is the Bloch vector. For a pure state, the vector has unit length, $|\vec{r}| = 1$. For a mixed state, it is shorter, $|\vec{r}| < 1$. The maximally mixed state has $\vec{r} = (0, 0, 0)$, yielding $\rho = \frac{1}{2}I$.

- For a mixed state representing an ensemble of pure states $\{|\psi_i\rangle\}$ with classical probabilities $\{p_i\}$, the density operator is the weighted average of their projectors:

$$\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|. \quad (3.29)$$

This geometric picture provides a complete map of the possible states of a single qubit. The next logical question is how to manipulate this state in a controlled manner; for instance, how to move it from one point on the sphere to another. This is the role of quantum gates, which implement the unitary evolution described in Postulate 2 as a set of precise, reversible operations.

3.3 Quantum Gates and Circuits

The evolution of a quantum state is governed by the Hamiltonian operator via the continuous-time Schrödinger equation, as described in Postulate 2. For a programmable computer, however, we need a way to implement any desired evolution using a finite set of controllable, discrete operations. This requirement leads to the paradigm of *digital quantum computing*. Unlike *analog quantum computers*, which aim to build a physical system whose Hamiltonian directly mimics a specific problem, the digital approach approximates any complex, continuous evolution by breaking it down into a sequence of simple, discrete operations called *quantum gates*.

This process of discretization is formally justified by product formulas like the Lie-Trotter-Suzuki decomposition [17, 18]. The first-order formula states that for two (not necessarily commuting) Hermitian operators A and B , the evolution over a time t can be expressed as:

$$e^{-i(A+B)t/\hbar} = \lim_{n \rightarrow \infty} \left(e^{-iAt/n\hbar} e^{-iBt/n\hbar} \right)^n. \quad (3.30)$$

This powerful result means that a complex evolution generated by a Hamiltonian $H = A + B$ can be approximated by rapidly alternating between simpler evolutions generated by A and B , each applied for a small time slice $\Delta t = t/n$. It is a non-trivial result because, in general, $e^{A+B} \neq e^A e^B$ unless the operators A and B commute [17]. These simpler evolutions correspond to our fundamental quantum gates. This thesis focuses exclusively on this *digital gate-based model*. Each quantum gate is a unitary operation, ensuring that the computation is, in principle, reversible.

3.3.1 Single-Qubit Gates

Gates that operate on individual qubits are represented by 2×2 unitary matrices. Their actions can be visualised as rotations of the state vector on the Bloch sphere. Table 3.1 provides a summary of the fundamental single-qubit gates.

The Identity and Pauli Gates. The most fundamental operation is the *Identity gate* (I), which leaves the qubit state unchanged. Alongside it, the *Pauli gates* (X, Y, Z) are foundational operations in quantum computing. These gates are the operational counterparts to the *Pauli matrices* ($\sigma_x, \sigma_y, \sigma_z$), a set of operators central to the mathematical description of spin in quantum mechanics. In a computational context, the X , Y , and Z gates implement rotations of π radians around the principal axes of the Bloch sphere.

- The *Pauli-X gate* is the quantum analogue of the classical NOT gate, performing a bit-flip. Its action on the basis states is:

$$X |0\rangle = |1\rangle, \quad X |1\rangle = |0\rangle. \quad (3.31)$$

- The *Pauli-Y gate* applies both a bit-flip and a phase-flip. Its action is:

$$Y |0\rangle = i |1\rangle, \quad Y |1\rangle = -i |0\rangle. \quad (3.32)$$

- The *Pauli-Z gate* is a phase-flip gate. It leaves $|0\rangle$ unchanged but applies a phase of -1 to the $|1\rangle$ state:

$$Z |0\rangle = |0\rangle, \quad Z |1\rangle = -|1\rangle. \quad (3.33)$$

It is a key result that the Pauli matrices, together with the identity matrix, form a basis for the space of all single-qubit operators.

The Hadamard Gate and Basis Changes. The *Hadamard gate* (H) is crucial for creating superposition and for changing basis. While we typically work in the computational basis $\{|0\rangle, |1\rangle\}$, it is often useful to switch to the diagonal basis, $\{|+\rangle, |-\rangle\}$, whose vectors are defined by the action of the Hadamard gate:

$$H |0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \equiv |+\rangle \quad (3.34)$$

$$H |1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \equiv |-\rangle \quad (3.35)$$

This change of basis is powerful because the action of certain gates becomes simpler to understand. For example, the Z gate, which applies a phase-flip in the computational basis, acts as a bit-flip in the diagonal basis ($Z|+\rangle = |-\rangle$), while the X gate acts as a phase-flip in this new basis ($X|+\rangle = |+\rangle, X|-\rangle = -|-\rangle$).

General Rotation Gates. While the Pauli gates perform fixed rotations, a more general set of operations are the continuous *rotation gates*, $R_x(\theta)$, $R_y(\theta)$, and $R_z(\theta)$, which perform a rotation of an arbitrary angle θ about an axis. These are essential for the parameterised circuits used in Quantum Machine Learning.

Gate	Symbol	Matrix	Description
Identity		$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	Leaves the state unchanged.
Pauli-X		$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	Performs a bit-flip; a π rotation around the x-axis.
Pauli-Y		$\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$	Performs a π rotation around the y-axis.
Pauli-Z		$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$	Performs a phase-flip; a π rotation around the z-axis.
Hadamard		$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$	Creates a uniform superposition.
Phase (S)		$\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$	A $\pi/2$ rotation around the z-axis. It is the square root of the Z gate ($S^2 = Z$).
$\pi/8$ (T)		$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$	A $\pi/4$ rotation around the z-axis.
Rotation-X		$\begin{pmatrix} \cos(\theta/2) & -i \sin(\theta/2) \\ -i \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$	Rotation of angle θ around the x-axis.
Rotation-Y		$\begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$	Rotation of angle θ around the y-axis.
Rotation-Z		$\begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}$	Rotation of angle θ around the z-axis.

Table 3.1: A summary of fundamental single-qubit gates, showing their circuit symbol, matrix representation, and description of their action.

3.3.2 Multi-Qubit Gates and Universality

While single-qubit gates allow for arbitrary state manipulation on individual qubits, they are fundamentally insufficient for universal quantum computation. Acting alone, they can only produce states that are separable, meaning the total state of an n -qubit system can be written as a tensor product of individual qubit states: $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes \cdots \otimes |\psi_n\rangle$. To unlock the full power of quantum mechanics, we require multi-qubit gates capable of generating *entanglement*.

The canonical two-qubit gate is the *Controlled-NOT (CNOT) gate*. It acts on two qubits, a *control* and a *target*, and performs a Pauli-X (NOT) operation on the target if, and only if, the control qubit is in the state $|1\rangle$. In circuit diagrams, the control is denoted by a solid dot (\bullet) and the target by a circled plus (\oplus). As it acts on a two-qubit system (a \mathbb{C}^4 space), it is represented by a 4×4 unitary matrix. In the computational basis ordered as $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, its circuit symbol and matrix representation are:

$$\begin{array}{c} \text{Control} \quad \bullet \\ \text{Target} \quad \oplus \end{array} \quad \text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Another important controlled gate is the *Controlled-Z (CZ) gate*. While seemingly different, the two are logically equivalent; a CNOT can be constructed from a CZ gate by applying Hadamard gates to the target qubit before and after the operation. The choice between them is often a matter of which is more *native* or easier to implement on a given hardware architecture. The concept of controlled operations can be extended to multiple controls, such as in the *Toffoli (CCNOT) gate*, which has two control qubits. The power of the digital gate-based model lies in the concept of *universality*. It is a remarkable result that any possible unitary operation can be approximated to an arbitrary degree of accuracy by composing gates from a small, finite set. A set with this property is called a *universal gate set*. While the set of all single-qubit gates and the CNOT gate is universal, a more powerful result is provided by the *Solovay-Kitaev theorem* [12], which guarantees that this approximation can be achieved efficiently. For practical purposes, universality is typically achieved using a finite set of gates, such as the Clifford+T gates³.

This principle of universality is what allows for the construction of powerful quantum subroutines, such as the Quantum Fourier Transform that enables Shor's

³The Clifford group consists of gates (H, S, CNOT) that are not universal on their own. As proven by the *Gottesman-Knill theorem* [19], any circuit composed solely of Clifford gates can be efficiently simulated on a classical computer, and thus cannot provide a quantum speedup. The addition of a non-Clifford gate, such as the T gate, is what makes the set universal and computationally powerful.

factoring algorithm [20]. The most immediate application of these tools, however, is the creation of entangled states. The simple circuit that combines a Hadamard gate with a CNOT gate, shown in Figure 3.3, is sufficient to transform a simple, separable state into a maximally entangled state. This process and its profound implications are so central to quantum computation that they merit their own dedicated discussion in the following section.

3.4 Entanglement: A Uniquely Quantum Resource

The concept of entanglement is not merely an extension of the superposition principle to multiple qubits; as Schrödinger famously described it, entanglement is not *one* but rather *the* characteristic trait of quantum mechanics [21]. It represents the core resource of the *second quantum revolution*: while the first revolution gave us lasers and transistors by understanding quantum ensembles, the second seeks to control individual quantum states to process information in ways impossible for classical physics.

To define it rigorously, it is best to first define its opposite: separability.

Definition 3.4.1 (Separable and Entangled States). *Let a composite system be described by a state space $V = V_1 \otimes V_2 \otimes \cdots \otimes V_n$.*

- A state $|\psi\rangle \in V$ is called *separable* if it can be written as a tensor product of states from the component subspaces, i.e., $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes \cdots \otimes |\psi_n\rangle$, where $|\psi_i\rangle \in V_i$.
- A state is *entangled* if it is not separable.

A separable state is one where each subsystem has its own well-defined state, independent of the others. An entangled state, by contrast, is one where the constituent parts have lost their individual identities in favour of a global, correlated state. The most direct way to understand how such a state is created is to see it emerge from the action of multi-qubit gates.

The simple circuit in Figure 3.3 uses just one Hadamard gate and one CNOT gate to transform a simple, separable state into a maximally entangled state known as a *Bell state*.

The process is as follows:

1. We start with two qubits in the separable state $|00\rangle$.
2. A Hadamard gate is applied to the first qubit, creating the superposition $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |10\rangle)$.

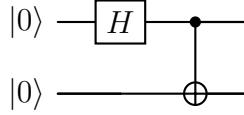


Figure 3.3: The quantum circuit for generating the Bell state $|\Phi^+\rangle$ from the initial state $|00\rangle$.

3. A CNOT gate is then applied, with the first qubit as control. The CNOT acts on the $|10\rangle$ component, flipping the second qubit to create $|11\rangle$.

The final state is the entangled Bell state:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \quad (3.36)$$

In this state, the qubits have lost their individual identities. If we measure the first qubit and find it to be 0, the system instantly collapses to $|00\rangle$, and we know with certainty that the second qubit is also 0. This perfect, non-local correlation is a fundamental resource that has no classical analogue. It lies at the heart of the *spooky action at a distance* that Einstein, Podolsky, and Rosen questioned in their famous EPR paper [22], which initiated the debate on the completeness of quantum mechanics.

It is crucial to note, however, that this correlation does not allow for faster-than-light communication, as the outcome of the first measurement is fundamentally random. This distinction can be illustrated by a thought experiment using classically correlated objects, famously popularised by John Bell [23]. Imagine we have two pairs of gloves, one red and one blue. We choose one pair at random, place each glove into a separate, identical box, and then separate the boxes by a vast distance. We know the gloves in the boxes are perfectly correlated: they are either both red or both blue.

Upon opening one box and finding a red glove, an observer instantly knows that the other box also contains a red glove. This knowledge is acquired instantaneously, yet no information has travelled from one box to the other at that moment. This classical correlation is based on *ignorance*: the state of the gloves was determined the moment they were packed. The measurement simply *reveals* a pre-existing reality, a worldview known as *local realism*.

In the quantum case, the reality is far stranger. Before a measurement is made, a property like the spin of an entangled particle is not simply unknown; it is genuinely *undefined*. The system exists in a true superposition of all possible outcomes. The act of measurement is not a passive revelation but an active process that forces the entire entangled system to collapse into a single, definite state. *Bell's theorem*

provides a mathematical framework to test this distinction, and decades of experiments, notably by Aspect, Clauser, and Zeilinger, have confirmed that quantum systems violate the predictions of local realism, proving that this superposition is real and its collapse is non-local [24, 25].

This ability to generate complex, entangled states through a simple sequence of gates is central to quantum computation. With these foundational concepts established, we now turn to the specific algorithms that leverage these resources for data analysis.

3.5 Elements of Quantum Machine Learning

Quantum Machine Learning originates from a foundational question: can the principles of quantum computing be applied to enhance the performance of classical machine learning algorithms? The term *learning* refers here to the iterative process of updating a model’s parameters to achieve a specific data-driven goal [26]. The central objective is to design quantum software capable of identifying patterns in data more efficiently or effectively than classical methods [27, 28].

Demonstrating a quantum advantage in machine learning remains a complex challenge. Unlike structured problems such as integer factorization, many learning tasks lack known optimal classical algorithms. Claims of advantage must therefore be supported by rigorous benchmarking, which may manifest as improvements in accuracy, generalisation, or the ability to sample from distributions inaccessible to classical systems [29].

Research in this field is currently shaped by the limitations of the NISQ era. The most viable strategy in this setting is the *hybrid quantum-classical model*, which assigns complementary roles to classical and quantum processors. Classical components handle data pre-processing and optimisation, while a Quantum Processing Unit (QPU) executes specific quantum circuits. The most prominent instantiation of this paradigm is the *Variational Quantum Algorithm* (VQA) [30], which forms the operational backbone of the models developed in this thesis.

This section builds the conceptual tools for this paradigm. Section 3.5.1 details the hybrid architecture underlying a VQA. Section 3.5.2 focuses on data encoding via quantum feature maps, which embed classical inputs into the Hilbert space. Section 3.5.3 describes the *parameterised quantum circuit*, the trainable core of variational models. Finally, Section 3.5.5 connects these concepts to recent industrial applications in finance.

3.5.1 The Hybrid Paradigm

Noisy Intermediate-Scale Quantum (NISQ) devices face significant limitations in coherence time and gate fidelity. To circumvent these constraints, the most viable strategy is the *hybrid quantum-classical model*. This approach assigns complementary roles to classical and quantum processors. Classical components handle data pre-processing and optimisation, while the Quantum Processing Unit (QPU) executes specific quantum circuits that are computationally hard for classical devices.

3.5.2 Data Encoding Strategies

Classical data must be mapped into the quantum state space before it can be processed. This step, known as *data encoding* or embedding, is handled by a quantum feature map $U(\mathbf{x})$. The choice of encoding determines the expressivity of the model and the complexity of the circuit.

Basis Encoding. This method associates each classical input with a computational basis state. An integer x (in binary representation) is mapped to the state $|x\rangle$. For example, the input $x = 5$ (binary 101) is encoded as $|101\rangle$. While intuitive,

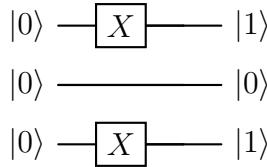


Figure 3.4: Circuit for Basis Encoding of the input $x = 5$ (101_2).

this method is inefficient for continuous data, as it requires discretisation and uses qubits linearly with the bit-precision of the input.

Amplitude Encoding. This strategy encodes a vector $\mathbf{x} = (x_0, \dots, x_{N-1})$ into the amplitudes of a quantum state:

$$|\psi_{\mathbf{x}}\rangle = \sum_{i=0}^{N-1} x_i |i\rangle, \quad (3.37)$$

where the input vector must be normalised such that $\sum |x_i|^2 = 1$. Amplitude encoding is exponentially efficient, storing N features in only $n = \log_2 N$ qubits. However, the state preparation circuit required to load arbitrary amplitudes is generally deep, scaling as $O(N)$, which often negates the advantage on NISQ hardware.

Angle Encoding. This method maps classical features to the rotation angles of quantum gates. For a data vector $\mathbf{x} = (x_1, \dots, x_n)$, the circuit applies rotations such as $R_y(x_i)$ to each qubit. Angle encoding is particularly suitable for financial

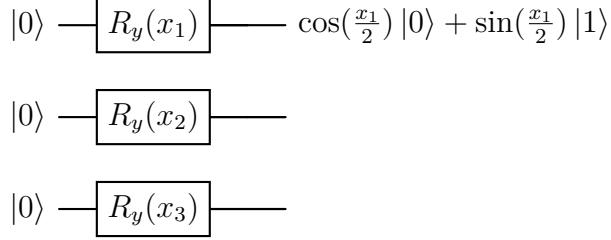


Figure 3.5: Circuit for Angle Encoding. Each feature x_i parameterises a rotation on a separate qubit. The top wire shows the resulting state superposition explicitly.

time series as it naturally handles continuous variables and requires constant circuit depth (one layer of rotations), preserving coherence.

3.5.3 Parameterised Quantum Circuits

Once encoded, the state is processed by a *Parameterised Quantum Circuit* (PQC). A PQC acts as the trainable block of the algorithm, analogous to the weights in a neural network. It consists of fixed entangling gates (to create correlations) and tunable rotation gates (parameterised by a vector θ). The goal of the hybrid loop is to find the optimal parameters θ that minimise a cost function derived from measurement outcomes.

3.5.4 The Quantum Kernel Trick

The most powerful application of PQCs in this thesis is the construction of a *Quantum Kernel*. Classical kernel methods rely on a function to compute similarity in a high-dimensional space without explicit mapping. Quantum computers perform this naturally: the Hilbert space of n qubits has dimension 2^n , offering an exponentially large feature space.

We define the quantum kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ as the fidelity between two encoded data states:

$$K(\mathbf{x}_i, \mathbf{x}_j) = |\langle \psi(\mathbf{x}_i) | \psi(\mathbf{x}_j) \rangle|^2. \quad (3.38)$$

This metric quantifies the similarity between two data points within the quantum feature space. While calculating this inner product is classically hard for deep circuits, it can be estimated efficiently on a quantum computer using interference tests, such as the SWAP test. This kernel forms the mathematical foundation for the statistical tests we employ in Chapter 5.

To understand the power of kernels, consider the conceptual illustration in Figure 3.6. In the original low-dimensional space, complex patterns may not be linearly separable. By mapping the data to a higher-dimensional feature space via a feature map $\phi(\mathbf{x})$, these patterns can become simple hyperplanes. Quantum computers exploit this principle by using the Hilbert space as an exponentially large feature space.

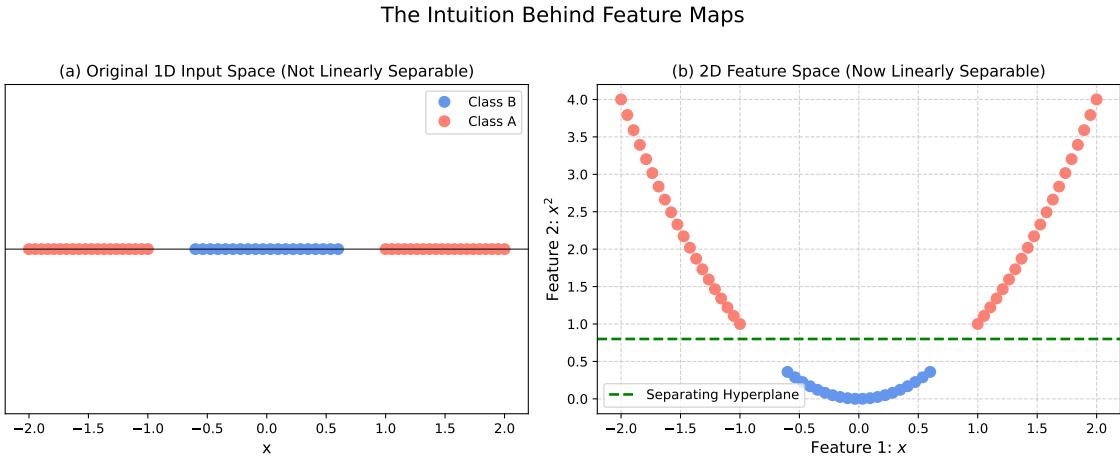


Figure 3.6: A conceptual illustration of the kernel trick. (a) In the original 1D space, the blue and red classes are not linearly separable. (b) After applying a feature map $\phi(x) = (x, x^2)$, the data is projected onto a 2D parabola where the classes become easily separable by a linear hyperplane.

3.5.5 State of the Art: Applications in Finance

The application of QML to real-world industrial problems is a rapidly growing field. Recent research has demonstrated the viability of these techniques in two primary domains relevant to finance: combinatorial optimisation and pattern recognition.

Quantum Optimisation (QUBO). A major line of research tackles large-scale asset allocation problems. While distinct from the circuit-based methods used in this thesis, this approach represents the most mature application of quantum technologies in finance. Recent studies have demonstrated real-world portfolio optimisation using Quadratic Unconstrained Binary Optimisation (QUBO) formulations, achieving results consistent with classical solvers [31]. This framework has been further expanded to include novel formulations for optimising the Sharpe Ratio [32]. Beyond finance, hybrid optimisation has been applied to energy scheduling, yielding improvements in algorithmic fairness [33].

Quantum Kernels and Classification. The second domain, which constitutes the focus of this thesis, uses quantum feature maps for classification and regime detection. The development of benchmarking frameworks has allowed for the rigorous comparison of quantum kernels against classical baselines [34]. Quantum-inspired techniques like Gaussian Boson Sampling have also been applied to detect hidden clusters in complex datasets [35]. Furthermore, practical guides for Amplitude Estimation have provided the necessary groundwork for implementing quantum risk analysis on near-term devices [36].

Our work builds directly upon this second stream of research, utilising the kernel-based approach to address the specific challenge of dynamic regime detection in financial time series.

Chapter 4

Quantum Probability Formulation

Classical probability theory serves as the standard language of quantitative finance. It describes market dynamics using real-valued measures over sample spaces defined by set theory. While this approach has supported the field for decades, Chapter 2 showed that it struggles to capture the full complexity of agent interactions. Specifically, it often fails to explain the emergence of the multimodal return distributions observed during financial crises.

This chapter presents *Quantum Probability* (QP) as a superior mathematical framework for modelling these dynamics. Drawing upon the work of Lin Li [1], we generalise classical probability by replacing real-valued probabilities with complex-valued amplitudes. Although originally developed for microscopic physics, QP provides a rigorous calculus for decision-making under uncertainty, allowing us to model the magnitude and phase of trading intentions simultaneously.

It is crucial to emphasise that QP is employed here as a mathematical framework rather than a physical hypothesis. We do not assert that financial markets adhere to the laws of atomic physics. Instead, we utilise the mathematical properties of QP, specifically the superposition principle and interference, to construct a model where the trading volume drives phase transitions between unimodal and multimodal regimes.

The chapter is organised as follows. Section 4.1 defines the formal structure of Quantum Probability and contrasts it with the classical Kolmogorov framework. Section 4.2 contextualises this choice by comparing it to alternative quantum approaches, specifically the Quantum Harmonic Oscillator model. Section 4.3 introduces the core concept of *Active Trading Intention* and its geometric representation. Finally, Section 4.4 derives the Schrödinger-like trading equation via Fourier decomposition, interpreting the potential function $V(r)$ as the driver of market regimes.

4.1 Quantum Re-interpretation of Classical Probability

Probability theory can be understood as a measure of uncertainty. To apply quantum principles to finance, we must first establish a correspondence between the classical and quantum definitions of this measure.

Definition 4.1.1 (Classical Probability Space). *A classical probability space, as formalised by Kolmogorov [37], is a triple consisting of a sample space Ω , a σ -algebra of events \mathcal{F} , and a probability measure \mathbb{P} .*

- Ω is the set of all possible fundamental and mutually exclusive outcomes of an experiment.
- \mathcal{F} is a collection of subsets of Ω representing all measurable propositions about the outcomes.
- \mathbb{P} is a function assigning a probability in the range $[0, 1]$ to each event.

In classical probability, if an event can occur in two mutually exclusive ways A and B , the total probability is simply the sum $P(A) + P(B)$. In contrast, the quantum framework is built upon the structure of complex Hilbert spaces defined in Chapter 3.

Definition 4.1.2 (Quantum Probability Space). *A quantum probability space is a triple defined by a Hilbert space \mathcal{H} , a set of projection operators \mathcal{E} , and a density operator ρ .*

- \mathcal{H} is a complex Hilbert space representing the state space of the system.
- \mathcal{E} is the set of all orthogonal projection operators on \mathcal{H} . Each projector E corresponds to an event or a measurable proposition.
- ρ is a density operator on \mathcal{H} which describes the state of the system.

This correspondence allows us to translate problems from the classical domain into the quantum domain.

Proposition 4.1.1 (Classical-Quantum Correspondence). *Let $\Omega = \{\omega_i\}$ be a discrete sample space and $\{|e_i\rangle\}$ be an orthonormal basis of a Hilbert space \mathcal{H} . A correspondence can be established as follows:*

- An elementary classical event $\{\omega_i\} \in \mathcal{F}$ is mapped to an elementary quantum event represented by the projection operator $E_i = |e_i\rangle\langle e_i| \in \mathcal{E}$.

- A general classical event $A = \bigcup_{i \in S} \{\omega_i\}$ is mapped to the projection operator $E_A = \sum_{i \in S} |e_i\rangle \langle e_i|$, which projects onto the subspace spanned by the corresponding basis vectors.

Under this mapping, mutually exclusive classical events correspond to mutually orthogonal projection operators.

For a system in a pure state $|\psi\rangle$, the probability of an event E is given by the *Born rule*:

$$P(E) = \text{tr}(E\rho) = \langle\psi|E|\psi\rangle = |\langle e|\psi\rangle|^2. \quad (4.1)$$

The use of complex amplitudes grants QP additional expressive power through the phenomenon of *interference*. According to the superposition principle, if a final state can be reached via two distinct paths represented by amplitudes ψ_1 and ψ_2 , the total state is $\Psi = \psi_1 + \psi_2$. The probability is the square of the sum, not the sum of the squares:

$$P_{tot} = |\Psi|^2 = |\psi_1 + \psi_2|^2 = |\psi_1|^2 + |\psi_2|^2 + 2\Re(\psi_1^* \psi_2). \quad (4.2)$$

The term $2\Re(\psi_1^* \psi_2)$ is the *interference term*. It can be positive or negative, a feature entirely absent in classical probability. This term provides the mathematical mechanism to model how conflicting information signals interact to produce the complex distributions observed in financial markets.

4.2 Alternative Quantum Approaches: The Harmonic Oscillator

Before detailing the specific framework used in this thesis, it is important to distinguish it from other prominent applications of quantum mechanics to finance. A notable alternative approach, proposed by Orrell [38], employs the *Quantum Harmonic Oscillator*, or QHO, to model asset dynamics.

In the QHO framework, the market is modelled using a quadratic potential $V(x) \propto x^2$. Physically, this represents a restoring force, analogous to a mechanical spring, generated by the entropic interplay of supply and demand. In this context, the discrete energy levels of the oscillator are often interpreted as price steps or the number of transactions involved.

While the QHO model has proven highly effective for tasks such as option pricing and reproducing the volatility smile, it typically describes systems tending towards equilibrium. A quadratic potential implies a system that seeks to return to a stable

mean when disturbed. Consequently, the ground state of a QHO is a Gaussian function, which aligns with the unimodal distributions of stable markets.

In contrast, the framework proposed by Li [1], which we adopt in this thesis, leaves the potential $V(r)$ undetermined a priori. This flexibility is essential for our specific objective: capturing the phase transition between regimes. As we will demonstrate, financial crises are characterised by a breakdown of equilibrium and the emergence of multimodality. Modelling this requires a potential that can deform from a simple quadratic shape into a complex multi-well structure, a feature central to Li's derivation.

4.3 The Asset Return Model: Active Trading Intention

Classical financial models describe what happened, namely the return r , but often fail to capture why. A return of zero could result from a day of no trading or a day of fierce battles between buyers and sellers. To capture this hidden dynamic, Li introduces the concept of *Active Trading Intention*.

Definition 4.3.1 (Active Trading Intention). *The Active Trading Intention, or ATI, quantifies the propensity of traders to transact at a certain return level r . It is characterised by two orthogonal components:*

- Intensity: *The magnitude or volume of the trading interest. A higher intensity implies a greater likelihood of transactions occurring at that return level.*
- Property: *The qualitative nature of the intention, such as sentiment or strategic goal.*

The quantum probability framework is uniquely suited to represent the ATI by using a complex probability amplitude, often called the *return wave function*. Here, the modulus $\phi(r)$ represents the intensity, and its square yields the classical probability density $f(r)$. The phase $\theta(r)$ encodes the property of the ATI. Li visualises this phase on the *ATI Plane*, shown in Figure 4.1, a complex plane where the axes correspond to distinct trading motivations.

- *Real Axis:* Represents rational rebalancing or the strategic adjustment of positions. Positive values indicate position accumulation, while negative values indicate reduction.
- *Imaginary Axis:* Represents emotional sentiment. Positive values indicate bearish sentiment, while negative values indicate bullish sentiment.

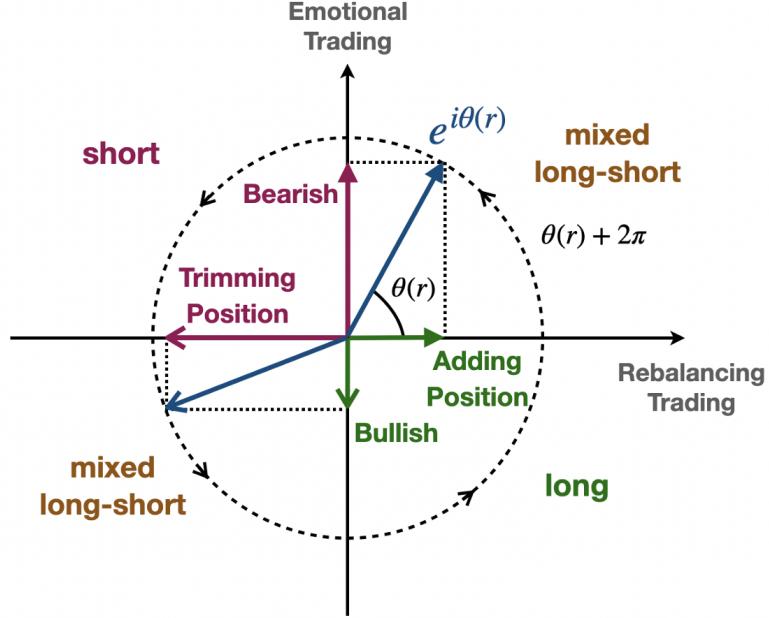


Figure 4.1: The ATI Plane. The phase $\theta(r)$ of the wave function encodes the qualitative nature of trading. The Real axis captures rational position rebalancing, while the Imaginary axis captures emotional sentiment. Adapted from Li [1].

This geometric decomposition allows us to model complex mixed states. For instance, an intention vector in the second quadrant represents behaviour combining short positions with bearish sentiment, which corresponds to panic selling. Conversely, the fourth quadrant represents speculative buying.

The true power of this framework emerges when considering the interaction between multiple traders. Unlike classical probabilities which simply add, quantum amplitudes superpose.

Proposition 4.3.1 (Principle of ATI Aggregation). *If a market consists of N traders, and the ATI of the i -th trader is given by the amplitude $\Psi_i(r)$, the total market ATI is the linear superposition of the individual ATIs:*

$$\Psi(r) = \sum_{i=1}^N \Psi_i(r). \quad (4.3)$$

The resulting probability density contains interference terms involving the cosine of the phase difference between traders:

$$f(r) = |\Psi_1 + \Psi_2|^2 = |\Psi_1|^2 + |\Psi_2|^2 + 2\phi_1\phi_2 \cos(\theta_1 - \theta_2). \quad (4.4)$$

The final term represents the effect of information interaction. It can be positive or negative, providing a formal mechanism to model phenomena like herd behaviour or the no-trade theorem.

4.4 The Schrödinger-like Trading Equation

Having established the representation of ATI through a complex wave function, we now derive the governing equation of the system. The derivation proposed by Lin Li constructs the equation from financial principles by decomposing the total market ATI into components representing different trading behaviours.

4.4.1 Fourier Decomposition: The ω -market

The total market ATI is expressed as a superposition of plane waves via the Fourier inverse transform:

$$\Psi(r) = \int_{-\infty}^{+\infty} c(\omega) e^{i\omega r} d\omega. \quad (4.5)$$

This mathematical decomposition allows us to view the market not as a monolith, but as a continuum of trading strategies. Each frequency component ω corresponds to a conceptual market populated by traders sharing a common behavioural characteristic.

Definition 4.4.1 (The ω -market). *An ω -market is a conceptual market populated by a group of traders, called ω -traders, who share a common behavioural characteristic encoded by the frequency parameter ω . Specifically:*

- The collective ATI of all traders within a given ω -market is represented by the term $c(\omega) e^{i\omega r}$.
- The complex coefficient $c(\omega)$ represents the aggregate intensity and initial phase of that specific group of traders.
- The term $e^{i\omega r}$ describes how the property (phase) of their ATI evolves as the asset return r changes.

This frequency parameter ω offers a direct financial interpretation regarding the time horizon of the traders. A high value of $|\omega|$ corresponds to a short cycle, describing a short-term speculator who reacts quickly to small price changes. Conversely, a low value of $|\omega|$ corresponds to a long cycle, describing a long-term investor whose sentiment evolves slowly and is sensitive only to large shifts in return. This Fourier decomposition effectively separates the fast noise of speculators from the slow signal of investors.

4.4.2 The Governing Equation and the Potential

By combining the Fourier superposition with the continuity condition for probability current, the central result of the framework is derived: a Schrödinger-like equation

that governs the return wave function.

$$i \frac{\partial \Psi(r, t)}{\partial t} = \left(-\frac{1}{2m} \frac{\partial^2}{\partial r^2} + V(r) \right) \Psi(r, t). \quad (4.6)$$

This equation describes the evolution of the market state as a dynamic balance between a kinetic term, representing diffusion of information, and a potential term $V(r)$.

The potential $V(r)$ is the critical variable that determines the shape of the return distribution. It represents the landscape of market consensus (Figure 4.2).

- *Ground State or Calm Regime:* In low-volume conditions, the potential is approximately quadratic. The solution to the Schrödinger equation in a quadratic potential is a Gaussian wave packet. This theoretically explains why classical finance models, such as Black-Scholes, work well in calm markets.
- *Excited States or Turbulent Regime:* As trading volume, which acts as the energy of the market, increases beyond a critical threshold E_0 , the potential deforms. It may develop double wells or complex shapes to accommodate the higher energy. The solution becomes a multimodal wave function, representing a market fractured into competing beliefs.

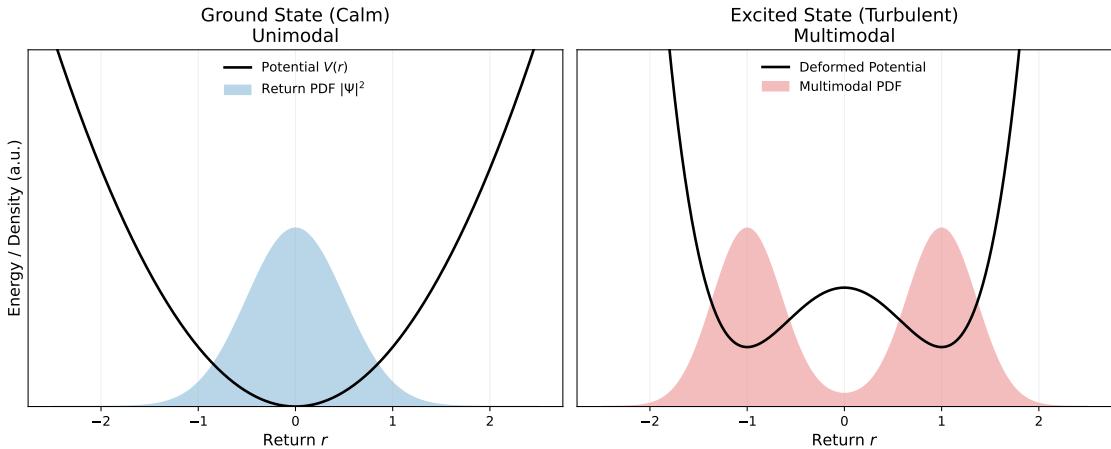


Figure 4.2: Transition of Market Regimes. Left: In the ground state (Calm), the potential is quadratic, confining returns to a unimodal Gaussian distribution. Right: In the excited state (Turbulent), high trading volume deforms the potential into a double-well structure, allowing for a multimodal probability distribution typical of financial crises. The probability density $|\Psi(r)|^2$ (coloured fill) is superimposed on the potential $V(r)$ (black line) and scaled for visual comparison.

This theoretical insight provides the physical justification for the regime-switching methodology we apply in the subsequent chapters. By monitoring the market vol-

ume, we can estimate the energy of the system and detect the phase transition from the ground state to excited states. The specific algorithm for identifying the threshold E_0 using the Hall-York test for multimodality will be detailed in Chapter 5.

Chapter 5

Methodology: The Experimental Pipeline

The theoretical framework established in Chapter 4 posits that financial markets are dynamic entities capable of undergoing phase transitions. Specifically, the Lin Li model predicts that as trading volume surpasses a critical threshold E_0 , the return distribution transitions from a unimodal ground state to a multimodal excited state. This chapter translates that theoretical insight into a concrete, reproducible experimental pipeline.

The methodology is structured in three sequential phases, moving from physical regime identification to quantum statistical validation, and finally to predictive application:

1. *Phase 1: Regime Identification.* We empirically determine the energy threshold E_0 for the SZSE Component Index using the Hall-York statistical test. This data-driven calibration allows us to segment the historical timeline into distinct *Calm* and *Turbulent* regimes based on market activity.
2. *Phase 2: Quantum Statistical Validation.* We implement the Quantum Two-Sample Test to rigorously verify that these heuristically labelled regimes are statistically distinct. By mapping data into a quantum feature space, we aim to detect structural differences that are invisible to classical tests.
3. *Phase 3: Hybrid Forecasting.* We leverage this distinction to construct a regime-switching forecasting model. This hybrid engine dynamically adapts to market conditions, switching between specialised models to optimise predictive performance.

The complete codebase implementing this pipeline is available in the repositories listed in Appendix A.

5.1 Phase 1: Regime Identification

The first objective is to ground the abstract concept of market energy in empirical data. We require a robust, data-driven method to identify the specific volume level E_0 where the market structure fractures.

5.1.1 Data Pre-processing

The primary dataset consists of the *Shenzhen Stock Exchange (SZSE) Component Index* from January 4, 2011, to December 31, 2020. Data were sourced from publicly available financial databases [39]. Consistent with the energy derivation in Chapter 4, we focus on the *intraday log-return* as the primary variable for modality testing:

$$r_t^{intra} = \ln(\text{Close}_t) - \ln(\text{Open}_t). \quad (5.1)$$

This metric isolates the specific trading activity of a single session, removing the overnight gap which acts as external noise. The corresponding daily trading volume V_t serves as the proxy for the kinetic energy of the system.

5.1.2 The Statistical Engine: Hall-York Test

Visual inspection is insufficient for formally detecting the transition from unimodality to multimodality. Rigorous detection requires a statistical test capable of quantifying the probability that a given distribution has more than one peak. While Silverman's original bandwidth test [40] is a foundational method for this purpose, it is known to be conservative and prone to false negatives. To address this limitation, we employ the *Hall-York Test* of 2001, a calibrated enhancement of Silverman's method [41].

The test operates on the principle of the *Critical Bandwidth*, denoted as h_{crit} . In Kernel Density Estimation (KDE), the bandwidth h controls the smoothness of the distribution. A large h produces a single smooth peak, while a small h reveals fine-grained noise. The critical bandwidth is defined as the smallest smoothing parameter such that the resulting density estimate has exactly one mode. Figure 5.1 illustrates a density estimate computed at this critical bandwidth for our high-energy regime; the emergence of two distinct peaks is the visual signature of the multimodality we aim to detect.

The Hall-York algorithm proceeds as follows:

1. Compute h_{crit} for the original data sample.
2. Generate B bootstrap samples from the density estimate computed at h_{crit} .

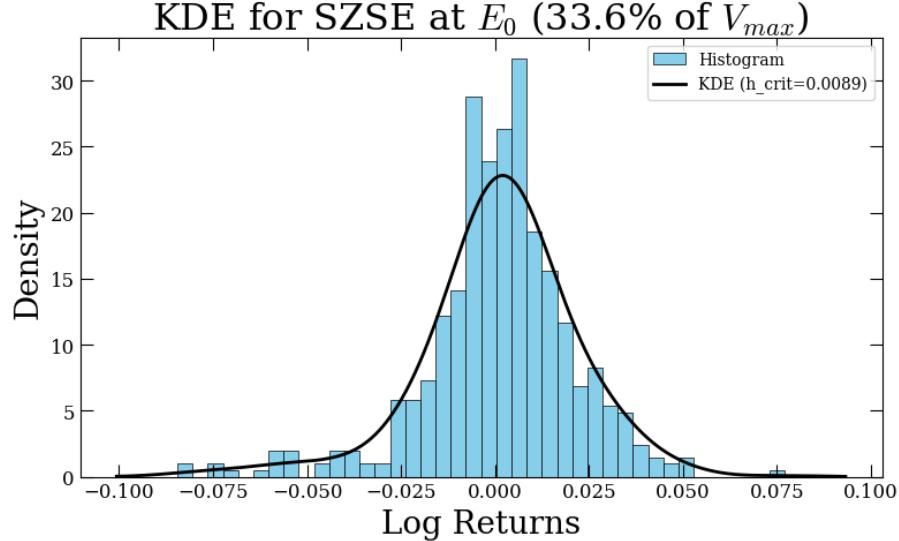


Figure 5.1: Kernel Density Estimate of intraday log-returns for the high-energy regime ($\text{Volume} \geq E_0$). The density exhibits clear multimodality, with two distinct peaks emerging at the critical bandwidth $h_{crit} = 0.0089$ identified by the Hall-York test. This confirms the presence of a phase transition, validating the regime shift predicted by the theoretical framework.

In our sensitivity analysis, we varied B to ensure convergence, typically employing $B \geq 100$.

3. Calculate the critical bandwidth h_{crit}^* for each bootstrap sample.
4. The *p-value* is the proportion of bootstrap samples where $h_{crit}^* < h_{crit}$.

A low p-value, typically $p < 0.05$, indicates that the original data requires a significantly larger amount of smoothing to appear unimodal than would be expected by chance, leading to the rejection of the null hypothesis of unimodality.

5.1.3 The Search for the Threshold E_0

Equipped with the Hall-York test, we implement an iterative search algorithm to locate the critical energy threshold E_0 . The procedure scans the volume space to find the breaking point of the market structure:

1. *Initialization:* We define a candidate volume threshold v_{star} , starting at a fraction of the minimum observed volume (e.g., 5% of the range).
2. *Filtering:* We construct a subsample of trading days where the daily volume $V_t \geq v_{star}$. This represents the high-energy subset of the market history.
3. *Testing:* We apply the Hall-York test to the intraday returns of this subset.

4. *Iteration:* If the null hypothesis of unimodality cannot be rejected ($p \geq 0.05$), we increment v_{star} by a step δ and repeat the process. The algorithm terminates when the p-value drops below the significance level $\alpha = 0.05$, indicating significant multimodality.

For the SZSE dataset (2011–2020), our analysis converged to a threshold of:

$$E_0 \approx 1.219 \times 10^{10} \text{ CNY (Chinese Yuan).} \quad (5.2)$$

5.1.4 Regime Labelling and Persistence

The identification of E_0 allows us to segment the time series into two distinct regimes based on market activity:

- *Calm Regime (Ground State):* Days where $V_t < E_0$. The market is dominated by low-frequency investors, and returns follow a unimodal distribution.
- *Turbulent Regime (Excited State):* Days where $V_t \geq E_0$. The market energy is sufficient to deform the potential, leading to multimodal distributions characteristic of conflicting trader beliefs.

To validate the temporal consistency of these regimes, we analysed specific market windows identified as multimodal. Figure 5.2 visualises one such window during the 2015 crisis.

The plot overlays two distinct metrics on a dual axis. The blue line, corresponding to the left axis, tracks the intraday log-returns and exhibits high volatility clustering. The red line on the right axis tracks the trading volume normalised by the threshold E_0 . The horizontal dashed line at $y = 1$ marks the critical energy transition. The persistence of volume above this threshold corresponds to a cluster of high volatility. This confirms that the transition to an excited state is a robust, persistent phenomenon rather than a momentary outlier.

However, identifying these regimes heuristically is not enough; we must now rigorously prove that they represent statistically distinct probability distributions.

5.2 Phase 2: Quantum Statistical Validation

Identifying potential regimes via the Hall-York test provides a necessary but insufficient condition for validation. Classical univariate tests often fail to distinguish between complex financial time series that share similar summary statistics, such as mean and variance, but differ fundamentally in their temporal structure. To prove

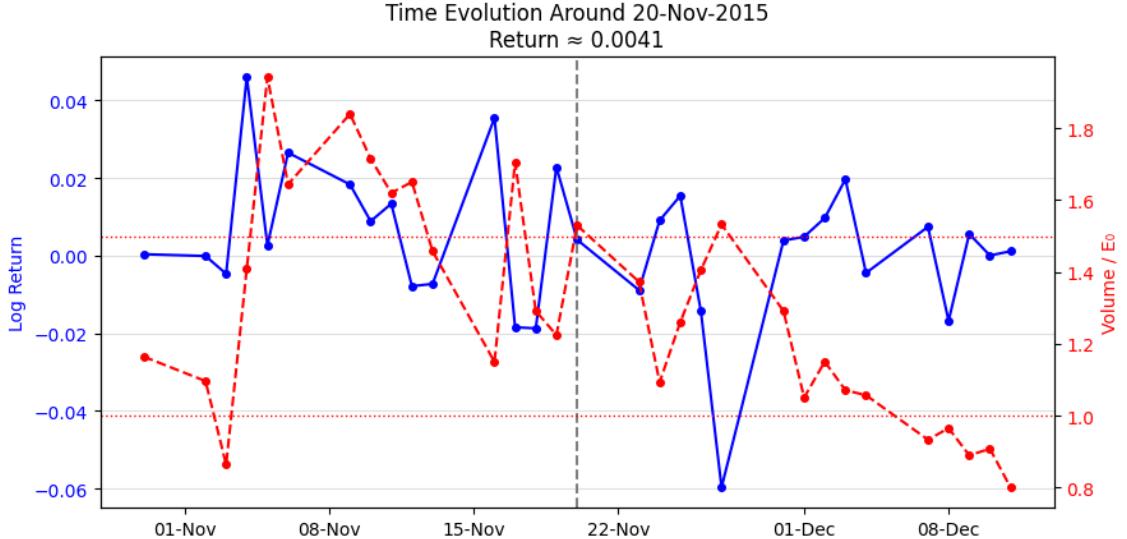


Figure 5.2: Time evolution of a high-energy market window from November 2015. The plot overlays two metrics: the intraday log-returns (blue line, left axis) and the trading volume normalised by the threshold E_0 (red line, right axis). The red trace remains consistently above the critical threshold at $y = 1$, confirming that the market was locked in a high-energy *excited state* characterised by persistent instability.

rigorously that the identified regimes are distinct, we implement a multivariate approach based on the Quantum Maximum Mean Discrepancy (MMD) framework proposed by Garvin et al. [2].

5.2.1 Synthetic Validation Model

To test our implementation in a controlled environment, we generate synthetic financial time series using a model that captures key features of real-world asset prices. The model is based on a stochastic differential equation that includes a drift term, a volatility term, and an autocorrelation term. The discrete version of this model, as implemented in our study, is given by:

$$S_t = S_{t-1} \exp \left\{ \left(\mu - \frac{1}{2} \sigma^2 \right) \Delta t + \sigma \sqrt{\Delta t} \left(\varrho z_{t-1} + \sqrt{1 - \varrho^2} z_t \right) \right\} \quad (5.3)$$

where:

- S_t is the asset price at time t .
- μ is the annualised drift, representing the long-term expected return. A positive drift implies growth, while negative implies decline.
- σ is the annualised volatility, measuring the magnitude of price fluctuations (risk).

- ϱ is the autocorrelation parameter, controlling the serial dependence. A value of $\varrho \approx 1$ implies momentum (trends), while $\varrho \approx 0$ implies a random walk consistent with the Efficient Market Hypothesis.
- $z_{t-1}, z_t \sim \mathcal{N}(0, 1)$ are random variables drawn from a standard normal distribution.
- Δt is the time step, typically set to 1/252 for daily trading data.

This model allows us to generate time series with distinct temporal structures by varying ϱ . Crucially, while the generative process differs, the resulting univariate distributions of daily log-returns have identical first and second moments (mean and variance) for all values of ϱ . This ensures that any detected difference by a multivariate test is due to higher-order moments or temporal correlations, providing a rigorous benchmark for the quantum kernel's sensitivity. We generated three

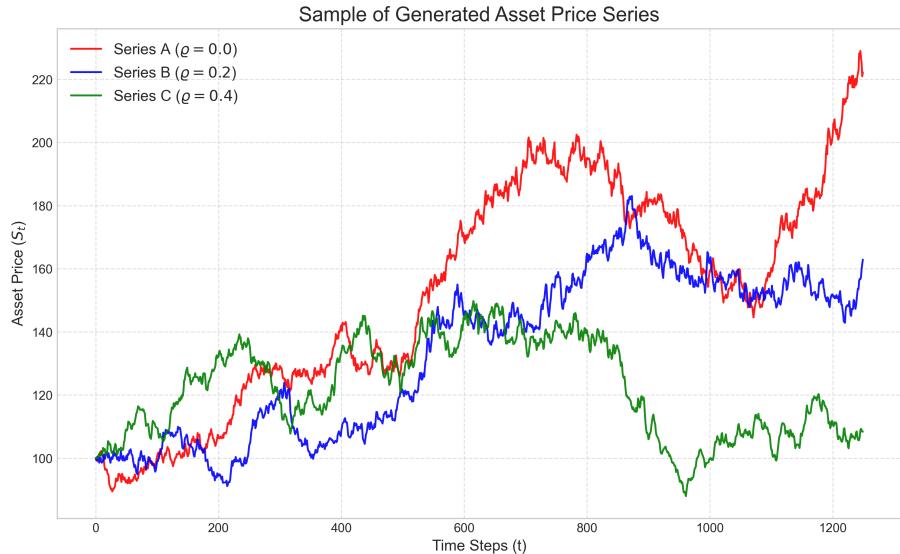


Figure 5.3: Sample of Generated Asset Price Series. Series A ($\varrho = 0$) behaves as a random walk, while Series C ($\varrho = 0.4$) exhibits clear momentum trends.

datasets: Series A with $\varrho = 0.0$ (random walk), Series B with $\varrho = 0.2$, and Series C with $\varrho = 0.4$ (momentum). We confirmed this indistinguishability by overlaying the histograms of log-returns for Series A and C in Figure 5.4. A Kolmogorov-Smirnov test yielded a high p-value of 0.61, failing to reject the null hypothesis. This result demonstrates that classical univariate tests cannot distinguish between these series, confirming the need for a quantum test sensitive to temporal correlations.

5.2.2 Feature Engineering: Capturing Temporal Dynamics

To capture the serial dependence missed by univariate tests, we transform the scalar time series into a sequence of feature vectors. Following Garvin et al., we construct

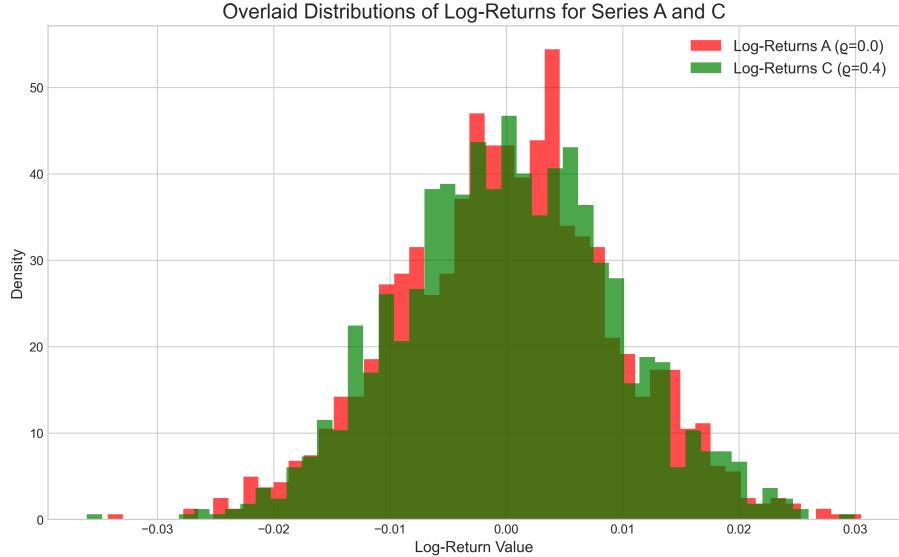


Figure 5.4: Overlaid Distributions of Log-Returns for Series A ($\rho = 0$) and C ($\rho = 0.4$). The distributions are nearly identical, as confirmed by the high p-value (0.61) of the KS test.

a 4-dimensional feature vector \mathbf{x}_t for each time step t , specifically designed to encode correlations between consecutive returns:

$$\mathbf{x}_t = \left[\ln \frac{S_t}{S_{t-1}}, \quad \ln \frac{S_{t-1}}{S_{t-2}}, \quad \left| \ln \frac{S_t}{S_{t-1}} + \ln \frac{S_{t-1}}{S_{t-2}} \right|, \quad \left| \ln \frac{S_t}{S_{t-1}} - \ln \frac{S_{t-1}}{S_{t-2}} \right| \right]. \quad (5.4)$$

These components correspond to the current return, the previous return, the absolute magnitude of their sum (momentum indicator), and the absolute magnitude of their difference (mean reversion indicator). This transformation converts the problem into distinguishing two clouds of points in a 4-dimensional space, which is suitable for processing by a 4-qubit quantum circuit.

5.2.3 The Classical Two-Sample Test (MMD)

As a baseline for our study, we employ the classical Maximum Mean Discrepancy (MMD) test [42]. This test provides a robust statistical framework for determining whether two samples, $P = \{\mathbf{x}_i\}_{i=1}^m$ and $Q = \{\mathbf{y}_j\}_{j=1}^n$, are drawn from different distributions.

The core idea is to map the data into a high-dimensional Reproducing Kernel Hilbert Space (RKHS), denoted as \mathcal{F} , via a feature map $\phi(\cdot)$. In this space, the probability distributions are represented by their mean embeddings, μ_P and μ_Q . The MMD is defined as the distance between these means:

$$\text{MMD}(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{F}}. \quad (5.5)$$

If the RKHS is characteristic, then $\text{MMD}(P, Q) = 0$ if and only if $P = Q$.

Calculating the coordinates in \mathcal{F} explicitly is often computationally prohibitive. However, by applying the *kernel trick*, we can compute the squared MMD solely using the kernel function $k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle$. The empirical estimator for the squared MMD is given by:

$$\text{MMD}^2(P, Q) = \frac{1}{m^2} \sum_{i,j=1}^m k(\mathbf{x}_i, \mathbf{x}_j) - \frac{2}{mn} \sum_{i=1, j=1}^{m,n} k(\mathbf{x}_i, \mathbf{y}_j) + \frac{1}{n^2} \sum_{i,j=1}^n k(\mathbf{y}_i, \mathbf{y}_j). \quad (5.6)$$

Intuitively, this formula compares the average intra-sample similarity (how similar data points are to others in the same set) with the inter-sample similarity (how similar they are to points in the other set). If the samples are distinct, the intra-sample similarity will significantly exceed the inter-sample similarity.

For our classical benchmark, we use the standard Gaussian Radial Basis Function (RBF) kernel:

$$k_{\text{RBF}}(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\sigma^2}\right), \quad (5.7)$$

where σ is a hyperparameter controlling the width of the kernel, typically set to the median heuristic of the data distances.

5.2.4 The Quantum Two-Sample Test

The quantum two-sample test is an analogue of the classical MMD test that leverages the principles of quantum mechanics to potentially achieve more powerful discrimination. The process involves four key steps:

1. *Data Encoding*: The classical data vector \mathbf{x} is encoded into the parameters of a quantum circuit. This crucial step translates classical information into a quantum state in the Hilbert space.
2. *Quantum Circuit Evolution*: A Parametric Quantum Circuit is applied to the qubits. This sequence of operations evolves the state based on the input data.
3. *Measurement*: We measure the expectation values of the system, providing a classical description of the final quantum state.
4. *Distance Calculation*: These values are used to compute the quantum kernel, defined as the fidelity (squared overlap) between states:

$$k_Q(\mathbf{x}_i, \mathbf{x}_j) = |\langle \psi(\mathbf{x}_i) | \psi(\mathbf{x}_j) \rangle|^2. \quad (5.8)$$

This kernel is then plugged into the standard MMD formula to compute the test statistic.

5.2.5 Parametric Quantum Circuit Architectures

The expressiveness of the quantum feature map depends on the structure of the Parametric Quantum Circuit, or PQC, which encodes the classical data into the quantum state (see Section 3.5.3). A PQC is a quantum circuit containing a set of tunable rotation gates, whose angles act as the trainable parameters of the model. We benchmarked four distinct architectures: the *Paper Circuit*, the *TFIM Circuit*, the *LTFIM Circuit*, and a *Hardware-Efficient Circuit*. Detailed diagrams and specifications for each architecture are provided in Appendix B. As detailed in the results, the *Paper Circuit* (Figure 5.5) demonstrated superior expressivity for temporal correlations and was selected for the final analysis.

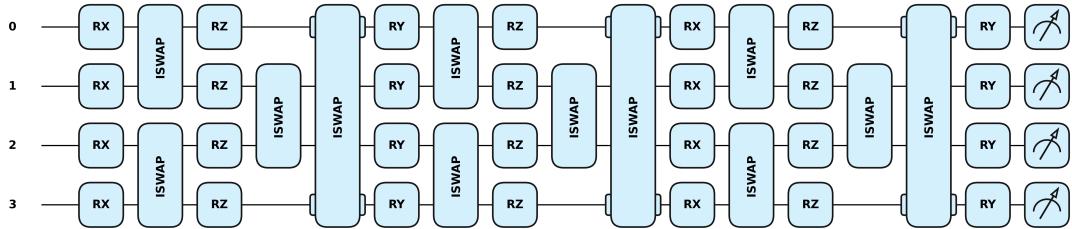


Figure 5.5: The Paper Circuit architecture (Garvin et al.). It features alternating layers of single-qubit rotations for data encoding and iSWAP gates for entanglement. This design proved most robust in our synthetic benchmarks.

5.2.6 Circuit Selection: Synthetic Benchmarking

Before applying the quantum pipeline to real market data, it is essential to validate the discriminatory power of the chosen ansatz in a controlled environment. We benchmarked the *Paper Circuit* against alternative architectures (detailed in Appendix B) using the synthetic dataset described in Section 5.2.1.

Figure 5.6 presents the results of this validation. The test compares a *Random Walk* process ($\varrho = 0$, Series A) against a *Momentum* process ($\varrho = 0.4$, Series C). The quantum kernel constructed with the Paper Circuit successfully distinguishes the two distinct processes, yielding a statistical distance (red bar) that is clearly separated from the baseline noise (blue/green bars), with a Signal-to-Noise Ratio (SNR) of approximately 1.7. This result confirms that the Paper Circuit possesses the necessary expressivity to detect temporal correlations in financial time series, justifying its selection for the analysis of real-world data in the following chapter.

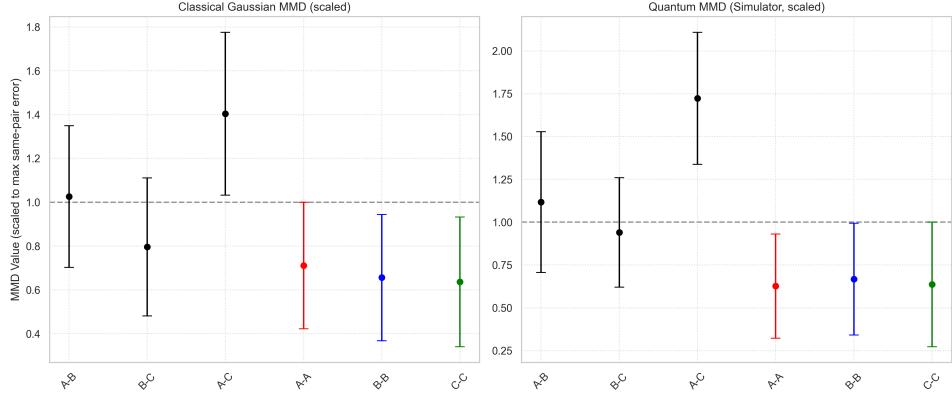


Figure 5.6: Benchmark results for the *Paper Circuit* on synthetic data. The test successfully distinguishes the distinct distributions (A vs C, red bar) from the baseline noise (blue and green bars), with the error bar lying entirely above the noise threshold ($y = 1.0$).

5.2.7 Data Sampling and Validation Strategy

The input features for the quantum two-sample test are log-returns derived from raw price data. A key advantage of the MMD framework is its flexibility: unlike element-wise metrics, it allows for the comparison of samples with unequal sizes ($N \neq M$). This capability is explicitly leveraged in our analysis to compare market regimes of varying duration (e.g., annual windows vs. shorter crisis bursts), verifying the robustness of the quantum kernel across different temporal scales.

Our validation logic follows a hierarchical approach:

1. *Baseline Check*: We first compare two non-overlapping samples drawn from the same stable, unimodal market period. The expected result is an MMD score statistically consistent with zero.
2. *Detection Test*: We then compare a sample from a stable period against a sample from a high-volatility period to test for statistical separability.

5.2.8 Bootstrap Resampling

A single MMD score calculated from a small sample of historical data is statistically noisy. To obtain a robust estimate of the distance and its uncertainty, we employ a *bootstrap resampling* methodology. Instead of a single comparison, we perform $B = 100$ independent test runs for each pair. In each run, we draw smaller random sub-samples (with replacement) from the original time windows. The final reported MMD is the mean of this distribution, and the error bars represent its standard deviation.

5.3 Phase 3: Hybrid Forecasting Strategy

The final phase of our pipeline evaluates the practical utility of the identified market regimes. If the *Calm* and *Turbulent* states are indeed structurally distinct, as indicated by the Quantum Two-Sample Test, then a forecasting model specialised on a specific regime should theoretically outperform, or at least match, a generalist model, even when trained on significantly less data.

To test this hypothesis, we implemented a *Walk-Forward Validation* experiment focused on the 2015 market crash.

Handling Regime Flickering

Our analysis of the regime blocks (see Appendix C) reveals that prolonged crisis periods are often interrupted by single days of lower volume, leading to a unimodal classification. For instance, the 2015 crash consists of several major multimodal blocks (Blocks 2, 4, 6, 8) separated by brief intervals of 1-2 days where volume dipped slightly below E_0 .

From a physical perspective, the market system does not instantly *cool down* to a stable ground state for 24 hours during a crash. These interruptions are likely artefacts of using a static threshold E_0 on a noisy variable like volume. Consequently, for the purpose of training the forecasting model, we adopt a *merging strategy*: adjacent turbulent islands separated by negligible gaps are treated as a single contiguous *Crisis* training set. This prevents the model from losing critical memory of the crisis dynamics due to transient noise.

5.3.1 Model Specification

We constructed two competing forecasting engines using the `AutoARIMA` framework combined with GARCH volatility modelling:

1. *The Baseline Model (Generalist)*: This model is trained on the full available history up to the test point. It represents the standard econometric approach of using maximum historical depth to estimate parameters. For instance, in the 2015 walk-forward validation detailed in Chapter 6, this training set comprised 1135 trading days.
2. *The Crisis Model (Specialised)*: This model is trained *exclusively* on the high-energy *islands* of volatility identified by the threshold E_0 in the period preceding the test. For the 2015 experiment, this training set consisted of only 76 trading days formed by merging two adjacent multimodal windows. This

represents the regime-switching approach: discarding *calm* data as noise when predicting a *crisis* event.

5.3.2 Validation Protocol: Walk-Forward vs Cross-Validation

Standard cross-validation (k-fold) is unsuitable for time series because it randomly shuffles data, allowing the model to train on future data to predict the past (data leakage). Instead, we implemented a strict *Walk-Forward Validation*, also known as a rolling-origin evaluation.

The process works as follows:

1. The model is trained on a historical window $t_0 \rightarrow t_k$.
2. It forecasts the next step t_{k+1} .
3. The window slides forward: the actual value at t_{k+1} is added to the training set, and the model is retrained (or updated) to predict t_{k+2} .

This simulation perfectly mimics the real-world constraints of a trader who receives data sequentially, ensuring that at no point does the model access future information.

5.3.3 Validation Metric

The performance is evaluated using two complementary metrics. The *Root Mean Square Error* (RMSE) measures the accuracy of the point forecast (the expected return):

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{t=1}^N (y_t - \hat{y}_t)^2}. \quad (5.9)$$

The *Volatility MSE* measures the accuracy of the risk prediction, comparing the forecasted variance $\hat{\sigma}_t^2$ against the squared return (a proxy for realized volatility):

$$\text{Vol-MSE} = \frac{1}{N} \sum_{t=1}^N (y_t^2 - \hat{\sigma}_t^2)^2. \quad (5.10)$$

5.3.4 Generalisation and Model Robustness

Finally, to assess the stability of market regimes over long time horizons, we perform a *Generalisation Test*. A model trained on a specific historical crisis is tested against a subsequent, distinct crisis period occurring several years later. This test determines whether the statistical properties of *Turbulence* are universal and static, or if they evolve over time. This phenomenon, known as *concept drift*, would imply that a model trained on a past crisis may not be optimal for a future one, necessitating a dynamic retraining strategy where the model is updated as new regimes are detected.

5.4 Implementation and Scalability Notes

This section details the computational implementation of the experimental pipeline. We justify the choice of exact state-vector simulation over noisy quantum hardware and clarify the implications of this approach regarding algorithmic scalability and performance.

5.4.1 Simulation Strategy

The calculation of the quantum kernel $k_Q(\mathbf{x}_i, \mathbf{x}_j)$ is designed to be performed on a quantum computer. However, to robustly validate the effectiveness of the *method* separately from the limitations of current hardware (noise and decoherence), this project utilises a high-performance classical simulation. We employed the `PennyLane` framework with the `lightning.qubit` backend (C++), which allows for exact state-vector simulation.

5.4.2 Scalability Considerations

It is important to clarify the scalability implications of this choice.

- *Quantum Scalability*: On real quantum hardware, the kernel estimation can be performed efficiently, scaling linearly with the number of samples $O(N)$.
- *Simulated Scalability*: Simulating the quantum kernel on a classical computer incurs an exponential cost in the number of qubits. Therefore, the runtime of our experiments scales poorly compared to classical kernels.

This thesis does not aim to demonstrate a computational speedup, which is impossible in simulation. Instead, the objective is to demonstrate a *quantum advantage* in *discriminatory power*: proving that the quantum model is “smarter” (more sensitive) than the classical baseline, regardless of the execution time.

Chapter 6

Results and Discussion

This chapter presents the empirical findings of the thesis. Building on the experimental pipeline defined in Chapter 5, we demonstrate that the quantum probability framework provides not only a theoretical description of market regimes but a measurable advantage in detection and forecasting.

The chapter is organised to mirror the methodological phases:

- Section 6.1 visualises the identified market regimes over the full decade of analysis, validating the physical meaningfulness of the energy threshold E_0 against historical market events.
- Section 6.2 presents the statistical validation. We first benchmark the quantum circuits on synthetic data and then apply the Quantum Two-Sample Test to real-world windows, demonstrating its superior sensitivity compared to its classical counterpart.
- Section 6.3 evaluates the practical utility of these findings. We discuss the results of a walk-forward forecasting experiment, focusing on the informational efficiency of regime-specialised models.

6.1 Phase 1: Regime Identification (SZSE 2011-2020)

The application of the iterative threshold search described in Section 5.1 successfully located the critical energy level E_0 for the SZSE Component Index. The analysis consistently converged to a stable threshold of:

$$E_0 \approx 1.219 \times 10^{10} \text{ CNY.} \quad (6.1)$$

This value is consistent with the findings reported by Lin Li [1] for the same index, confirming the correct implementation of the detection pipeline.

6.1.1 Macro-Analysis and Historical Validation

Figure 6.1 visualises the resulting segmentation of the 10-year timeline. The plot combines three perspectives to validate the regimes: price action, trading volume relative to E_0 , and a standard volatility metric. The 21-day rolling volatility, showed in the bottom panel, is computed as the annualised standard deviation of log-returns over a sliding window, serving as a classical proxy for risk.

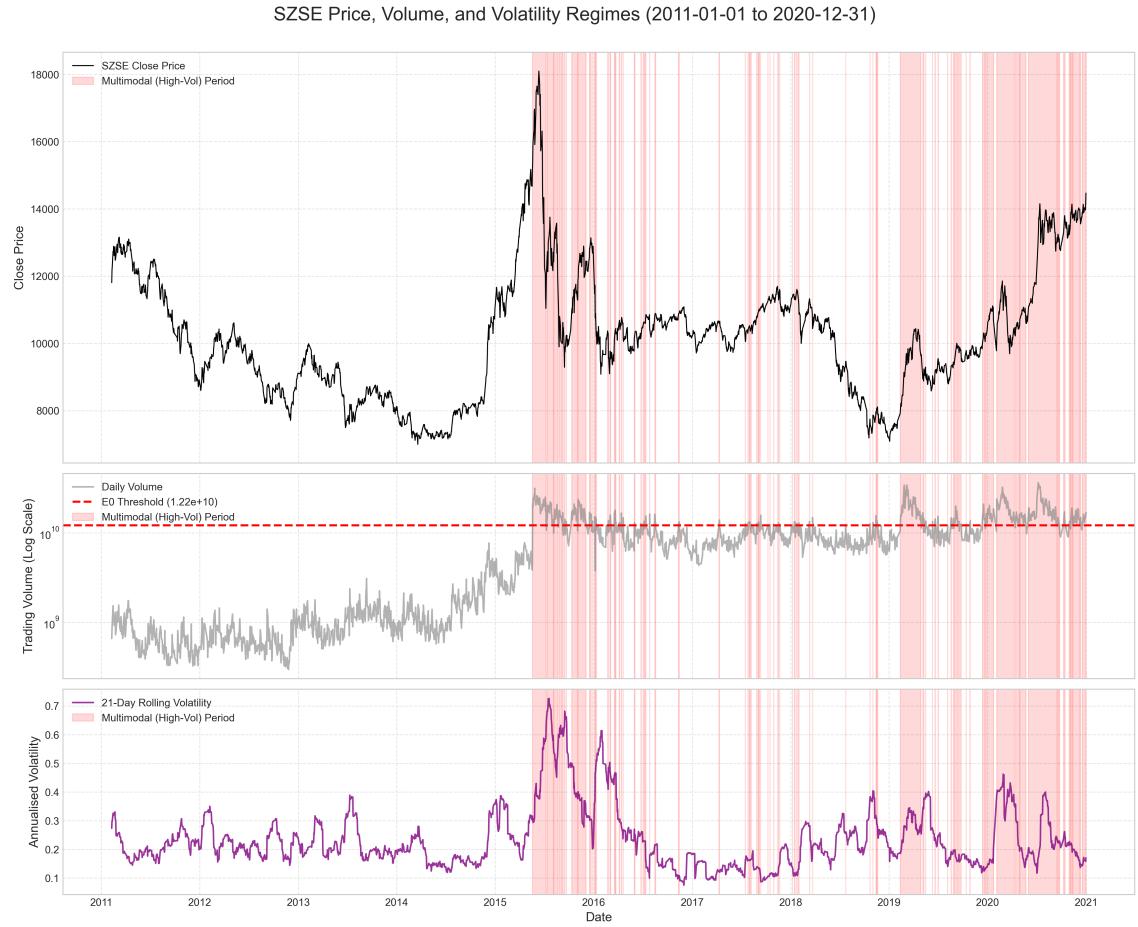


Figure 6.1: Macro-Analysis of Market Regimes. *Top Panel:* The closing price of the SZSE index. Red shaded regions indicate the Turbulent regimes. These regimes consistently align with major market corrections. *Middle Panel:* Daily trading volume relative to the threshold E_0 (dashed red line). *Bottom Panel:* The 21-day rolling volatility. The strong correlation between the identified regimes and volatility spikes confirms that the volume-based energy metric serves as a reliable proxy for market risk.

The distribution of turbulent regimes, highlighted in red, aligns with major macroeconomic shocks affecting the Chinese market:

- **2011–2014:** The market remains largely in the calm regime (Block 1 in Appendix C). This long period of stability validates the existence of a robust ground state where prices evolve without sufficient energy to fracture the return distribution.
- **2015:** A dense cluster of turbulent days appears mid-year, corresponding to the bursting of the 2015 Chinese stock market bubble [43].
- **2019–2020:** A prolonged period of instability is visible, culminating in the market volatility associated with the COVID-19 pandemic in early 2020.

The visual correlation between the red zones ($\text{Volume} > E_0$) and the spikes in the rolling volatility metric confirms that the volume-based energy metric successfully captures periods of elevated market risk.

6.1.2 Regime Dynamics and Persistence

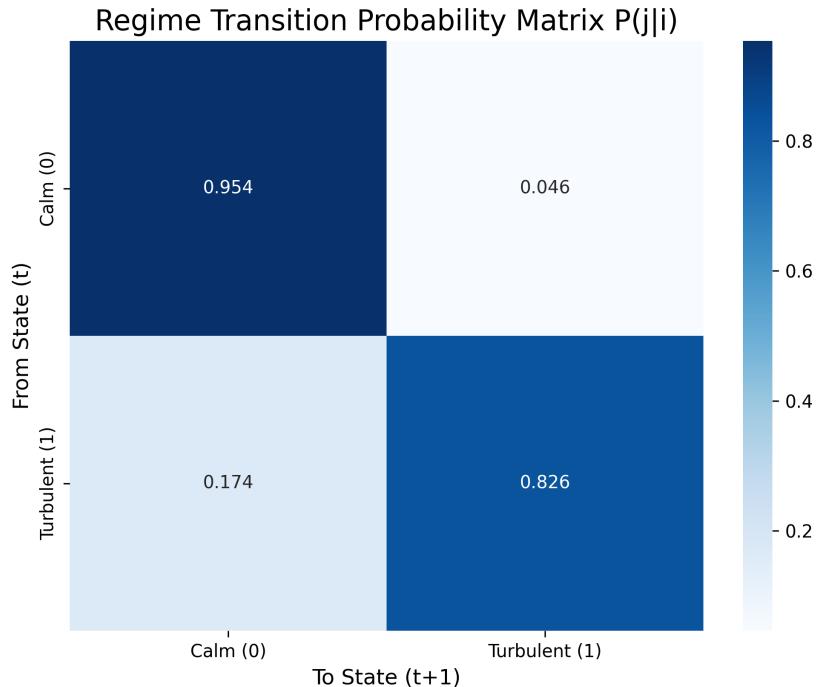


Figure 6.2: Regime Transition Probability Matrix. The high values on the diagonal (95.4% and 82.6%) indicate that market regimes are persistent states. Once the market enters a turbulent phase, there is an 82.6% probability it will remain turbulent the next day.

To confirm that these regimes represent stable market states rather than transient noise, we analysed the transition probabilities between daily states. Figure 6.2

presents the transition matrix as a heatmap. The probability of remaining in a *Calm* state is 95.4%, while the probability of remaining in a *Turbulent* state is 82.6%.

This contrasts sharply with a memoryless random process. If regime switching were independent of history, the probability of observing a turbulent day tomorrow would simply equal the long-term prevalence of the regime ($\approx 21\%$). The fact that the conditional probability (83%) is four times higher than the unconditional probability confirms that volatility clusters in time. This *stickiness* is the fundamental property that validates the use of regime-switching models: since the market state persists, identifying the current regime provides predictive information about the future.

6.2 Phase 2: Quantum Statistical Validation

Having identified the candidate regimes based on the volume threshold E_0 , we now evaluate the sensitivity of the Quantum Two-Sample Test in distinguishing between them. The validation follows a hierarchical logic:

1. *Baseline Stability (Sanity Check)*: We first verify that the test does not produce false positives when comparing distinct periods that belong to the same regime.
2. *Regime Discrimination (Detection)*: We then test whether the quantum circuit can successfully distinguish between a *Calm* period and a *Turbulent* period, quantifying the quantum advantage over the classical benchmark.

All tests were performed using the *Paper Circuit* architecture, which demonstrated the highest robustness in the synthetic benchmarks (see Appendix B).

6.2.1 Baseline Stability Analysis

The first step is to ensure that the quantum test is stable over time within the ground state. The first unimodal block identified by our pipeline (Block 1, see Appendix C) presents a unique challenge: it spans 1060 trading days (2011–2015). Comparing such a long window directly against shorter blocks introduces statistical artefacts due to the disparity in sample size.

To address this, we segmented this long stable period into three annual windows: 2011, 2012, and 2013. Since these years all fall well below the critical threshold E_0 , they should theoretically represent the same statistical manifold. A valid test must therefore yield a quantum MMD distance close to zero.

Figure 6.3 presents the results of this intra-regime comparison. The analysis reveals two key insights:

- *High Self-Distance Variance:* The error bars for the self-distances (the coloured bars representing A vs A') are relatively large. This indicates that even within a single year, the market exhibits a degree of natural variability or noise.
- *Indistinguishability:* Crucially, the cross-distance bars (e.g., 2011 vs 2012, , shown in black) are statistically indistinguishable from the self-distance bars (baseline noise). The mean MMD scores for the cross-pairs lie comfortably below the $y = 1.0$ noise threshold ratio.

This result confirms the stability of the method: despite the temporal separation of entire years, the quantum feature map correctly identifies the underlying market structure as unchanged.

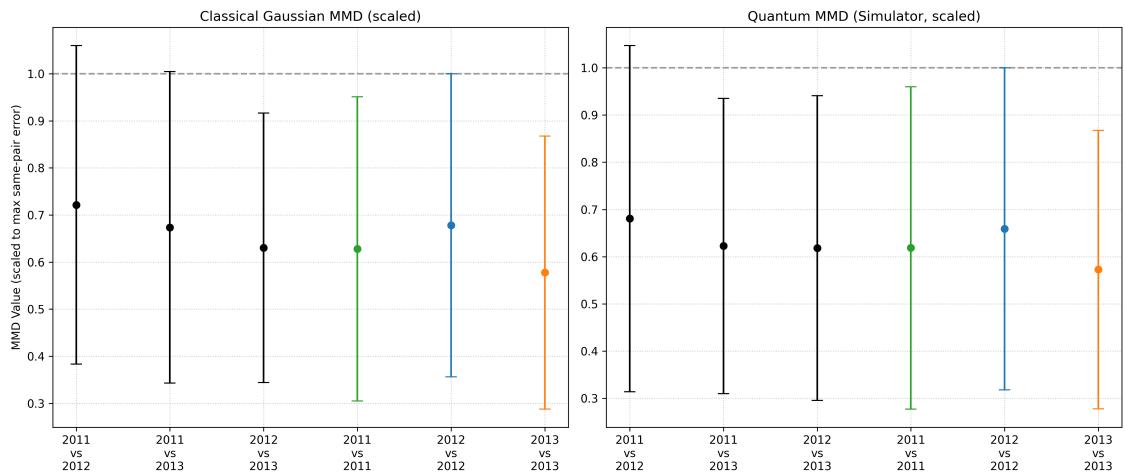


Figure 6.3: Baseline Stability Test (2011–2013). The plot compares the Quantum MMD distances between three consecutive years classified as Calm. The cross-distances (black bars) overlap significantly with the self-distances (coloured bars) and remain below the noise threshold. (For the detailed price and volume evolution of these annual windows, see Figure D.1 in Appendix D).

6.2.2 Robustness to Time and Sample Size

Having confirmed stability within the initial long block (2011–2013), we extended the validation to later market periods. We selected three distinct unimodal “islands” identified by the Hall–York algorithm as the most significant stable periods post-2015:

- *Block 63 (2016-2017):* Duration of 96 trading days.
- *Block 65 (2017):* Duration of 64 trading days.
- *Block 109 (2018):* Duration of 81 trading days.

This selection introduces two stress factors: it compares regimes separated by multiple years (testing for concept drift in the ground state) and compares samples of unequal length (testing for sample size bias).

Figure 6.4 displays the results. Despite the temporal distance and length disparity, the quantum MMD distances between these blocks remain low. Specifically, the cross-distances (e.g., Block 63 vs Block 109) are comparable to the self-distances. This confirms that the Calm regime retains a consistent statistical structure over the years, further validating it as a stable ground state regardless of the specific window size used.

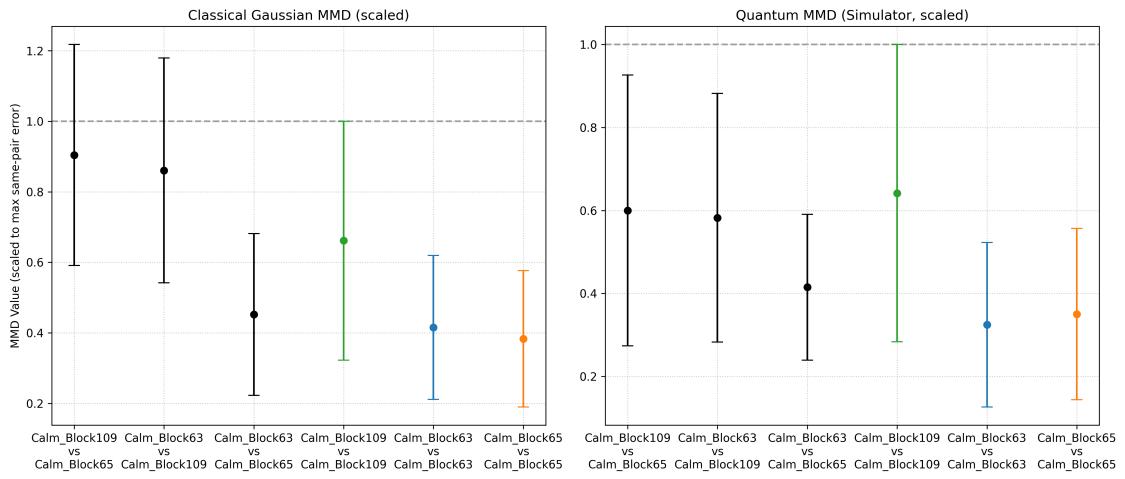


Figure 6.4: Robustness Check on Unimodal Blocks (2016–2018). The test compares three distinct calm periods of varying lengths. The low cross-distances (black bars) confirm that the unimodal regime is structurally stable over time and robust to sample size variations.

6.2.3 Regime Discrimination Test (Detection)

Finally, we evaluate the core capability of the quantum implementation: detecting the transition from a Calm to a Turbulent regime. We applied the Paper Circuit to a set of three canonical market windows ($N=50$ days) identified in Phase 1:

- *Calm '17 (Green)*: A stable period in 2017.
- *Calm '18 (Blue)*: A stable period in 2018.
- *Crisis '20 (Orange)*: The onset of the COVID-19 market crash.

Figure 6.5 illustrates the context of these samples. The goal is to demonstrate that the distance between two Calm periods is negligible (Stability), while the distance to the Crisis period is significant (Detection).

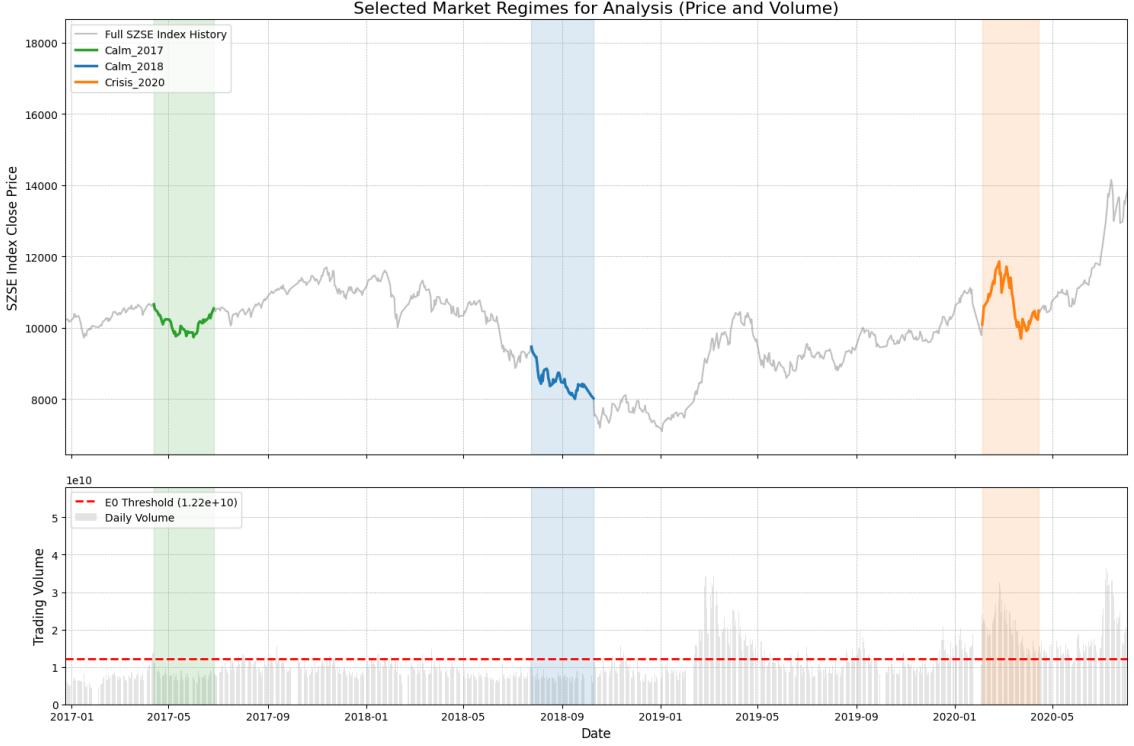


Figure 6.5: Context analysis for the Regime Discrimination Test. The shaded regions indicate the 50-day samples extracted for analysis: Calm '17 (green), Calm '18 (blue), and Crisis '20 (orange). The statistical separation of these windows is quantified in Figure 6.6.

Results and Quantum Advantage

Figure 6.6 presents the Quantum MMD scores. The analysis yields two decisive findings:

First, the stability of the test is confirmed, as the distance between the two calm periods (Calm_17 vs Calm_18) is low (0.0076 ± 0.0029) and its confidence interval overlaps with the baseline noise, indicating no false positives. Second, the test demonstrates successful quantum detection: the distance between the calm and turbulent regimes (Calm_17 vs Crisis_20) is significantly higher (0.0253 ± 0.0075), with its cross-distance clearly separated from the noise floor.

This result verifies that the quantum feature map captures high-order temporal correlations in the crisis data, providing a clear detection signal where classical univariate tests fail.

6.3 Phase 3: Forecasting Application

The ultimate test of the identified regimes is their predictive utility. If the “Turbulent” state represents a distinct physical phase, a model trained exclusively on this

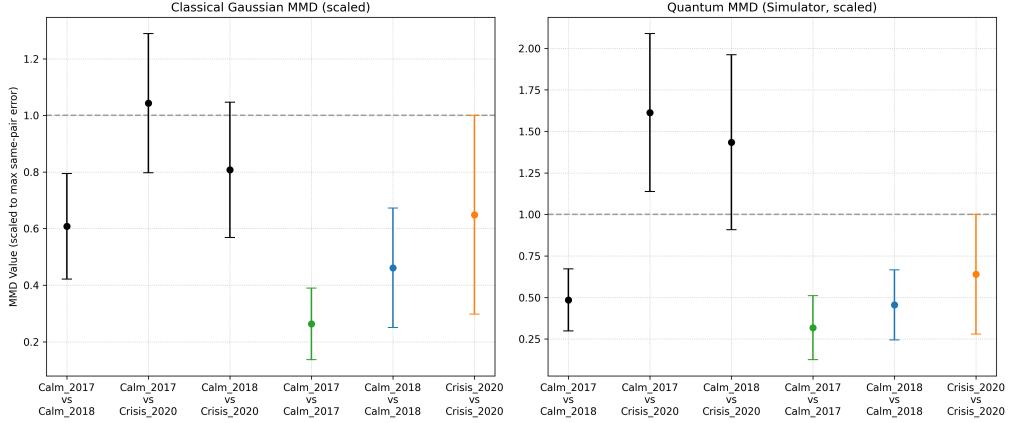


Figure 6.6: Regime Discrimination Results. The comparison between two Calm periods (leftmost black bar) is negligible. The comparison between Calm and Crisis (middle black bar) shows a sharp spike, confirming successful detection of the regime shift illustrated in Figure 6.5. This separation is significantly clearer than that achieved by the classical benchmark test, highlighting the quantum advantage.

regime should capture the crisis dynamics efficiently, without the dilution of years of calm data.

To validate this, we implemented a rigorous *Walk-Forward Validation* on the 2015 market crash.

6.3.1 Methodology: The Merged Crisis Model

As detailed in the methodology, the 2015 crisis consists of several high-energy islands. We merged the contiguous multimodal blocks (Blocks 2, 4, 6, and 8) to create a specialised *Crisis Training Set* of 76 trading days. As described in Chapter 5, we compared a *Baseline Model* trained on the full 4-year history against a specialised *Crisis Model* trained exclusively on these 76 days. Both models forecasted the subsequent independent crisis island in Late 2015 (Block 18).

6.3.2 Results: Informational Efficiency

Table 6.1 summarizes the performance metrics.

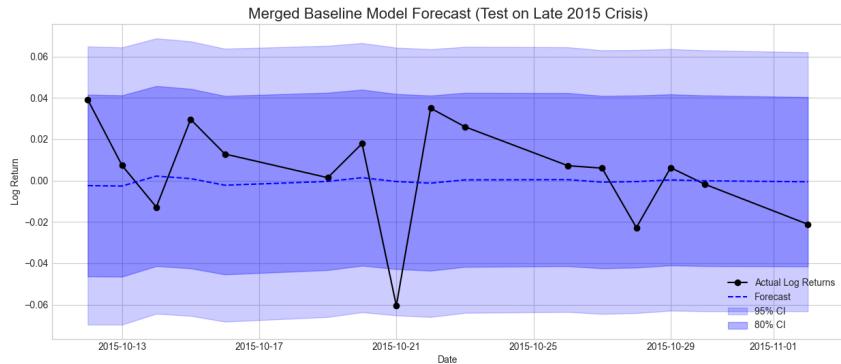
Model	RMSE (Point Forecast)	Volatility MSE (Risk)
Merged Crisis Model	0.02465	0.00000109
Fair Baseline Model	0.02501	0.00000105

Table 6.1: Forecasting Performance on the Late 2015 Crisis Period.

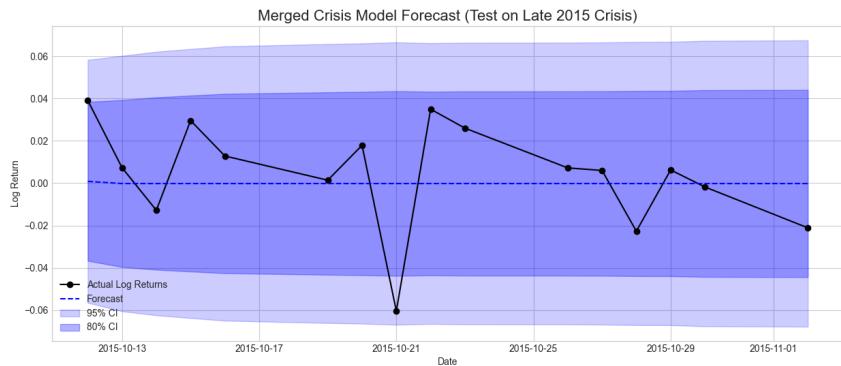
The Crisis Model achieved an RMSE of 0.0246, comparable to the Baseline (0.0250). While there is no drastic numerical improvement, the result demonstrates

high *informational efficiency*: the specialised model matched the performance of the generalist baseline using less than 7% of the training data. This confirms that during a crisis, the relevant signal is concentrated in the high-energy regimes, while long periods of calm act as noise.

Figure 6.7 compares the forecasts visually. Both models correctly identify the mean-reverting nature of the returns (point forecast ≈ 0). However, the Crisis Model (Bottom) produces confidence intervals that are robust and conservative, adequately capturing the envelope of the crash volatility.



(a) Baseline Model Forecast



(b) Crisis Model Forecast

Figure 6.7: Comparison of forecasts for the Late 2015 Crisis. *Top*: The Baseline model, trained on 4 years, captures the volatility cluster. *Bottom*: The Crisis model, trained on only 76 days, generates a virtually identical risk envelope. This proves that the essential volatility dynamics were fully contained within the short Crisis training set.

6.3.3 Concept Drift and Dynamic Adaptation

Finally, we tested the robustness of the 2015-trained model on the 2019 crisis. In this cross-temporal test, the specialised model failed to outperform the baseline, a

negative result that is theoretically significant. We tested the models on the 2019 crisis period (Block 120), with the performance metrics summarised in Table 6.2.

<i>Model</i>	<i>RMSE</i>	<i>Volatility MSE</i>
2015 Crisis Model (76 days)	0.02198	0.00000093
2015 Baseline Model (1135 days)	0.02189	0.00000090

Table 6.2: Generalisation Test: Performance of 2015-trained models on the 2019 crisis.

As the results show, the Baseline Model achieves a slightly lower error in both point and volatility forecasting. This confirms the presence of significant *Concept Drift*: the market structure during the 2015 deleveraging crash is fundamentally different from the 2020 pandemic crash. The quantum distance measured between these two periods was high (0.0152 ± 0.0062), providing a physical validation for this statistical observation. This confirms that “Crisis” is not a single static state; the market evolves. Therefore, a robust trading system cannot rely on a static historical model but requires a dynamic detection tool, like the Quantum Two-Sample Test, to trigger retraining on the most recent relevant regime.

Chapter 7

Conclusion and Future Directions

This thesis investigated whether quantum-inspired methodologies could offer a tangible advantage in the analysis and forecasting of financial time series. Starting from the observation that classical models often struggle to capture the abrupt regime shifts characteristic of financial crises, we explored the framework of Quantum Probability as a means to model the complex dynamics of market participants.

7.1 Summary of Findings

Our research followed a structured path from theoretical formulation to empirical validation, yielding three primary contributions:

1. *Validation of Physical Market Regimes.* By applying the Lin Li framework to the SZSE Component Index (2011–2020), we confirmed that trading volume acts as a reliable proxy for market energy. The Hall-York statistical test identified a stable threshold E_0 , effectively separating the market into *Calm* (unimodal) and *Turbulent* (multimodal) regimes. This validates the physical hypothesis that financial crises represent a distinct phase of the system, characterised by high persistence (82% probability of staying in a turbulent state), rather than merely being statistical outliers in a static distribution.
2. *Quantum Advantage in Detection.* We adapted and implemented a Quantum Two-Sample Test using parameterised quantum circuits. On real-world data, this quantum-circuit-based approach demonstrated a superior sensitivity compared to the classical Maximum Mean Discrepancy (MMD) benchmark. Specifically, it successfully distinguished between regime shifts (e.g., Calm '17 vs Crisis '20) with a higher Signal-to-Noise Ratio than the classical MMD benchmark, confirming that quantum feature maps can capture temporal correlations invisible to standard statistical tests.

3. *Informational Efficiency in Forecasting.* The application of this regime-awareness to forecasting produced a significant result regarding data efficiency. In a rigorous walk-forward validation on the 2015 crash, a “Crisis Model” trained on less than 7% of the available historical data (specifically, the high-energy islands) matched the predictive accuracy (RMSE) of a baseline model trained on the full history. This suggests that during turbulent periods, the predictive signal is highly concentrated in high-energy states, while long periods of calm effectively act as noise for crisis prediction.

7.2 Implications: Concept Drift and Dynamic Adaptation

Beyond the numerical results, this study highlights a critical structural weakness in static thresholding methods. Our analysis revealed that the energy threshold E_0 is subject to long-term drift, likely due to market growth and inflation. An absolute threshold calibrated in 2011 generates false positives in 2020 simply because nominal volumes have increased, not because the regime has changed physically.

This validates the necessity of the *Quantum Two-Sample Test*. Unlike a static threshold, the quantum test is a *relative* measure: it computes the statistical distance between the current market window and a reference baseline. This makes it inherently robust to secular trends like volume inflation. Furthermore, by detecting the structural distance between different crises (as observed between 2015 and 2020), the quantum test acts as a “Concept Drift Detector”, signaling precisely when a forecasting model has become obsolete and triggering a retraining cycle on fresh data.

7.3 Future Research Directions

To bridge the gap between this proof-of-concept and industrial application, we propose three avenues for future research:

- *Dynamic Thresholding.* Our analysis assumed a static energy threshold E_0 . However, the detection of frequent, short-lived multimodal islands in later years suggests that this threshold is subject to drift, likely driven by long-term trends in market liquidity. Developing an adaptive algorithm that updates E_0 based on a rolling window would ensure the regime classification remains robust over decades.

- *Data Augmentation via Regime Clustering.* A major challenge in training specialised models is data scarcity (e.g., short crisis durations). Future research should develop mathematical methods to aggregate non-contiguous but *statistically similar* regime islands (identified via the quantum test) into consolidated training sets. Solving the issue of artificial price jumps at the concatenation points would allow for the creation of robust “Cluster Models” trained on larger datasets without violating stationarity assumptions.
- *Dynamic Model Selection via Regime Similarity.* Our analysis revealed that while different crises are structurally distinct (e.g., 2015 vs 2020), intra-crisis periods can be statistically identical. A promising future direction is to implement a “Regime Library”: by computing the quantum statistical distance between the current unfolding market window and various historical reference periods in real-time, the system could automatically select and deploy the forecasting model (or Cluster Model) trained on the most structurally similar past regime. This would evolve the system from simple regime detection to fully adaptive model selection.
- *Implementation on Quantum Hardware.* While this study relied on classical simulation to isolate theoretical performance, the ultimate goal is scalability. Future work should focus on implementing the kernel estimation on actual NISQ devices, leveraging error mitigation techniques to verify if the quantum advantage persists in the presence of hardware noise.

In conclusion, this thesis has moved beyond the theoretical proposal of Quantum Finance to provide a concrete, empirical validation on real-world data. We have demonstrated that financial markets exhibit distinct, persistent physical regimes that are invisible to standard univariate statistics but detectable via quantum feature maps.

While the forecasting application highlights the challenge of non-stationarity, our findings establish a clear path forward: the future of algorithmic trading lies not in static “super-models” trained on decades of history, but in agile, regime-aware systems that use quantum sensitivity to adapt dynamically to the shifting landscape of market risk. This work provides the foundational methodology for building such systems.

Appendix A

Code and Data Availability

To ensure reproducibility, the codebase is organised into two primary repositories covering the methodology and the statistical testing framework.

Repository A: Regime Identification Framework

Implements the methodology for detecting the multimodality threshold E_0 using the Hall-York test, as derived from Lin Li's framework.

URL: <https://github.com/giorgiamazzaro/thesis-quantum-finance>

Repository B: Quantum Two-Sample Test

Contains the PennyLane implementation of the Quantum Maximum Mean Discrepancy (MMD) test used for regime validation.

URL: <https://github.com/giorgiamazzaro/quantum-two-sample-test>

Appendix B

Additional Quantum Circuit Architectures

This appendix details the alternative Parametric Quantum Circuit (PQC) architectures evaluated during the model selection phase but ultimately discarded due to inferior performance compared to the Paper Circuit.

B.1 Evaluated Architectures

In addition to the reference architecture (Paper Circuit) detailed in Chapter 5, we benchmarked three alternative designs:

- *TFIM Circuit*: Inspired by the Transverse Field Ising Model, a standard Hamiltonian in condensed matter physics. It employs layers of R_x rotations and ZZ coupling gates [44].
- *LTFIM Circuit*: A layered extension of the TFIM ansatz (Layered Transverse Field Ising Model) incorporating additional R_z rotations to increase expressivity.
- *Hardware-Efficient Circuit*: An architecture optimised for current superconducting processors. It uses native gates (R_y, R_z, CZ) to minimise circuit depth and gate errors [45].

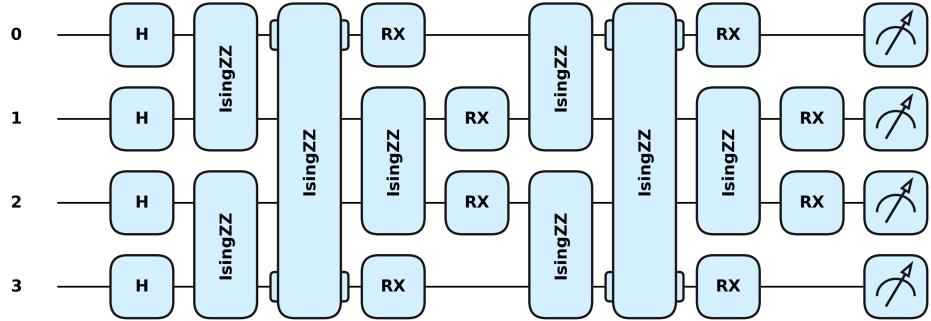


Figure B.1: The TFIM Circuit architecture. It uses a pattern of ZZ-interactions (entanglement) followed by local X-rotations.

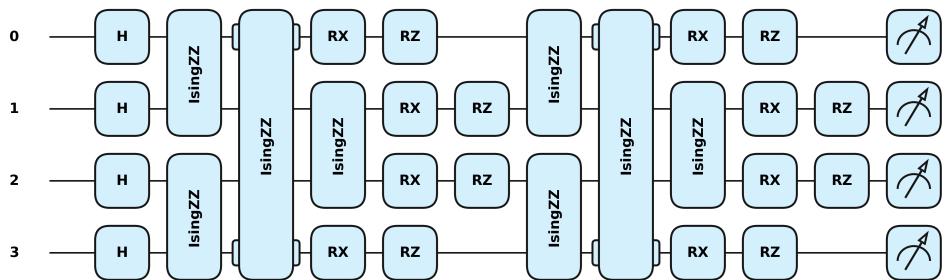


Figure B.2: The LTFIM Circuit architecture. An enhancement of the TFIM design with additional parameterized Z-rotations.

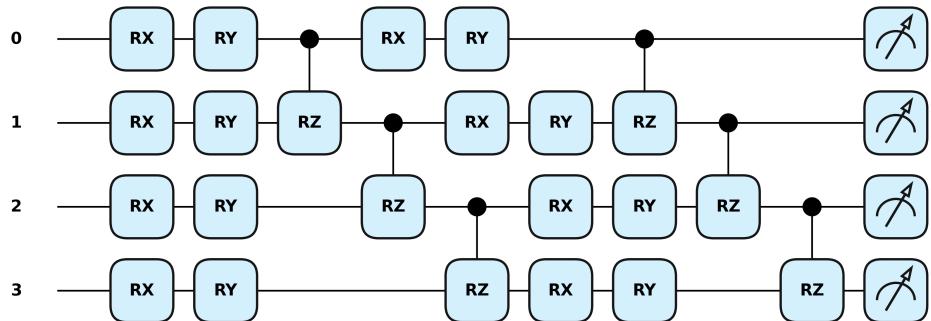


Figure B.3: The Hardware-Efficient Circuit architecture. Designed to fit the connectivity of NISQ devices, using nearest-neighbor CZ gates.

B.2 Benchmark Comparison

Figure B.4 illustrates the raw distribution of MMD scores for all four circuits on the synthetic dataset. The plot highlights the critical issue with the TFIM-based architectures: their output MMD scores are orders of magnitude smaller than the Paper and Hardware-Efficient circuits, indicating a failure to map the temporal features into a distinguishable quantum state (kernel collapse). While the Hardware-Efficient circuit produces larger values, the overlap between the "Different" (Salmon) and "Same" (Blue) distributions is substantial, resulting in a Signal-to-Noise Ratio < 1 . Only the Paper Circuit achieves a clean separation.

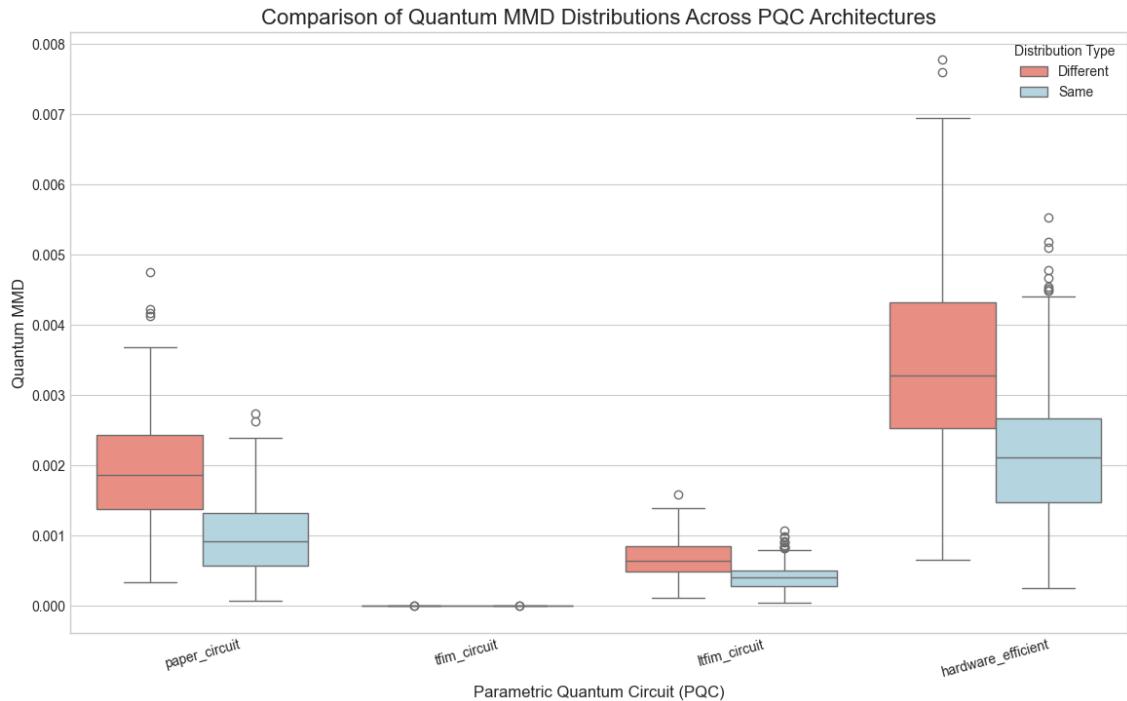


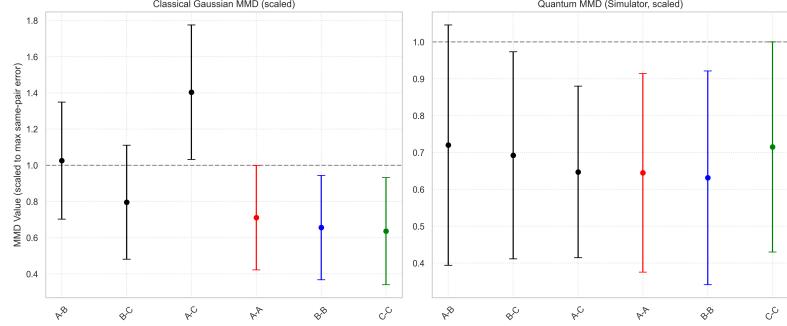
Figure B.4: Boxplot of Quantum MMD distributions for all tested architectures. Note the scale difference on the y-axis. The TFIM and LTFIM circuits produce negligible distances. The Hardware-Efficient circuit (right) shows high variance and overlap. Only the Paper Circuit (left) shows a robust separation between signal and noise.

B.3 Detailed Benchmark Results

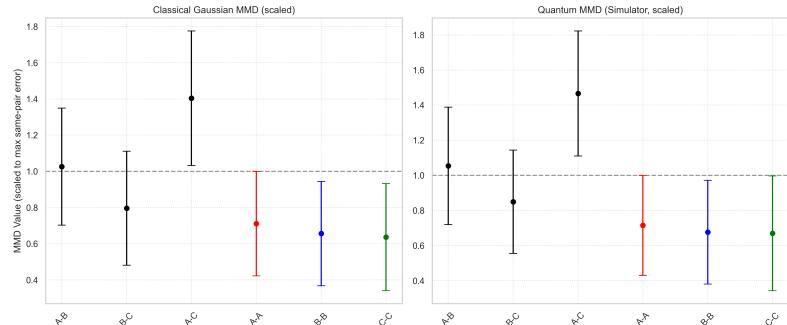
We compared the discriminatory power of each circuit using a scaled Signal-to-Noise metric. Figure B.5 displays the results for the architectures that failed to meet the selection criteria.

- *TFIM Circuit*: Produced MMD scores orders of magnitude smaller than the signal threshold (kernel collapse).

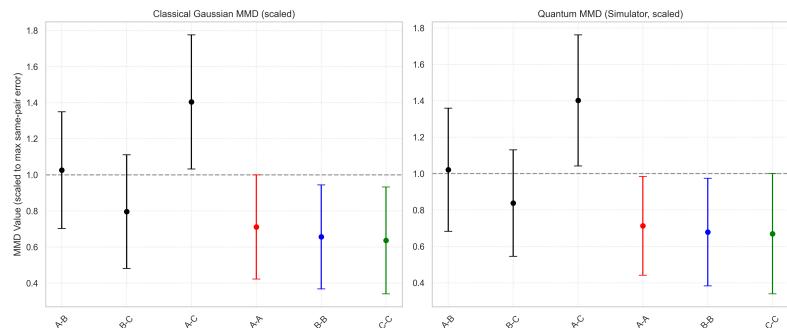
- *LTFIM Circuit*: Showed marginal improvement but failed to achieve a robust separation.
- *Hardware-Efficient Circuit*: Exhibited high variance, resulting in significant overlap between the signal and noise distributions.



(a) TFIM Circuit (Failed)



(b) LTFIM Circuit (Marginal)



(c) Hardware-Efficient Circuit (Marginal)

Figure B.5: Benchmark results for the discarded architectures. The coloured bars represent the baseline self-distances (noise): Red (A-A), Blue (B-B), and Green (C-C). The black bars represent the cross-distances (signal): A-B, B-C, and A-C. A robust test requires the A-C signal (third black bar) to be clearly separated from the noise ceiling, represented by the dashed horizontal line at $y = 1.0$.

Appendix C

Identified Market Regimes

The Hall-York thresholding process identified 178 distinct contiguous market regimes. Table C.1 lists the ten longest contiguous periods for each regime type, which were prioritised for the forecasting analysis.

Block ID	Regime	Start Date	End Date	Duration (Days)
Top 10 Longest Unimodal (Calm) Periods				
1	Unimodal	2011-01-04	2015-05-19	1060
63	Unimodal	2016-11-15	2017-04-07	96
109	Unimodal	2018-03-26	2018-07-23	81
65	Unimodal	2017-04-13	2017-07-14	64
111	Unimodal	2018-07-25	2018-10-19	57
61	Unimodal	2016-08-19	2016-11-10	53
119	Unimodal	2018-11-21	2019-02-11	52
93	Unimodal	2017-11-22	2018-01-05	32
151	Unimodal	2019-10-30	2019-12-12	32
45	Unimodal	2016-04-21	2016-05-30	27
Top 10 Longest Multimodal (Turbulent) Periods				
162	Multimodal	2020-06-01	2020-09-14	74
178	Multimodal	2020-12-17	2021-03-30	68
120	Multimodal	2019-02-12	2019-04-29	54
156	Multimodal	2020-02-04	2020-04-10	48
250	Multimodal	2023-03-16	2023-05-15	39
212	Multimodal	2022-05-31	2022-07-25	39
2	Multimodal	2015-05-20	2015-07-09	36
273	Multimodal	2023-09-05	2023-10-24	30
20	Multimodal	2015-11-04	2015-12-04	23
154	Multimodal	2019-12-25	2020-01-23	21

Table C.1: The longest contiguous market regimes identified by the Hall-York thresholding process. Note the exceptional stability of the 2011-2015 period (Block 1) compared to the more fragmented dynamics of later years.

Appendix D

Supplementary Context Plots

This appendix provides the detailed price and volume context for the validation experiments discussed in Chapter 6. The context plots in this appendix follow a consistent visualisation style. The top panel always displays the closing price, which is plotted as a line chart to represent its continuous evolution over time. The bottom panel displays the daily trading volume, which is plotted as a bar chart to reflect its discrete, transactional nature on each trading day.

D.1 Baseline Stability (2011–2013)

Figure D.1 illustrates the market conditions for the three consecutive years used in the Baseline Stability Test. The trading volume remains consistently below the threshold E_0 , confirming that the entire period belongs to the unimodal Calm regime used to verify the test's false-positive rate.

D.2 Robustness Check (2016–2018)

Figure D.2 shows the context for the robustness check performed on disjoint unimodal blocks (Blocks 63, 65, 109). These periods, occurring after the 2015 crash, confirm that the market returns to a stable ground state characterized by low volume, comparable to the 2011 baseline.

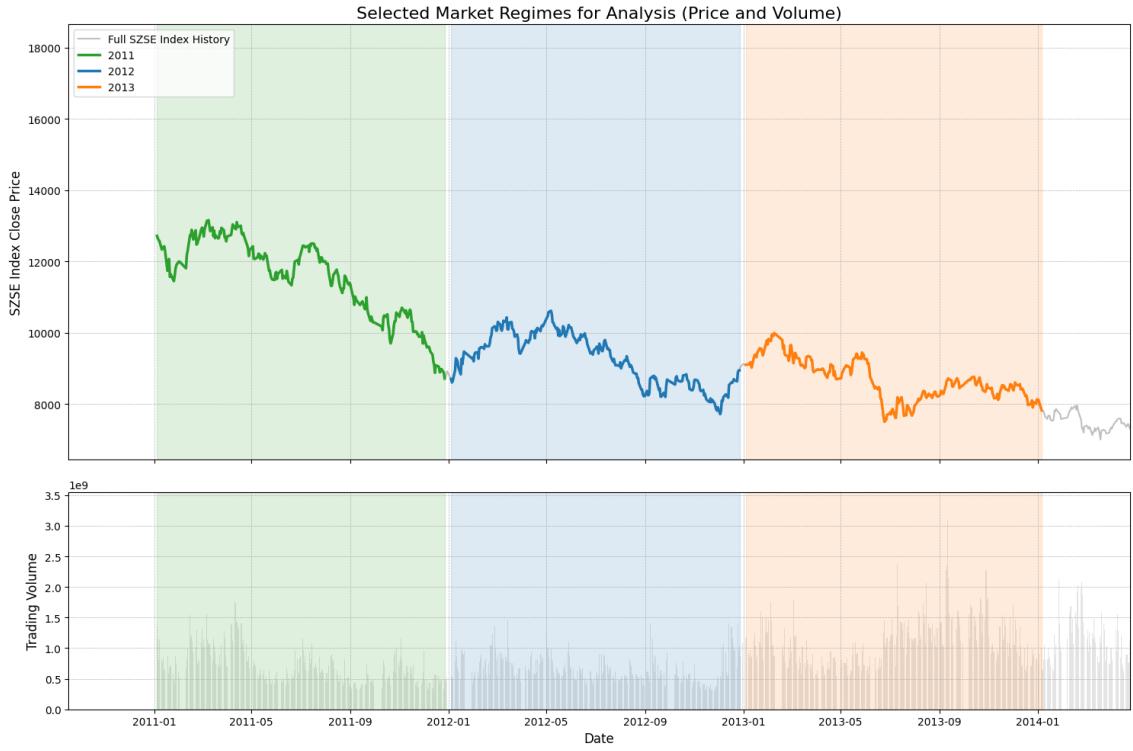


Figure D.1: Context analysis for the Baseline Validation (2011–2013), shown in Figure 6.3. The trading volume consistently remains below the critical threshold E_0 , confirming that the entire period belongs to the stable Calm regime.

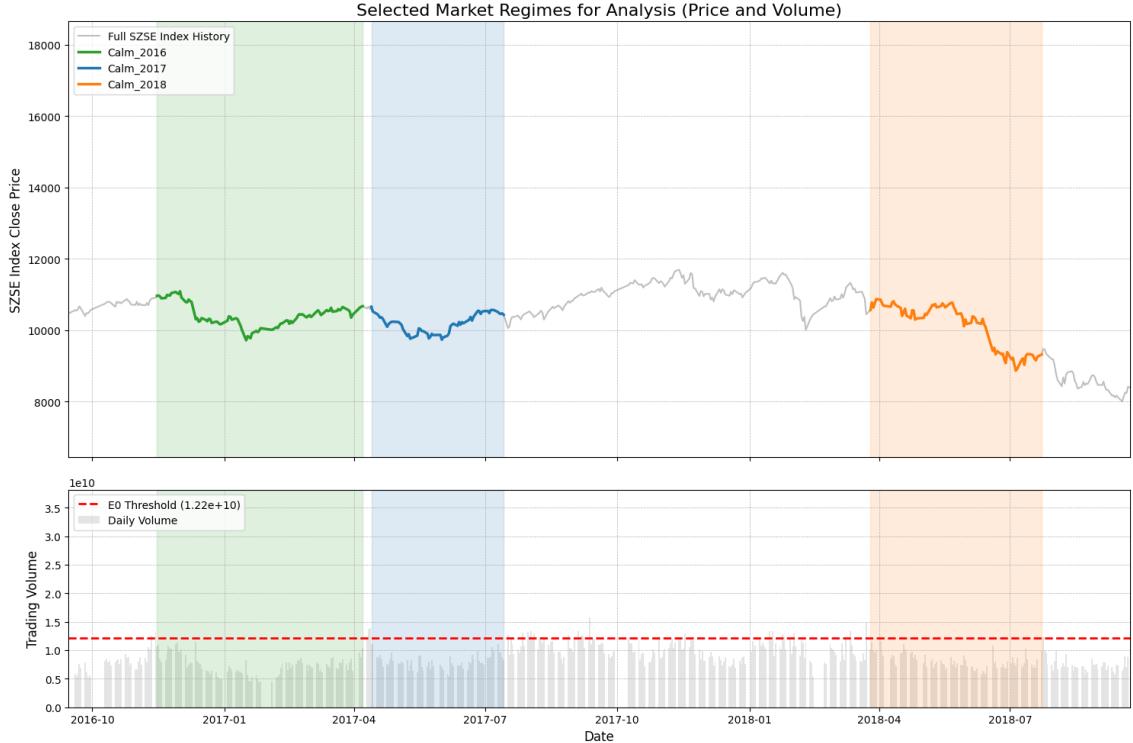


Figure D.2: Context analysis for the Robustness Check. The selected windows (Blocks 63, 65, 109) represent stable unimodal islands in the post-2015 era.

Bibliography

- [1] Li Lin. “Quantum Probability Theoretic Asset Return Modeling: A Novel Schrödinger-Like Trading Equation and Multimodal Distribution”. In: *SSRN Electronic Journal* (2024). DOI: [10.2139/ssrn.4691510](https://doi.org/10.2139/ssrn.4691510) (cit. on pp. 4, 36, 39, 40, 58).
- [2] David Garvin et al. “Quantum Two-Sample Test for Investment Strategies”. In: *arXiv preprint* (2024). arXiv: [2404.06213](https://arxiv.org/abs/2404.06213). URL: [http://dx.doi.org/10.2139/ssrn.4789400](https://dx.doi.org/10.2139/ssrn.4789400) (cit. on pp. 5, 48).
- [3] Louis Bachelier. “Théorie de la spéculation”. In: *Annales Scientifiques de l’École Normale Supérieure* 17 (1900), pp. 21–86 (cit. on p. 7).
- [4] Fischer Black and Myron Scholes. “The Pricing of Options and Corporate Liabilities”. In: *Journal of Political Economy* 81.3 (1973), pp. 637–654 (cit. on p. 7).
- [5] Eugene F. Fama. “Efficient Capital Markets: A Review of Theory and Empirical Work”. In: *The Journal of Finance* 25.2 (1970), pp. 383–417 (cit. on p. 7).
- [6] Chris Brooks. *Introductory Econometrics for Finance*. 3rd. Cambridge University Press, 2014 (cit. on p. 8).
- [7] Rama Cont. “Empirical properties of asset returns: stylized facts and statistical issues”. In: *Quantitative Finance* 1.2 (2001), pp. 223–236 (cit. on p. 9).
- [8] Benoit Mandelbrot. “The variation of certain speculative prices”. In: *The Journal of Business* 36.4 (1963), pp. 394–419 (cit. on p. 9).
- [9] George E. P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day, 1970 (cit. on p. 10).
- [10] Tim Bollerslev. “Generalized Autoregressive Conditional Heteroskedasticity”. In: *Journal of Econometrics* 31.3 (1986), pp. 307–327 (cit. on p. 11).
- [11] John Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (2018), p. 79. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79) (cit. on p. 13).

- [12] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition. Cambridge University Press, 2010 (cit. on pp. 14, 15, 28).
- [13] David J. Griffiths. *Introduction to Quantum Mechanics*. 3rd. Cambridge University Press, 2018 (cit. on pp. 14, 19).
- [14] P. A. M. Dirac. *The Principles of Quantum Mechanics*. Oxford: Clarendon Press, 1930 (cit. on p. 14).
- [15] Erwin Kreyszig. *Introductory Functional Analysis with Applications*. New York: John Wiley & Sons, 1989 (cit. on p. 14).
- [16] W. Heisenberg. “Über den anschaulichen Inhalt der quantentheoretischen Kinetik und Mechanik”. In: *Zeitschrift für Physik* 43.3-4 (1927), pp. 172–198 (cit. on p. 20).
- [17] Hale F. Trotter. “On the product of semi-groups of operators”. In: *Proceedings of the American Mathematical Society* 10.4 (1959), pp. 545–551 (cit. on p. 25).
- [18] Masuo Suzuki. “Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems”. In: *Communications in Mathematical Physics* 51.2 (1976), pp. 183–190 (cit. on p. 25).
- [19] Daniel Gottesman. “The Heisenberg representation of quantum computers”. In: *arXiv preprint quant-ph/9807006* (1998) (cit. on p. 28).
- [20] Peter W. Shor. “Algorithms for Quantum Computation: Discrete Logarithms and Factoring”. In: *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (1994), pp. 124–134 (cit. on p. 29).
- [21] E. Schrödinger. “Discussion of probability relations between separated systems”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 31 (4 1935), pp. 555–563 (cit. on p. 29).
- [22] Albert Einstein, Boris Podolsky, and Nathan Rosen. “Can quantum-mechanical description of physical reality be considered complete?” In: *Physical review* 47.10 (1935), p. 777 (cit. on p. 30).
- [23] J. S. Bell. “Bertlmann’s socks and the nature of reality”. In: *Journal de Physique Colloques* 42.C2 (1981), pp. 41–62. DOI: [10.1051/jphyscol:1981202](https://doi.org/10.1051/jphyscol:1981202) (cit. on p. 30).
- [24] John F. Clauser et al. “Proposed Experiment to Test Local Hidden-Variable Theories”. In: *Physical Review Letters* 23.15 (1969), pp. 880–884. DOI: [10.1103/PhysRevLett.23.880](https://doi.org/10.1103/PhysRevLett.23.880) (cit. on p. 31).

[25] Alain Aspect, Philippe Grangier, and Gérard Roger. “Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell’s Inequalities”. In: *Physical Review Letters* 49.2 (1982), pp. 91–94. DOI: [10.1103/PhysRevLett.49.91](https://doi.org/10.1103/PhysRevLett.49.91) (cit. on p. 31).

[26] Marcello Benedetti et al. “Parameterized quantum circuits as machine learning models”. In: *Quantum Science and Technology* 4.4 (2019), p. 043001. DOI: [10.1088/2058-9565/ab4eb5](https://doi.org/10.1088/2058-9565/ab4eb5) (cit. on p. 31).

[27] Jacob Biamonte et al. “Quantum machine learning”. In: *Nature* 549.7671 (2017), pp. 195–202. DOI: [10.1038/nature23474](https://doi.org/10.1038/nature23474) (cit. on p. 31).

[28] Maria Schuld and Francesco Petruccione. *Supervised Learning with Quantum Computers*. Springer International Publishing, 2018. ISBN: 978-3-319-96423-2. DOI: [10.1007/978-3-319-96424-9](https://doi.org/10.1007/978-3-319-96424-9) (cit. on p. 31).

[29] Vojtch Havlík et al. “Supervised learning with quantum-enhanced feature spaces”. In: *Nature* 567.7747 (2019), pp. 209–212. DOI: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2) (cit. on p. 31).

[30] M. Cerezo et al. “Variational Quantum Algorithms”. In: *Nature Reviews Physics* 3.9 (2021), pp. 625–644. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9) (cit. on p. 31).

[31] Wolfgang Sakuler et al. “A real world test of Portfolio Optimization with Quantum Annealing”. In: *arXiv preprint arXiv:2303.12601* (2023). URL: <https://arxiv.org/abs/2303.12601> (cit. on p. 34).

[32] Mirko Mattesi et al. “Diversifying Investments and Maximizing Sharpe Ratio: A Novel Quadratic Unconstrained Binary Optimization Formulation”. In: *Quantum Reports* 6.2 (2024), pp. 262–278. DOI: [10.3390/quantum6020018](https://doi.org/10.3390/quantum6020018) (cit. on p. 34).

[33] David Bucher et al. “Incentivizing Demand-Side Response Through Discount Scheduling Using Hybrid Quantum Optimization”. In: *IEEE Transactions on Quantum Engineering* 5 (2024), pp. 1–13. DOI: [10.1109/TQE.2024.3411019](https://doi.org/10.1109/TQE.2024.3411019) (cit. on p. 34).

[34] Francesco Di Marcantonio et al. “Quantum Advantage Seeker with Kernels (QuASK): a software framework to speed up the research in quantum machine learning”. In: *Quantum Machine Intelligence* 5.1 (2023), p. 13. DOI: [10.1007/s42484-023-00107-2](https://doi.org/10.1007/s42484-023-00107-2) (cit. on p. 35).

[35] Nicolò Bonaldi et al. “Boost clustering with Gaussian Boson Sampling: a quantum-classical hybrid approach”. In: *Quantum Machine Intelligence* 6.2 (2024), p. 65. DOI: [10.1007/s42484-024-00185-w](https://doi.org/10.1007/s42484-024-00185-w) (cit. on p. 35).

[36] Marco Maronese et al. “The Quantum Amplitude Estimation Algorithms on Near-Term Devices: A Practical Guide”. In: *Quantum Reports* 6.1 (2024), pp. 1–13. doi: [10.3390/quantum6010001](https://doi.org/10.3390/quantum6010001) (cit. on p. 35).

[37] Andrey Nikolaevich Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Berlin: Springer, 1933 (cit. on p. 37).

[38] David Orrell. “A Quantum Oscillator Model of Stock Markets”. In: *Quantum Economics and Finance* 1.1 (2024), pp. 74–81. doi: [10.1177/29767032231226220](https://doi.org/10.1177/29767032231226220) (cit. on p. 38).

[39] Investing.com. *SZSE Component Index Historical Data*. Accessed: 2025-04-22. 2025. URL: <https://it.investing.com/indices/szse-component-historical-data> (cit. on p. 45).

[40] B. W. Silverman. “Using Kernel Density Estimates to Investigate Multimodality”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 43.1 (1981), pp. 97–99. URL: <https://www.jstor.org/stable/2985156> (cit. on p. 45).

[41] Peter Hall and Matthew York. “On the Calibration of Silverman’s Test for Multimodality”. In: *Statistica Sinica* 11.2 (2001), pp. 515–536 (cit. on p. 45).

[42] Arthur Gretton et al. “A Kernel Two-Sample Test”. In: *Journal of Machine Learning Research*. JMLR.org, 2012, pp. 723–730. doi: [10.5555/2188385.2188410](https://doi.org/10.5555/2188385.2188410) (cit. on p. 50).

[43] Shiyu Liu. “The 2015 Chinese stock market crash”. In: *Available at SSRN* 2764204 (2016) (cit. on p. 59).

[44] A. Smith et al. “Exact Ising model simulation on a quantum computer”. In: *Quantum* 2 (2018), p. 109 (cit. on p. 71).

[45] R. Wiersema et al. “Hardware-efficient variational quantum algorithms for time evolution”. In: *Physical Review Research* 3.4 (2021), p. 043094 (cit. on p. 71).

Acknowledgements

This thesis represents the culmination of a challenging yet rewarding journey. Reaching this milestone would not have been possible without the guidance, encouragement, and support of those who have been by my side.

First and foremost, I would like to express my sincere gratitude to my academic supervisor, Prof. Giovanna Turvani, for her availability and support throughout the administrative process. My deepest appreciation goes to my company supervisor, Dr. Davide Tezza. Thank you for your mentorship, for guiding me through the technical challenges, and for always being available to discuss ideas and solutions. I also extend my thanks to Dr. Davide Caputo for the opportunity to work in such a stimulating environment.

I would like to mention Prof. Riccardo Adami: thank you not only for the knowledge shared but for the passion and enthusiasm you put into teaching, which have been truly contagious. A special acknowledgment goes to Stefano Biondi: our collaboration in building the *BeQuantum* community and organising its events paved the way for this thesis opportunity, and I am grateful for the direction it gave to my career. A thanks also to Antonio Gallerati for the irony, the laughs, and the pearls of wisdom shared when I needed clarity.

Desidero dedicare un ringraziamento speciale alla mia famiglia, il mio porto sicuro. A mamma, per l'infinita pazienza e per avermi sopportato nei momenti di stress. A papà, per la sicurezza che mi trasmette e per avermi insegnato a non arrendersi mai. A mio fratello, per avermi sempre spinto a guardarmi intorno, a non accontentarmi e a cercare sempre nuove prospettive. Ai miei nonni Lodovico, Giulio e Silvana, agli zii e ai cugini, per l'affetto che mi fate sentire anche a distanza. Un pensiero speciale va a Mariano, per avermi spronato fin da piccola a sognare in grande.

A significant part of this thesis was written within the walls of *Collegio Einaudi*. There is a saying: *Si scrive Collegio, si legge Famiglia* (It is written Collegio, it is read Family). For me, it has truly been a home where I felt accepted and valued. My gratitude goes to the *La Manica* zone at Sezione Crocetta. Thank you to my friends and floor-mates: Davide, Paola, Irene, Giorgia, Irving, Aldo, Matteo, Angelo,

Giulia, Nicole, Roberta, Gabriele, Alessia, Andrea, and all the other Erasmus friends, among whom Aum and Mikołaj stand out: thank you for dragging me out for walks when I was too tired to work but too stubborn to rest. I am also grateful to my longtime friends from the 4th floor: Chiara, Giulia, Vito, Matteo, and Kekko. A special mention goes to Rossella, for the morning chats and the constant kindness.

I want to thank my fellow travelers in the world of Physics and Engineering: Francesca, my constant companion since the Bachelor's degree; my course mates Federico, Sofia, Riccardo, Lorenzo, Matteo, Marco, Pietro, Giacomo, and Eva; and Daevid, Marco, Raffaele, and Renzo. I also want to thank my Quantum friends: Fabio, Giorgio, and Alessandro from the Technical University of Munich (TUM) for being a vital part of my quantum network and for the constant exchange of insights, and Riccardo for the motivating chats on research life.

I cannot forget my roots. A heartfelt thank you to my friends from Eboli and surroundings, who are always there despite the distance: Alessandra, Michele, Giovanni, Francesco, Domenico, Domitilla, Serena, and Marianna. Thank you for making me feel at home whenever I return.

Finally, a necessary thank you goes to the AI tools that supported me in organising my work and structuring my thoughts during the most chaotic moments of this project.

To everyone I might have forgotten to list but who has shared a coffee, a smile, or a word of encouragement with me: thank you.