An Examination of Share Buyback Program Execution

A Major Qualifying Project (MQP) Report
Submitted to the Faculty of
WORCESTER POLYTECHNIC INSTITUTE
in partial fulfillment of the requirements
for the Degree of Bachelor of Science in

Computer Science

By:

Robert Chiocchio

Project Advisors:

Marcel Blais
Randy Paffenroth
Stephan Sturm

Sponsored By:

Jörg Osterrieder (Berner Fachhochschule, University of Twente)

Date: 2024-03-22

An Examination of Share Buyback Program Execution

Robert Chiocchio
Department of Computer Science
Worcester Polytechnic Institute
Worcester, MA 01609
rmchiocchio@wpi.edu

2024-03-22

Abstract

This study analyzes the efficacy and risks of various share buyback strategies, with a particular focus on the fixed participation rate approach. This strategy aligns buyback volumes with a fixed percentage of daily trading volume, aiming to minimize market impact, mitigate financial risk, and optimize shareholder value. Using Value at Risk (VaR) for risk assessment and Long Short-Term Memory (LSTM) models for trading volume prediction, we compare the fixed participation rate strategy against traditional methods. Our findings emphasize the need for a shift from strategies focused on short-term gains to those prioritizing market stability and long-term shareholder value creation. This research offers insights for corporate decision-making, policy development, and the advancement of sustainable financial management practices.

Table of Contents

Figures

1 Introduction

In a financial landscape where corporate strategies pivot towards enhancing shareholder value, share buybacks emerge as a focal point, sparking debates on their impact and effectiveness. These strategies, while popular, present challenges in their execution, particularly in minimizing market influence and financial risk. This paper examines the risk and efficacy of share buybacks, exploring their strategic implications.

This strategy aligns buyback volumes with a fixed percentage of daily trading volume, aiming to minimize market impact and price manipulation risk, and optimize shareholder value under volatile conditions. Our analysis will compare the fixed participation rate strategy against traditional buyback methods, focusing on metrics related to risk, market impact, and shareholder value creation.

Our research utilizes data from BNP Paribas' share buyback program as a real-world benchmark for examination and evaluation [3]. By integrating this data with additional market information, we simulate and assess alternative share buyback strategies. This approach allows for a detailed comparison with BNP Paribas' actual strategy, providing insights into the relative risks and efficiencies of different buyback methods in real-world conditions.

This research aims to enhance understanding of share buyback strategies in corporate finance, providing empirical insights into their risk profiles and effectiveness. It seeks to inform corporate decision-making on capital allocation for enhancing shareholder value and has implications for policy-making and corporate governance.

2 Background

Corporate share buybacks are complex financial transactions with far-reaching implications. Companies employ these strategies for various reasons, from boosting earnings per share to signaling confidence in their future. However, executing buyback programs effectively requires navigating operational challenges, regulatory constraints, and potential market impacts. This background provides context for our in-depth analysis of share buyback execution strategies.

2.1 An Overview of Share Buybacks

Share buybacks, also referred to as share repurchases or stock buybacks, have become a notable trend in corporate finance. This strategy involves a company buying back its own shares from the open market or directly from shareholders. Its growing popularity, especially among major corporations, reflects a significant shift in corporate financial strategies [24]. When a company undertakes a share buyback, it uses its available funds to purchase a certain number of its shares. These purchased shares are then canceled, reducing the total number of shares available for trading in the market. Consequently, this action increases the relative ownership stake of the remaining shareholders in the company.

This trend is particularly pronounced among S&P 500 companies in the United States. Over the past five years, these companies have allocated approximately \$3.8 trillion to share repurchase programs [15]. Apple Inc. serves as a prime example: in the 2021 fiscal year, it spent \$85.5 billion on share buybacks compared to \$14.5 billion on dividends [17]. This emphasizes a broader corporate trend of utilizing share buybacks as a means to reward shareholders, influence market perceptions, and manage earnings per share metrics.

2.2 Execution Costs

Effectively managing the expenses associated with share buybacks, particularly brokerage fees, is a critical consideration. These fees can vary significantly across different transactions, depending on the execution strategy and the specifics of the broker contract, often falling into two main categories: predetermined fixed commissions and variable fees, which are linked to the brokers' trading activities during the execution phase [24, 21]. Understanding and managing these costs is vital for companies engaging in share buybacks. Brokerage fees can account for a considerable share of the overall expenses of the buyback program, affecting its overall financial efficacy.

The Financial Times sheds light on the complexities of brokerage fees through Royal Mail's 2022 share buyback experience. Despite setting aside £200 million, Royal Mail managed to secure only £184 million in shares, leading to an unexpectedly high commission rate for the investment bank involved, over 8.5% [17]. This case serves as a cautionary tale on the impact of market volatility and broker advantage, emphasizing the critical need for strategic execution and cost management to protect shareholder interests.

2.3 A "Free Lunch" for Brokers

Benchmarking is a key tool for assessing the quality of share buyback execution, with the Volume-Weighted Average Price (VWAP) commonly serving as the standard metric. However, the use of VWAP as a measure of buyback performance introduces several challenges. The influence of a significant trade on the VWAP can distort the perception of trading efficiency. In cases where a single order constitutes the bulk of the day's trading volume, the trading cost measured against VWAP could misleadingly appear as zero, failing to reflect the true cost or impact of the trade [9].

Osterrieder and Seigne critiqued the use of what they termed the "Bogus Benchmark," arguing that it might not align with shareholder interests [24]. Coined to highlight its questionable utility, the "Bogus Benchmark" refers to the arithmetic average of daily VWAPs over the buy-back period, contrasting with the "Institutional" VWAP, which is calculated by taking each day's VWAP and weighting it by that day's trading volume. Unlike the "Institutional" VWAP, the "Bogus Benchmark" simplifies this calculation to a mere average, granting brokers undue influence over the benchmark setting period [24, 21].

The foundational issue with the "Bogus Benchmark" lies in its misalignment with the primary objectives of share buy-backs: to maximize the number of shares repurchased or to minimize the cost of capital transfer. Success against this benchmark does not necessarily equate to achieving these goals, signaling a disconnect between benchmark performance and shareholder benefits.

Siege and Osterreider argue that certain execution strategies, particularly those benchmarked against the "Bogus Benchmark," might offer a "Free Lunch," or illusory benefits at the expense of shareholder interests [21]. This term implies that while brokers might achieve nominal success in beating this benchmark, such victories do not translate into genuine value for shareholders, instead reflecting a misallocation of resources that could have been more effectively utilized in direct share repurchases. Brokers are incentivized to spread out large order executions to manipulate their performance against the "Bogus Benchmark" at the expense of additional risk for the shareholders. This discrepancy underscores the inappropriateness of the "Bogus Benchmark" for guiding buy-back strategies, as optimizing against this benchmark fails to ensure maximized shareholder value [24, 21, 9].

Conflict of interest is a significant concern, particularly regarding broker compensation tied to outperforming the "Bogus Benchmark". When brokers are incentivized based on this benchmark, their goals may diverge from those of shareholders. Given the broker's partial control over the benchmark setting, there is a risk that brokers may optimize for outcomes favorable to themselves rather than the shareholders, creating a misalignment of interests. This situation is further complicated when brokers guarantee buy-back prices relative to this arbitrary benchmark, exposing them to risks that are, paradoxically, within their power to influence [17, 24, 9].

2.4 Considerations for Buyback Strategies

The volume of shares a company repurchases can significantly influence their share price. Effective management of the participation rate—the proportion of shares executed relative to the total market volume—is instrumental in maintaining market stability and aligning with strategic financial objectives. A higher participation rate can drive up share prices by increasing demand, while a lower rate reduces the buyback's market influence, thus preventing undesired price volatility. Achieving the right balance is essential for companies looking to execute buybacks efficiently and with minimal market disruption, emphasizing the need to determine the optimal daily trading volumes [1, 23].

Furthermore, the injection of large orders into the market can pose challenges to immediate liquidity at current price levels. Liquidity refers to the ease with which an asset can be bought or sold without significantly impacting its price. For example, a significant buy order might consume all available sell orders at the lowest ask price, moving on to higher prices to fulfill the order. To mitigate the price pressure from large orders, traders and institutions often employ strategies such as breaking up large orders into smaller, discreet "iceberg orders" and using algorithmic trading designed to minimize market impact [9]. The flexibility in execution timelines, or temporal optionality, allows corporations to navigate through market volatility efficiently. This adaptability, essential for minimizing price risk and managing stock price volatility, emphasizes the strategic significance of execution timing in mitigating adverse impacts on share price [23].

Price risk, stemming from stock market volatility and exposure duration to market changes, poses a challenge in share buybacks. Companies use the Value at Risk (VaR) method to quantify and mitigate this risk. VaR is a measure of the potential loss in value of an asset or portfolio over a given time period, for a specified confidence level. VaR helps predict potential financial losses due to price volatility, enabling firms to strategize their buyback executions without significantly impacting the stock price [24, 10].

Existing studies, such as those conducted by Osterrieder and Seigne [22], have begun to shed light on the complexities of buyback execution. However, the financial community's understanding of how to optimize these elements to maximize shareholder value while minimizing costs and market impact remains incomplete. This gap indicates a pressing need for further research to develop and refine models of buyback execution that can guide corporations in their capital allocation decisions. As the market environment continues to evolve, with changing regulatory landscapes and financial innovation, the quest for optimal buyback strategies becomes even more critical. This study aims to contribute to closing this gap by challenging conventional buyback approaches and highlighting strategies that prioritize risk management, cost-effectiveness, and market stability.

3 Methodology

In this study, we examine BNP Paribas' (BNP.PA) share buyback program to understand the dynamics and strategic considerations of corporate share repurchase strategies. This section outlines the methodology employed to analyze the risks associated with share buyback strategies, with a particular focus on how a fixed participation rate strategy can provide insights into the risk profile of the company's original buyback program.

We source historical trading data, including daily trading volumes and closing prices, from Yahoo Finance [29]. Additionally, we reviewed weekly buyback reports from BNP Paribas from April 3rd, 2023, to August 3rd, 2023, detailing the daily quantities of shares repurchased and their average prices [3]. These reports provide a direct insight into the company's buyback activities.

Our methodology leverages Python to simulate and assess buyback strategies during the specified period, employing metrics such as Value at Risk (VaR). Trading volumes are forecast using various methods, including Long Short-Term Memory (LSTM) networks. Asset prices are simulated with Geometric Brownian Motion, allowing us to explore a wide range of potential market conditions.

Our primary objective is to understand how the risks of a fixed participation rate strategy compare to those of the company's original buyback strategy. This will help isolate risks specific to the original implementation and understand how different buyback behaviors might influence the overall risk profile.

3.1 Preliminary Analysis

Our preliminary analysis aims to provide an understanding of the strategic execution of buybacks in relation to company-disclosed targets, market conditions, and internal strategic objectives. This section outlines the approach adopted to assess the alignment between BNP Paribas' disclosed buyback plans and their actual implementation, alongside an analysis of the strategic timing and market impact of these activities.

3.1.1 Review of Buyback Disclosures and Strategic Objectives

Initially, we examine publicly disclosed information regarding the company's share buyback program. This includes annual reports, press releases, and regulatory filings where the company outlines its objectives for the share repurchases, such as intentions to return value to shareholders, optimize capital structure, or enhance earnings per share. Additionally, this review covers the specific targets and limits set by the company, encompassing the total number of shares targeted for repurchase, the budget allocated for these transactions, and any temporal constraints or conditions highlighted by the company.

3.1.2 Visualization of Buyback Activities

Following the review of buyback disclosures, we evaluate the execution of the share repurchase program by examining the quantity, cost, and timing of the shares repurchased. This involves a detailed examination of the buybacks' share of the daily trading volume, termed the participation rate. Understanding this rate is key to assessing the impact of buyback activities on market dynamics and maintaining stable prices. By investigating how much of the daily trading volume is made up of buybacks, we assess the strategy's influence on the market and its commitment to reducing market disturbances.

This examination is supplemented by visualizations created using Python in a Jupyter notebook. These visualizations provide insights into the pattern of buyback activities over the designated period. They aid in evaluating adherence to set targets and restrictions, identifying periods of unusual buyback activity, and examining any departures from the planned approach. Visualization tools uncover insights into the buyback strategy, showing its responsiveness to market conditions and alignment with objectives.

3.1.3 Comparison Against Market and Internal Targets

In the final stage of our exploratory analysis, we compare the actual buyback activities to the company's internal targets. This involves evaluating whether the executed buybacks remained within the prescribed budget, adhered to the specified time frame, and achieved the target share volume outlined in the company disclosures. Additionally, we assess the strategic timing of the buybacks, analyzing how the company responded to market conditions, such as stock price fluctuations and market volatility, to maximize the value of the repurchases.

3.2 Buyback Strategy Simulation

This section outlines our approach to creating and simulating buyback strategies. These strategies are simulated over the same time frame as the original buyback, facilitating direct comparison. To facilitate this analysis, we developed a Python module to generate and analyze buyback strategies. This custom library standardizes our performance evaluation tools and allows for consistent strategy simulation over a range of market conditions.

3.2.1 Fixed Participation Rate Strategy

Our analysis focuses on the Fixed Participation Rate strategy—also referred to as the Percentage of Volume (POV)—for share buybacks. This strategy aligns repurchase volumes with market conditions to maintain stability. It involves setting daily repurchase targets as a fixed percentage of estimated trading volumes, ensuring that buyback activity remains proportionate to market liquidity. The aim is to minimize market disruption and share price impact while avoiding undue risk [13].

To calculate daily repurchase targets for each strategy, a predetermined participation rate percentage is multiplied by the estimated trading volume. The amount of shares purchased is adjusted to align with budgetary limitations and buyback targets. While not modeled here, implementations of this strategy should consider regulatory constraints when setting daily targets.

To determine the participation rates, we first review the historical data. We investigate a range of rates akin to the mean and median of the original strategy's participation rate over different periods. These rates are tuned to achieve a risk profile aligned with that of the original strategy. This configuration serves as a benchmark, offering insight into the implications of buyback execution behavior. By comparing the original strategy against this benchmark, we can gauge how different behavior affects risk and performance.

To evaluate the performance of each strategy, it is necessary to assess the cost of shares purchased each day. Since implementing an execution model that accounts for market impact falls outside the scope of this study, we must approximate share prices. Where applicable, the daily average purchase price from the original strategy can be used as an analog for share prices in simulated strategies. Alternatively, strategies can be simulated over a range of synthetic market conditions.

3.2.2 Asset Price Modeling using Geometric Brownian Motion

In our methodology, we leverage the Monte Carlo simulation method, incorporating Geometric Brownian Motion (GBM) to simulate the future stock price paths of BNP Paribas. This approach allows us to analyze the risks associated with various share buyback strategies under diverse market conditions. Monte Carlo simulations are recognized for their ability to model the behavior of financial markets, which are inherently uncertain and complex. By employing random sampling techniques, these simulations generate a vast array of possible future outcomes based on probabilistic models [10].

Geometric Brownian Motion (GBM) serves as the foundation for simulating stock price paths in our Monte Carlo framework. GBM is favored in financial modeling for its simplicity and the realistic manner in which it represents stock price dynamics. It operates under the premise that the logarithms of stock price movements conform to a normal distribution. This characteristic ensures that the prices generated through simulation remain non-negative, reflecting the behavior of real-world stock prices. [10]. Geometric Brownian Motion can be mathematically represented as:

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t \tag{1}$$

Here, dS represents the change in stock price, S is the stock price, μ is the expected return, σ is the volatility of the stock returns, dt is a time increment, and dW is standard Brownian motion. This equation models stock price dynamics by incorporating both a directional trend component (μ) and a random shock component (σdW), simulating the inherent unpredictability of real-world markets.

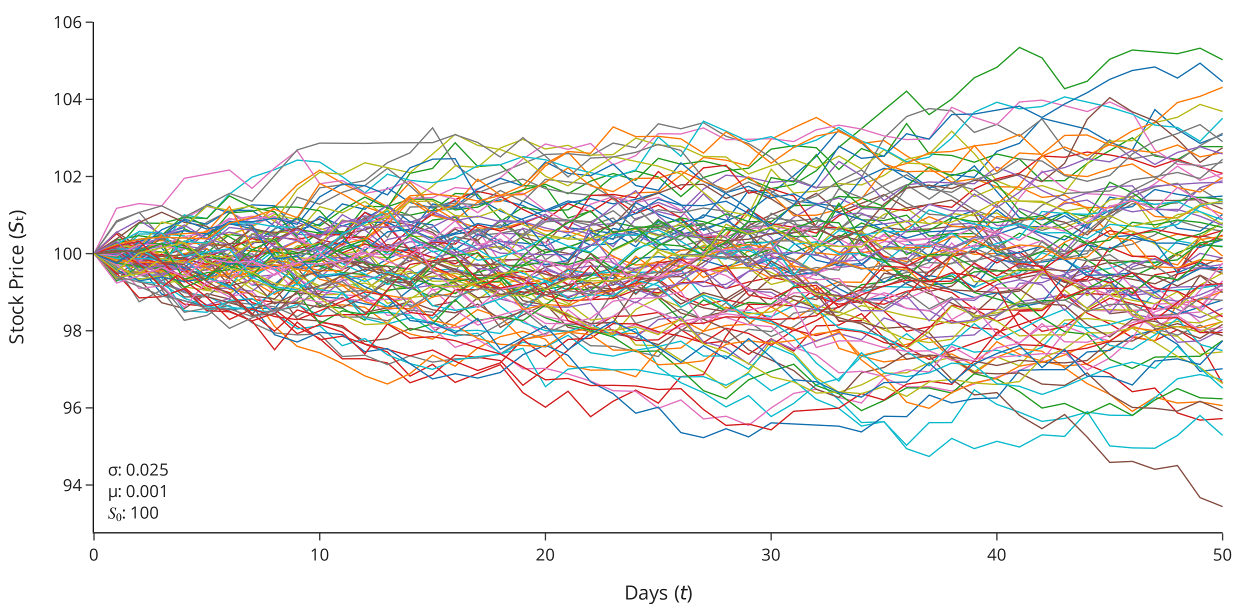

Figure 1: 50 day price simulation using geometric brownian motion.

To assess the risk associated with BNP Paribas's share buyback strategies, we simulate thousands of potential stock price paths using the GBM model within a Monte Carlo framework. As shown in Fig. 1, Each path represents a possible future scenario of stock price evolution over the buyback period. Since the pricing data inherently reflects price changes caused by the original strategy's execution, it cannot be used to reliably assess the performance of alternate strategies. By simulating a range of stock price paths, we can evaluate strategies without the influence of the original strategy's distinct market impact.

3.2.3 Predicting Trade Volume with Machine Learning

In the absence of real-time intraday volume data, we employ several alternative methods to estimate the daily trading volume. Initially, we establish the real daily trading volumes as our baseline. Following this naïve approach, we calculate the 21-day simple and exponential moving averages of these volumes to smooth out short-term fluctuations and identify longer-term trends in the trading activity. This serves as a more reliable basis for comparison and analysis, mitigating the impact of day-to-day volatility.

Following Libman et al. [16], we selected a novel machine learning approach for experimentation; this study employs a Long Short-Term Memory (LSTM) model to forecast the trading volume of BNP Paribas (BNP.PA). The Keras Python library was chosen to assist in the implementation and training of this model [5]. LSTM networks are a type of Recurrent Neural Network (RNN) designed for predicting sequences of data, such as time series data. Unlike in traditional RNNs, LSTM units are equipped with a "memory cell" to selectively regulate the flow of information. Each memory cell contains three gates: an input gate to decide what new information to add to the cell state, a forget gate to determine what information to discard, and an output gate to control what information gets passed along. This architecture allows LSTMs to overcome the vanishing gradient problem that plagues traditional RNNs, making them particularly well suited for long-term financial market predictions [14, 8, 11, 20].

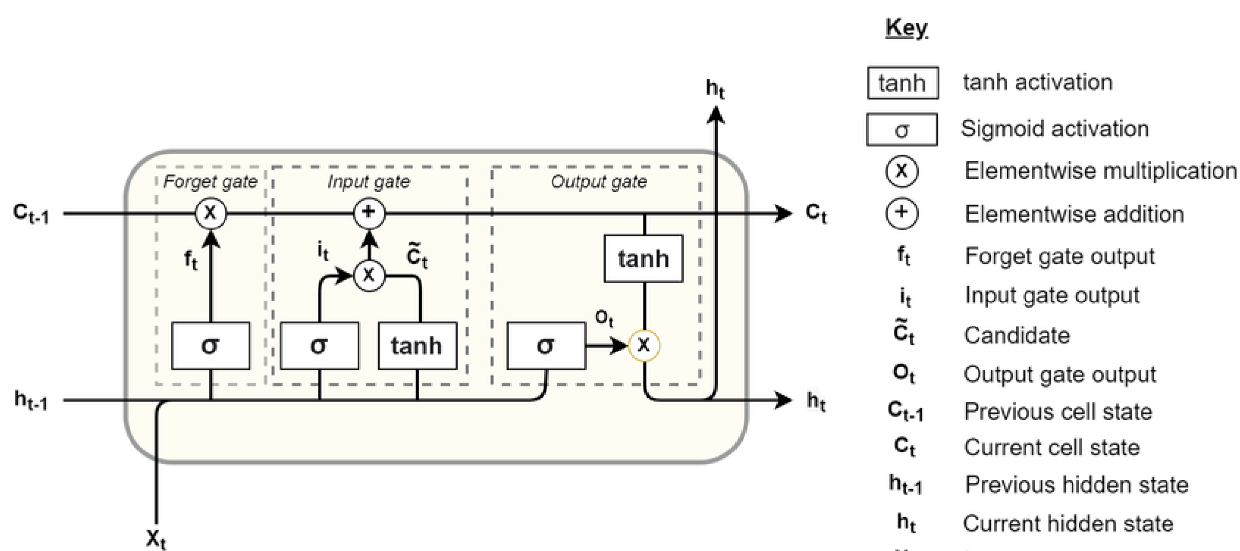

Figure 2: Visualization of the LSTM cell (Adapted from [7]). Licensed under CC BY 4.0.

Historical BNP Paribas trading volume data from Yahoo Finance, covering a three-year period up to the anticipated share buyback, forms the basis of our dataset. This dataset captures a wide range of market conditions, which is crucial for effective model training. To optimize the training process, this data is normalized to a range between 0 and 1 using the MinMaxScaler from the scikit-learn library [26]. This step promotes efficient learning within the neural network by maintaining consistency in the scale of input features. Following normalization, the dataset is segmented into training and validation sets to facilitate model evaluation and mitigate overfitting [6].

The model's architecture, as illustrated in Fig. 3, features two stacked LSTM layers consisting of 50 units each. The number of LSTM units on each layer was chosen to strike a balance between model complexity and computational efficiency. This configuration allows the model to learn from the data at multiple levels of abstraction. The first LSTM layer captures the basic patterns in the data, while the second layer, receiving the first's output as its input, can interpret these patterns in a more complex and nuanced manner. This multi-layer approach enhances the model's ability to identify intricate patterns across different timescales, which is vital for adapting to the dynamic nature of financial markets [25].

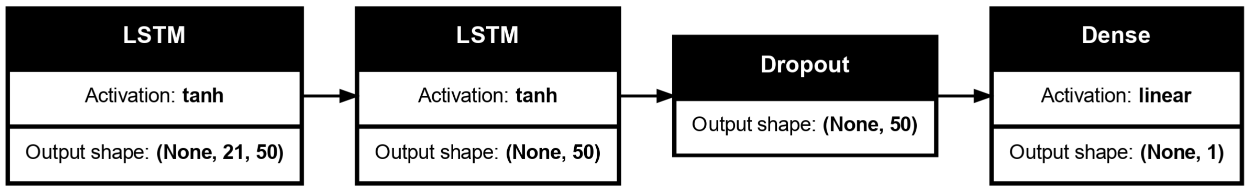

Figure 3: Architecture of the Long Short-Term Memory network.

To mitigate the risk of overfitting—where the model becomes overly attuned to the training data and performs poorly on new data—a dropout layer is introduced following the final LSTM layer. The dropout layer randomly sets a portion of the input units to 0 at each step during training time, which helps prevent the model from becoming overly dependent on any single feature or input pattern. This regularization technique ensures that our model generalizes better to unseen data, enhancing its predictive reliability [27, 6].

We utilize KerasTuner[19] to verify our model's hyperparameters. This tool systematically explores a defined search space to determine settings such as the number of LSTM units per layer, the dropout rate, and the learning rate. Combinations of these hyperparameters are randomly sampled and evaluated to find the optimal configuration. The selections are made to minimize prediction error, evaluated through cross-validation on the training dataset [12].

Model training performance is evaluated and tuned using metrics such as Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE). Described in the following equations, $\hat{y}$ is the predicted value, y is the actual value, and n is the number of samples. The MAE represents the average magnitude of the errors in a set of predictions, without considering their direction. A lower MAE suggests that the model's predictions are closely aligned with the actual trading volumes, indicating high accuracy.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|\hat{y}_i - y_i| \tag{2}$$

The MSE measures the average squared difference between the estimated values and the actual value. A lower MSE value points to a model with high prediction accuracy, and in this case, it indicates that the model's predictions deviate minimally from the actual volumes on average.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{3}$$

The RMSE is widely used in machine learning because it gives a relatively high weight to large errors, although scholars [28, 2] argue it is unintuitive and often unreliable for forecast evaluation.

$$\text{RMSE} = \sqrt{\text{MSE}} \tag{4}$$

The trained LSTM model is then used to forecast the trade volume of BNP Paribas throughout the buyback window. For each day in the program, the model iteratively predicts the trading volume by evaluating the previous month (21 trading days) of volume data. The prediction "hit rate" is used as a preliminary assessment of model performance. The hit rate measures the proportion of instances where the model successfully predicts the direction of volume changes from one day to the next. A high hit rate suggests that the model is adept at capturing underlying market trends, while a lower hit rate would signal a potential need for additional model refinement. Then, the predicted volumes are compared against the actual trading volume to ensure the accuracy is suitable for use in constructing our buyback strategies.

3.3 Evaluation

This section details the evaluation of the simulated strategies, focusing on metrics such as Value at Risk and performance relative to a buyback benchmark. Our analysis aims to contrast the fixed participation rate strategy against BNP Paribas' original program, offering insight into how different execution behavior impacts both performance and risk profiles.

3.3.1 Value at Risk for Risk Assessment

To evaluate the risk associated with our simulated share buyback strategies, we employ Value at Risk (VaR) as our primary metric. VaR offers a probabilistic estimate of the maximum expected loss over a specified time frame at a given confidence level, providing a quantifiable measure of financial risk. This approach allows us to compare the potential risks of each buyback strategy under varying market conditions [10].

Within the context of share buyback strategies, VaR is focused on quantifying the potential financial losses stemming from rising costs of shares yet to be acquired. This interpretation of VaR addresses the risk that the share price may rise before the completion of the buyback program, thereby increasing the cost of acquiring the remaining shares needed to fulfill the buyback target or reducing the overall amount of shares that can be purchased.

To calculate VaR, our methodology utilizes the Variance-Covariance Method. While this approach is favored for its clarity and computational efficiency, it relies on the assumption that stock returns are normally distributed [1].

The formula used to calculate VaR is described in Equation 5. The equation uses the following variables: Q, the number of shares yet to be bought; P, the market value of the stock; σ, the standard deviation; Z, the number of standard deviations from the mean; and T, the time horizon:

$$\text{VaR} = P_t \times Q \times \sigma \times Z \times \sqrt{T} \tag{5}$$

To evaluate each strategy's overall risk exposure, we first calculate the Value at Risk at $t = 0$. The total VaR for a strategy is derived from the length of the program and the overall number of shares acquired. Additionally, we employ a dynamic approach to calculate the VaR on a daily basis throughout the execution period of the buyback program. This approach is designed to capture the evolving risk profile associated with changing market conditions and the program's progress. Each day, VaR is calculated using the current share price, the standard deviation of the stock's returns (σ), the remaining amount of shares to be repurchased, and the number of days left in the program. This allows for a nuanced assessment of risk that reflects both the strategy's timeline and inherent market fluctuations.

The Monte Carlo method provides a powerful alternative for calculating VaR. This method simulates thousands of potential market scenarios, each representing a possible trajectory of the stock's price over the buyback period. By analyzing the distribution of simulated outcomes, we can determine the potential losses at a specified confidence level. This probabilistic approach accounts for the complex interplay of market forces and their impact on the buyback program, offering a more comprehensive view of the risks involved.

3.3.2 Average Purchase Price

To assess the performance of our simulated share buyback strategies, we analyze the average purchase price for the shares acquired. This indicator helps determine how well each strategy maximizes value during the buyback. Calculating this involves dividing the total expenditure on shares by the number of shares bought, providing insight into the cost-efficiency of each approach. Strategies yielding lower average purchase prices suggest superior value creation, as they secure shares at more advantageous prices.

3.3.3 Volume-Weighted Average Price Benchmark

This study compares share buyback strategies against the simple volume-weighted average price buyback benchmark, also referred to as the "Bogus Benchmark," to contextualize the perceived performance of each strategy. The VWAP benchmark is traditionally calculated by averaging the daily VWAP for each day in the buyback period. Daily VWAP refers to the average price at which a stock is traded throughout the day, weighted by the trade volume. For this experiment, the daily VWAP is approximated using the Typical Price. The Typical Price is computed as the mean of the high, low, and closing prices for a stock on a given day:

$$\text{Typical Price} = \frac{High + Low + Close}{3} \tag{6}$$

This indicator is used in place of the daily VWAP because it can be calculated with low level market data freely available from Yahoo Finance [18, 9, 24].

4 Results

This section analyzes BNP Paribas' 2023 share buyback program, focusing on its execution patterns and responsiveness to market conditions. We investigate the Fixed Participation Rate (Percentage of Volume) buyback strategy, comparing its risk profile and performance to BNP Paribas' original approach. Additionally, we evaluate various trading volume forecasting models, including traditional statistical methods and machine learning methods using Long Short-Term Memory networks.

4.1 Original BNP Paribas Buyback Program

According to their 2022 financial report, BNP Paribas launched their 2023 buyback program with the aim of boosting earnings per share in the wake of their sale of Bank of the West, Inc [4]. The original press release outlined their plan to execute the program in two parts, each with a maximum worth of €2.5 billion [3]. Besides a limit of €88 per share, no performance targets were established publicly.

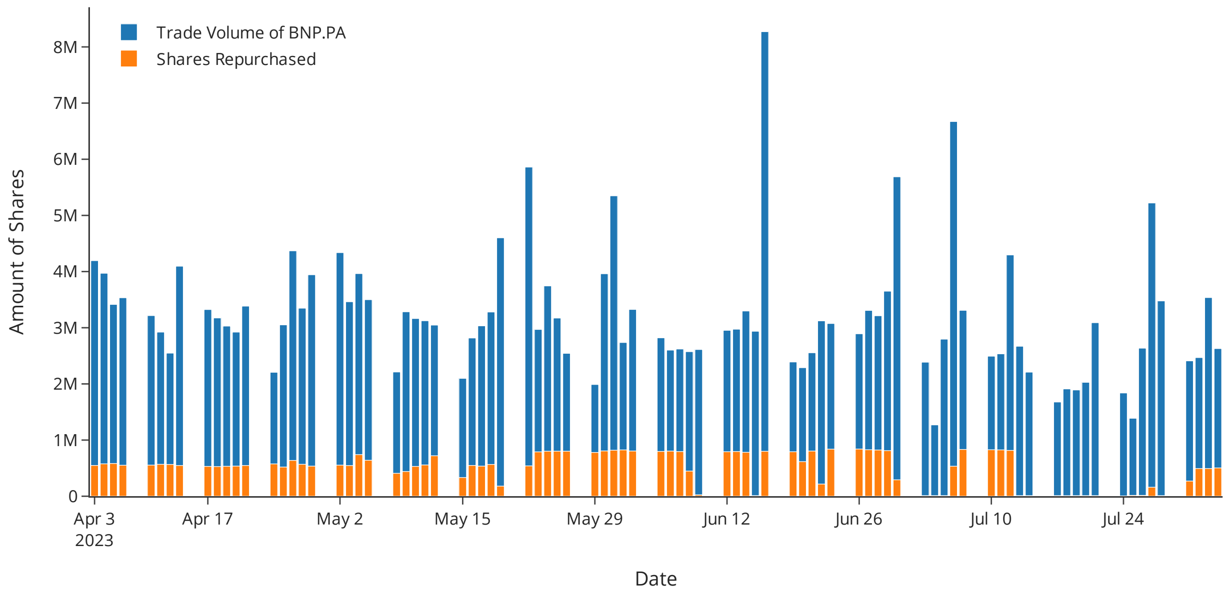

Figure 4: Shares repurchased overlaid on the daily trade volume of BNP.PA

Our analysis focuses on the first tranche of BNP Paribas' buyback program, running from April 3, 2023 to August 3, 2023. During this period, a total of 43,882,757 shares were repurchased for a sum of €2,500,003,409.70. The total cost exceeded the €2.5 billion budget by just €3409.70. These shares were purchased at an average price of €56.97 each.

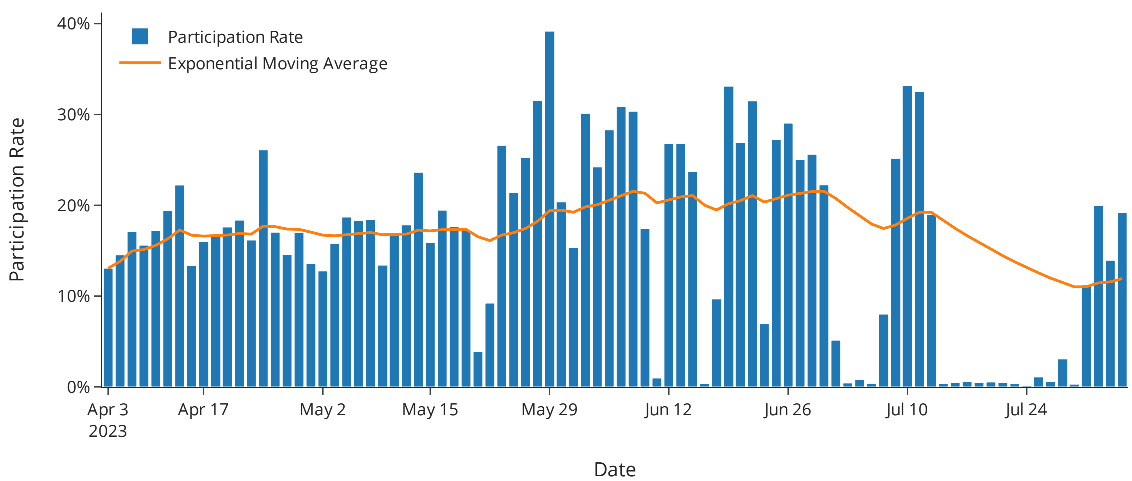

Figure 5: Original strategy participation rate over time with an exponential moving average (EMA).

Over the period analyzed, BNP Paribas had an average participation rate of 16.26%. The participation rate ranged from a minimum of 0.15% to a maximum of 39.17%, with a standard deviation of 10.20%. The correlation between the daily buyback amount and trade volume was 0.27. The high standard deviation and low correlation suggests that buyback targets may not have been adjusted for daily fluctuations in trade volume.

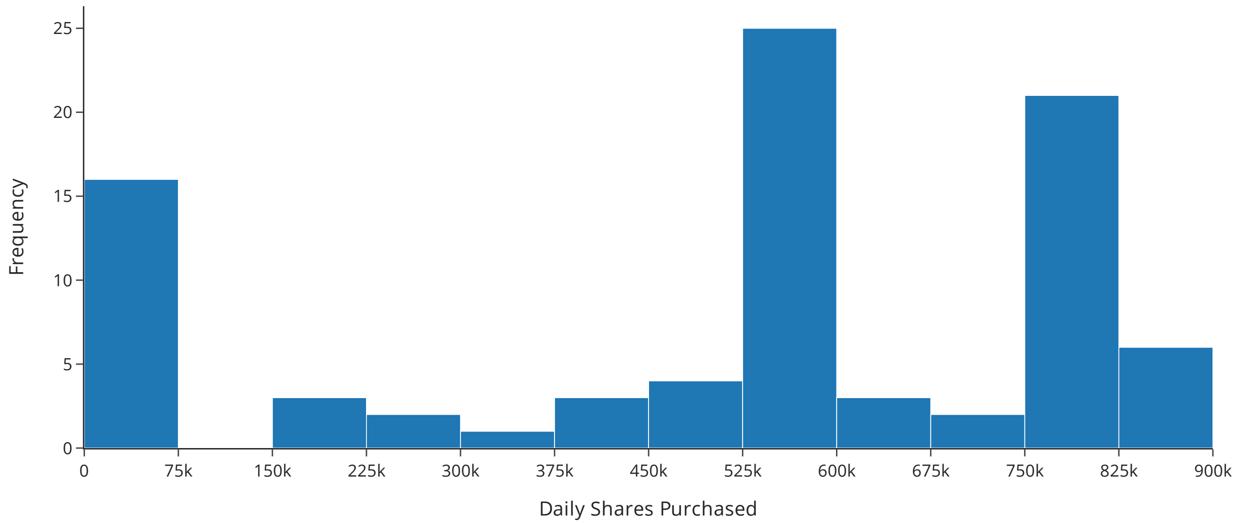

Figure 6: Histogram of daily share purchase volumes during BNP Paribas buyback execution.

Our examination of BNP Paribas' share buyback execution unveils a strategic pattern in the daily volumes of shares purchased. As shown in Fig. 6, the purchase volumes predominantly fall around three distinct levels: 10k, 550k, and 800k. This indicates a structured system of operation that may prioritize consistency over risk or market impact mitigation.

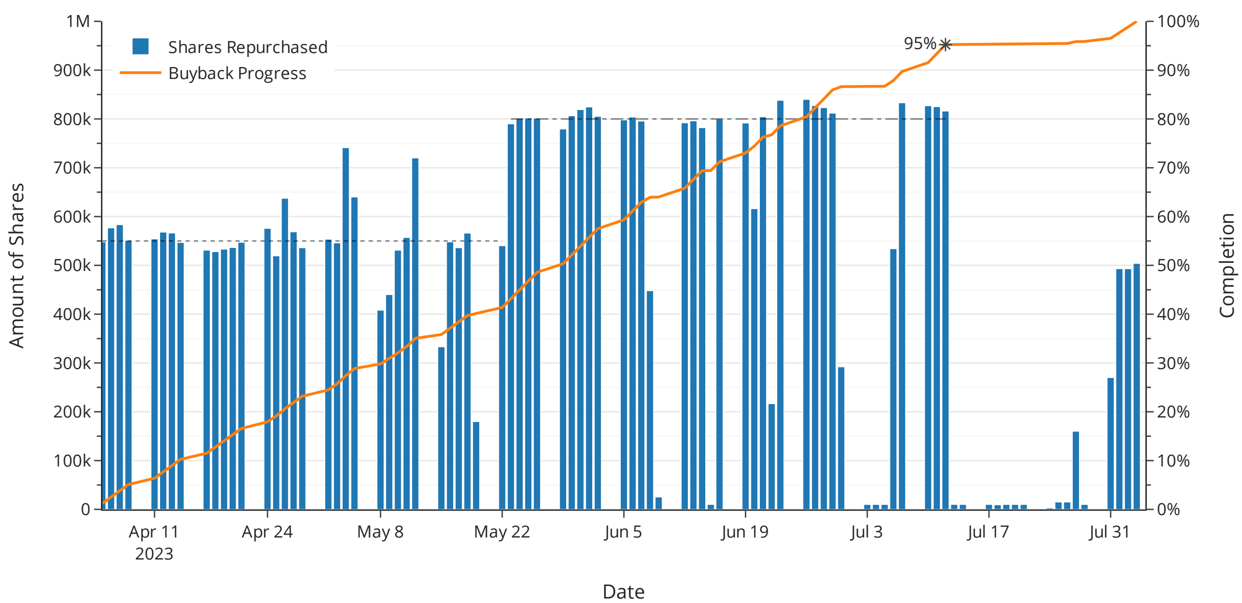

Figure 7: Shares repurchased over time with buyback completion progress.

From 03 April to 22 May 2023, the amount of shares repurchased each day remains around the 550k line. On 23 May 2023, share purchases suddenly increased to approximately 800k. Buyback levels remain at this level until 12 July 2023, with some exceptions. These trends can be seen in Fig. 7.

There was a significant period of time in which few shares were purchased. From 13 to 30 July 2023, only 272,008 shares were purchased with an average participation rate of 0.70%. Notably, 12 July 2023 is when buyback progress, measured as the percentage of the total budget spent, reached 95%. From then on, only about 10,000 shares were purchased each day until the last four days of the program, when the remaining shares were acquired.

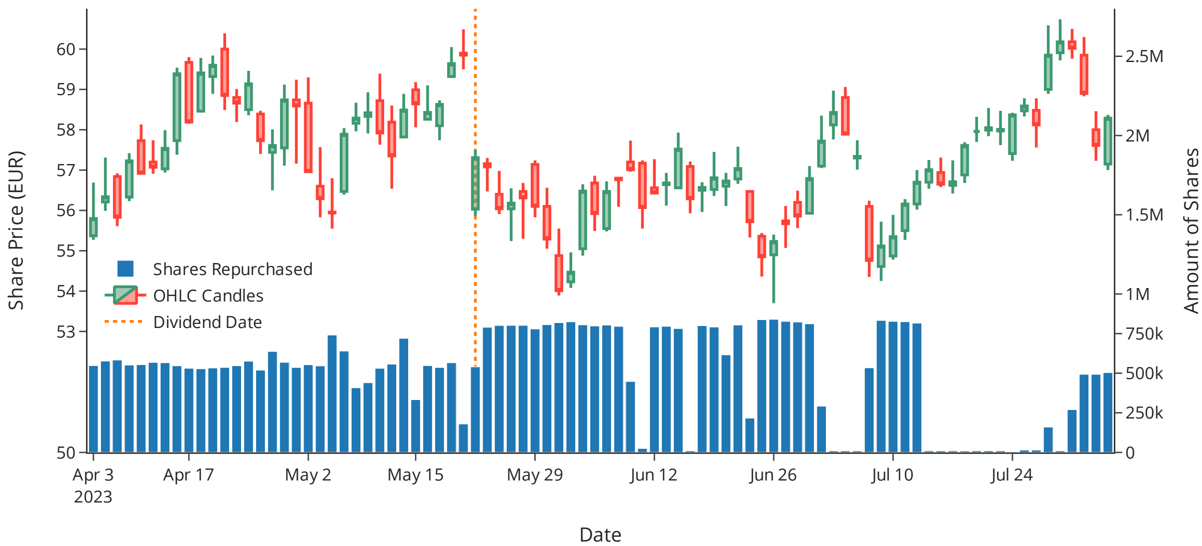

Figure 8: BNP.PA Open-High-Low-Close chart with repurchase volumes.

With the context of BNP.PA share prices, it becomes evident that the shift to increased repurchasing amounts corresponds to a significant drop in share price by 4.24% on 22 May 2023. It is important to note that 22 May 2023 was the ex-dividend date for BNP.PA [29]. Reduced buyback targets in the first phase may have been a response to this upcoming stock event.

Besides the period from 13 to 30 July 2023, there is a brief span from 30 June to 6 July 2023 where repurchasing volumes were significantly reduced. On the three days from 3 to 5 July 2023, only 10,000 shares were purchased each day. During this time, the share price for BNP.PA peaked at €59.06, an increase of 6.61% from the previous week. This reaction to price fluctuations alludes to a degree of temporal optionality in the buyback strategy execution.

	Median	**Min**	**Max**	**Mean Price**
2023/04/03 to 05/22	548,000	179,700	741,000	€58.03
2023/05/23 to 07/12	800,000	10,000	840,000	€56.31
2023/07/13 to 08/03	10,000	2721	503,922	€58.20
Overall	549,650	2721	840,000	€57.32

Table 1: Daily repurchasing volume statistics for BNP Paribas split over three periods, divided at the ex-dividend date and the date overall buyback progress reached 95%.

BNP Paribas' share buyback program follows a structured approach, with distinct activity phases at 550k, 800k, and 10k daily shares (Table 1). While this suggests an intentional strategy, its rigidity could hinder optimal price impact mitigation and risk management. The consistent daily volumes, regardless of prevailing market conditions, limit the strategy's responsiveness to real-time market dynamics. To fully optimize shareholder value while minimizing risk, BNP Paribas might benefit from more adaptable execution strategies.

4.2 Trade Volume Forecasting

This section evaluates the predictive performance of various forecasting models for trading volume predictions. To facilitate direct comparison between traditional and machine learning techniques, loss metrics were calculated using log volumes. We establish a baseline for comparison with the Naïve Forecast, the simplest approach that assumes future volume will replicate the previous day. The performance of each model is evaluated using key metrics: Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE), with results presented in Table 2.

	MAE	**MSE**	**RMSE**
Naïve Forecast	0.2969	0.1586	0.3983
21 Day SMA	0.3004	0.1593	0.3992
21 Day EMA	0.2627	0.1199	0.3463
Single LSTM	0.1515	0.0396	0.1990
Stacked LSTM	0.1285	0.0288	0.1697

Table 2: Volume forecasting performance metrics.

The naïve approach, while basic, sets a benchmark for evaluating more sophisticated techniques. This model yields an MAE of 0.2969, an MSE of 0.1586, and an RMSE of 0.3983 (Table 2). The 21-day simple moving average (SMA) slightly lags behind the naïve forecast, with an MAE of 0.3004, MSE of 0.1593, and an RMSE of 0.3992, indicating minimal improvement in prediction accuracy through simple past volume averaging. Conversely, the 21-day exponential moving average (EMA), which emphasizes more recent data, shows better performance with an MAE of 0.2627, MSE of 0.1199 and an RMSE of 0.3463.

4.2.1 Long Short-Term Memory Predictor

Having established benchmarks with traditional forecasting techniques, we now assess the effectiveness of Long Short-Term Memory (LSTM) networks in predicting trading volume. To demonstrate the benefits of using stacked LSTMs, we first examine the performance of using a single LSTM layer. This model already improves significantly over the conventional techniques, achieving an MAE of 0.1515, MSE of 0.0396, and RMSE of 0.1990. However, the stacked LSTM model achieved the lowest error—an MAE of 0.1285, MSE of 0.0288, and RMSE of 0.1697, as detailed in Table 2. The superior performance of the stacked LSTM model suggests that the multi-layer architecture is more effective at capturing the complex patterns inherent in the trading volume data.

Hyperparameter	**Search Space**	**Selected Value**
Lookback window	5, 10, 21, 30, 42, 50	21
LSTM layers	1, 2, 3	2
Units per layer	10, 25, 40, 50, 75, 100	50
Dropout rate	0.0, 0.1, 0.2, 0.3	0.1
Learning rate	1×10^{-2}, 1×10^{-3}, 1×10^{-4}	1×10^{-3}
Optimizer	Adam, RMSProp, SGD	Adam
Batch size	32, 64, 128	64

Table 3: Hyperparameters selected for the LSTM trade volume prediction model.

After extensive tuning with KerasTuner[19], we identified an optimal combination of hyperparameters to minimize loss. Table 3 provides an overview of the search space and selected values for each hyperparameter, providing a reference for the model's architecture and training configuration.

Layer	**Output Shape**	**Trainable Parameters**
Input	(21, 1)	0
LSTM 1	(21, 50)	10,200
LSTM 2	50	20,200
Dropout	50	0
Dense	1	51
Total		30,451

Table 4: Summary of the LSTM model layers.

The stacked LSTM model contains 30,451 trainable parameters, with the majority concentrated within its LSTM layers. As detailed in Table 4, the two LSTM layers have 10,200 and 20,200 parameters, respectively. The parameter count for each LSTM layer is calculated using the formula:

$$\begin{aligned} |\text{Parameters}_{\text{LSTM}}| &= 4 \times (|\text{input}| + |\text{Units}| + 1) \times |\text{units}| \\ 10{,}200 &= 4 \times (21 + 50 + 1) \times 50 \\ 20{,}200 &= 4 \times (50 + 50 + 1) \times 50 \end{aligned} \tag{7}$$

where |Input| is the input dimension and |Units| is the number of hidden units. The factor of 4 represents the cell state and the three internal gates in the LSTM architecture: input, forget, and output. While the Input and Dropout layers have no trainable parameters, the Dense layer has 51. The Dense layer finalizes the model's predictions by linearly transforming inputs received from the previous layer into the desired output. During the transformation process, each of the 50 input values are multiplied by their corresponding weight. The resulting values are summed together, and then a bias term is added to the sum. This operation condenses the learned features into a single predictive value [12, 6].

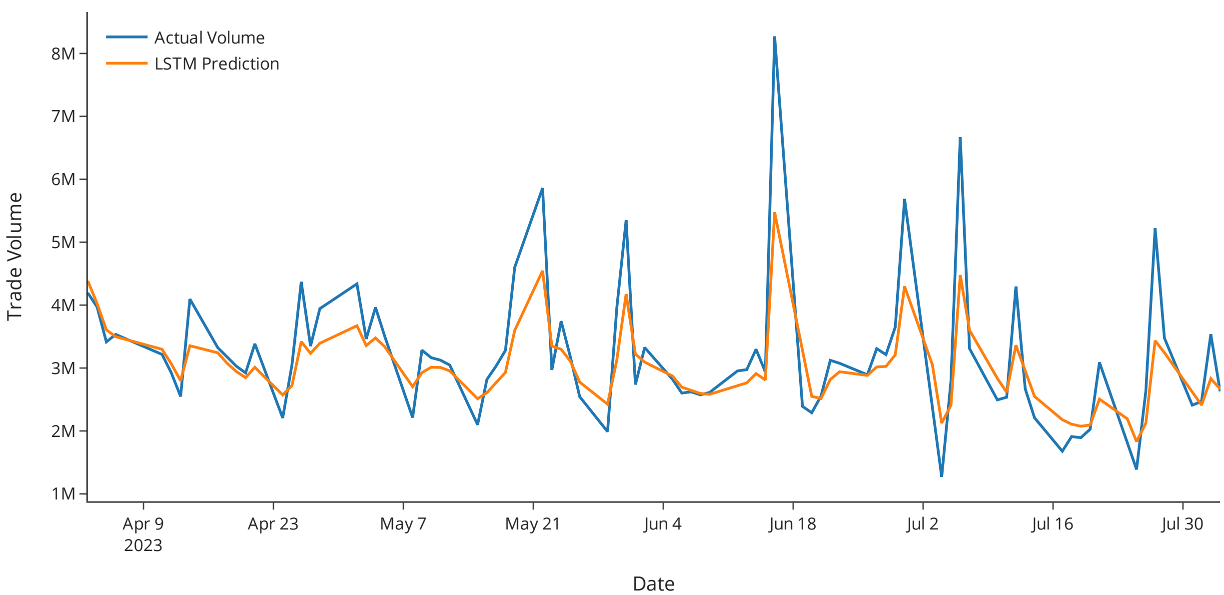

Figure 9: LSTM prediction of trade volume vs. actual volume over the buyback period.

The LSTM model exhibits a promising ability to capture general trends in trading volume, frequently aligning with market activity (Figure 9). Moreover, the model's hit rate of 91.10% demonstrates its reliability in correctly predicting the trend of trading volume shifts. The model's high predictive accuracy indicates its potential value in financial forecasting.

Despite its strengths, the model faces challenges in accurately predicting volume during periods of increased volatility. It tends to underpredict in high-activity phases and overpredict when the market is quieter. This suggests the model may struggle to fully capture the market dynamics leading to extreme fluctuations, possibly due to scaling issues in the data preprocessing phase.

However, while the LSTM model exhibits a commendable ability to trace general market trends and boasts a high hit rate, it encounters difficulties in periods of pronounced market volatility. The model's tendency to underpredict in times of high activity and overpredict during quieter periods raises critical considerations regarding its capacity to fully understand market dynamics. These challenges, potentially rooted in scaling issues during the data preprocessing stage, highlight areas for future enhancement.

Despite the encouraging performance of the LSTM model in forecasting trade volumes, its nuanced performance underscores the need for additional enhancement. Prior to implementing the model in any share buyback strategies, significant effort is required to refine its architecture and address the underprediction issues.

4.3 Fixed Participation Rate Strategy

The design of our simulated share buyback strategy prioritizes alignment with historical execution patterns. We implement a fixed participation rate of 18% based on the median participation rate observed in BNP Paribas' original strategy up to 12 July 2023. This rate is pivotal in aligning the simulated strategy closely with historical execution patterns, ensuring a relevant and comparable analysis.

The 18% strategy's average daily purchase amounts to 509,800 shares, closely mirroring the original strategy's average of 510,265 shares. The similarity in averages indicates that the 18% rate is a reasonable approximation of the original strategy's behavior.

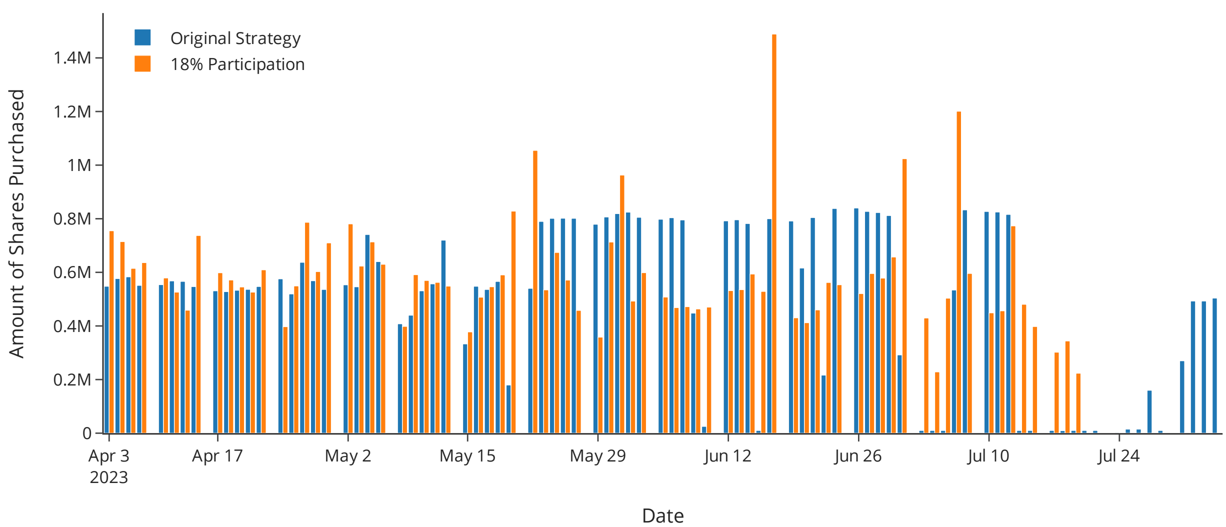

Figure 10: Daily shares purchased, original strategy vs. 18% participation rate strategy.

The comparison of median daily shares purchased reveals a close similarity between the two strategies. The 18% strategy has a median purchase volume of 534,966 shares, slightly lower than the original strategy's 549,650 shares. This difference suggests the 18% strategy closely mirrors the original strategy's overall purchase pattern while demonstrating greater consistency in daily purchase volumes, resulting in less variation around the median.

A notable divergence between the two strategies is observed in their respective end dates. The 18% fixed participation rate strategy concluded on 19 July 2023, whereas the original strategy extended an additional

11 days, ending on 3 August 2023 This difference in duration highlights a key consideration in the implementation of fixed participation rate strategies: the potential for variability in the time frame required to execute the buyback plan fully. The earlier completion date of the 18% strategy may indicate a more efficient capital deployment, assuming it achieves its buyback objectives within a shorter period. However, this expedited timeline could also reflect differences in market conditions or liquidity that could affect the strategy's execution dynamics.

The comparative analysis of the 18% fixed participation rate strategy and the original buyback strategy demonstrates a key advantage: aligning the fixed rate with historical participation rates can achieve similar daily purchase volumes within a shorter timeline. This offers a less complex strategic approach while potentially reducing risk exposure due to a shortened presence in the market.

4.4 Strategy Evaluation

At the outset of our share buyback analysis, the original strategy exhibited a VaR of €617,701,726.70, reflecting the potential market risk under standard market conditions. In comparison, the simulated strategy with a fixed 18% participation rate demonstrated a slightly reduced VaR of €571,083,316.78. This represents a noteworthy difference of €46,618,409.92, equating to a 7.5% reduction in VaR when compared to the original strategy.

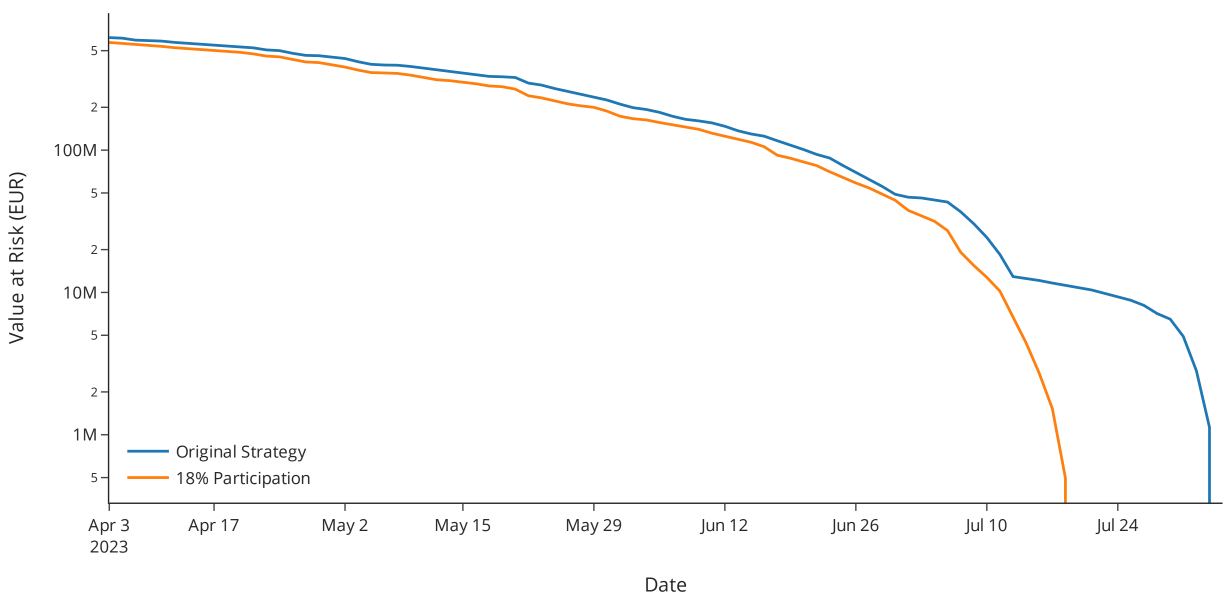

Figure 11: Value at Risk (95%) comparison on a logarithmic scale, original strategy vs. 18% participation rate strategy.

This reduction in VaR suggests that the 18% strategy potentially offers a more risk-averse approach to executing share buybacks, with a lower expected maximum loss under normal market conditions. The significance of this finding cannot be overstated, as it implies that a fixed participation rate strategy not only achieves similar purchase volumes to the original strategy but does so with a reduced exposure to market risk at the outset of the strategy.

4.4.1 Volume-Weighted Average Price

The study utilized Geometric Brownian Motion (GBM) to simulate a variety of share buyback strategies, generating a distribution of Volume-Weighted Average Prices (VWAPs) to assess the potential efficiency and effectiveness of these strategies under different market conditions. The analysis of these simulations revealed a mean VWAP of €56.18, indicating the average cost per share across all simulated strategies. This outcome, coupled with a relatively narrow standard deviation of €0.51, suggests that the majority of simulated buyback strategies would result in purchase prices closely clustered around this average value.

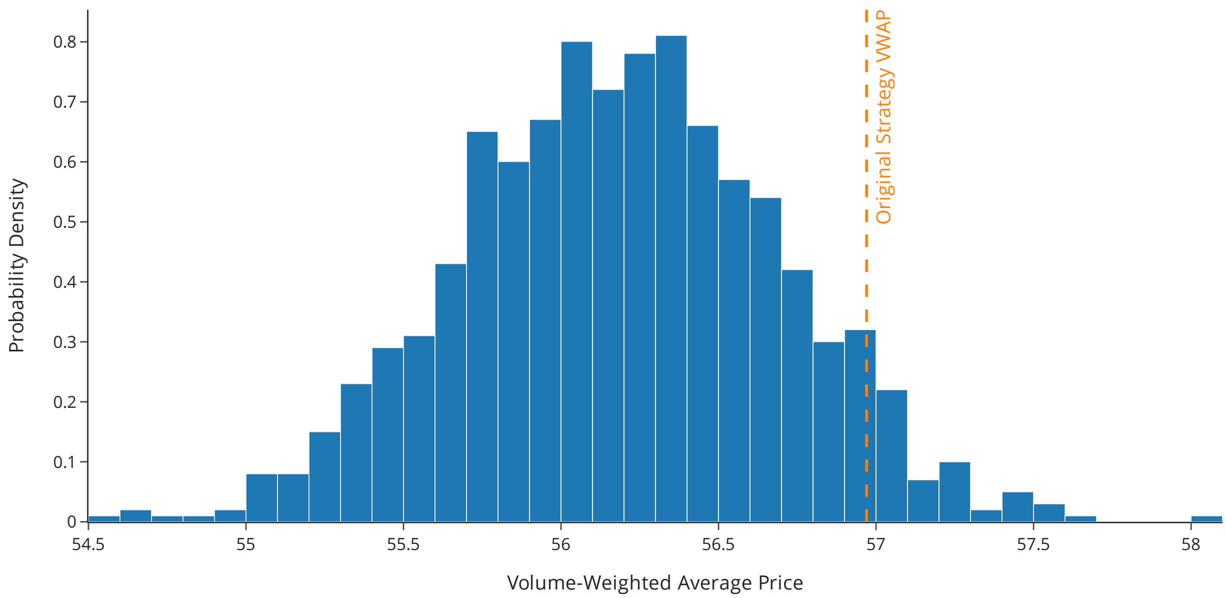

Figure 12: Distribution of simulated VWAPs under GBM, compared to the original strategy's VWAP.

Comparing these results with the original buyback strategy executed by BNP Paribas, which achieved an overall VWAP of € €56.97, indicates that the actual strategy led to a higher average cost per share than the mean of the simulated strategies. This discrepancy highlights the potential for optimizing buyback strategies to secure shares at more favorable prices, thus enhancing shareholder value.

GBM provides a theoretical framework for simulating share price movements, yet its application in analyzing share buyback strategies is constrained by several limitations. GBM operates under simplified assumptions that overlook critical market dynamics such as liquidity and the impact of large trades on share prices. It presupposes markets without friction, where transactions do not affect share prices, a scenario that diverges from the reality of significant buyback operations potentially elevating share prices. Additionally, GBM's foundation on constant volatility and log-normal price distributions might not accurately mirror the market's complexity, where investor sentiment and external factors significantly influence volatility and price trajectories.

Given these limitations, interpreting GBM simulation outcomes for share buybacks necessitates caution. While providing theoretical perspectives on strategy performance, these simulations might not encompass the nuanced interplay of real-world market forces. For instance, the simulation-derived mean VWAP of €56.18 points to possible cost efficiencies that practical constraints related to liquidity and trading impact could limit.

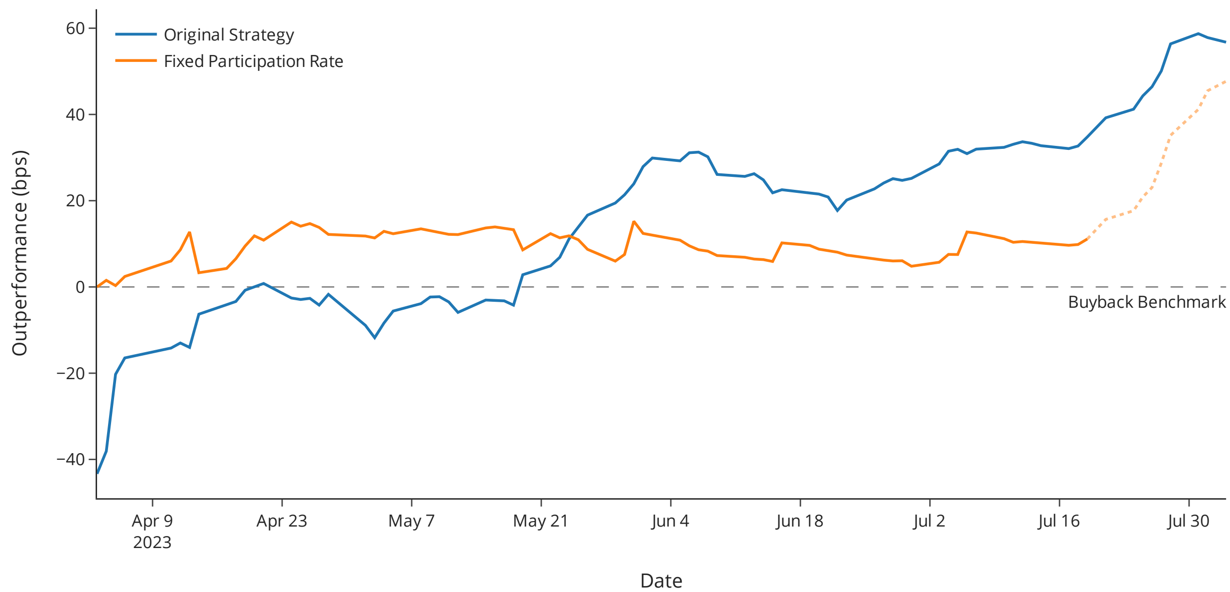

Figure 13: Strategy performance normalized against the buyback benchmark.

Our analysis compares the outperformance of the original share buyback strategy against a fixed participation rate strategy. Outperformance is measured in basis points (bps). This metric is determined by the deviation between the average buyback price and a simple VWAP benchmark.

The original strategy achieved a peak outperformance of 58.73 bps. Conversely, the fixed participation rate strategy demonstrated an outperformance of 15.23 bps. While at first glance, this may appear modest in comparison, it's crucial to contextualize this performance within the operational time frame. Visible in Fig. 13, The fixed participation rate strategy was not executed over as long a period as the original strategy. Adjusting for the duration, a projection indicates that had the fixed participation rate strategy been extended to match the original strategy's timeline, its outperformance could potentially reach 47.69 bps.

Extending the execution period of the fixed participation rate strategy introduces considerations of potential benchmark manipulation. Delaying or strategically timing buybacks could influence the VWAP, thus affecting the calculated outperformance. While this study does not imply intentional manipulation, it highlights the sensitivity of outperformance metrics to the timing and duration of buyback strategies.

5 Discussion

The analysis of the original BNP Paribas buyback strategy reveals that it was not optimally configured for mitigating risk and minimizing market impact. The notable variability in the participation rate underscores the strategy's ability to significantly influence market dynamics, especially in terms of liquidity and price movements. On occasions characterized by high participation rates, the buyback activities could substantially reduce market liquidity. This reduction in liquidity, in turn, could lead to tighter spreads and, potentially, a temporary increase in volatility.

The investigation into the extended duration of the BNP Paribas buyback process suggests a focus on enhancing benchmark performance, potentially at the expense of core objectives like financial stability and market impact reduction. The intrinsic flaws and susceptibility to manipulation of standard VWAP-centric approaches starkly contrast with the principal aims of maximizing shareholder value and executing fiscally prudent repurchases. The conventional buyback benchmark overlooks critical factors like market liquidity, volatility, and the strategic timing of repurchases, all of which are crucial for accurately evaluating the true effectiveness and consequences of buyback operations. Consequently, there is a clear necessity for companies to adopt more comprehensive evaluation methodologies that encompass a broader array of market factors, thereby aligning buyback strategies more effectively with the goals of long-term shareholder value enhancement and risk reduction.

While a fixed participation rate as a buyback strategy demonstrates limited practical applicability due to its inherent simplicity and rigidity, it offers valuable insights as a benchmarking instrument. This approach provides a novel perspective for assessing the efficacy of share repurchase strategies across diverse market conditions, offering insights that traditional methods might overlook.

However, it is essential to contextualize these findings within the limitations of our data. The absence of comprehensive market data means that the observed effects of the buyback on the market could differ if analyzed within a more extensive dataset, including transactions not captured in our current scope.

When assessing the original buyback against a fixed participation benchmark, it seems that extending the buyback duration for manipulative gains yields diminished benefits. This observation encourages a shift toward more stable participation rates in buyback activities, with the aim of minimizing market disruption and reducing price fluctuations.

In conclusion, the findings of this study underscore the need to transition away from strategies favoring ephemeral gains towards those fostering market stability. Implementing benchmarks that incentivize more uniform participation rates in buyback programs not only minimizes the potential for market manipulation but also aligns the buyback process with the overarching goals of shareholder value enhancement and market integrity.

5.1 Improving the Trading Volume Prediction Model

The exploration of Long Short-Term Memory (LSTM) networks for predicting trading volumes represents a promising avenue within the broader landscape of financial market analysis. Based on our initial findings, we propose several directions for future work to enhance the predictive accuracy, generalizability, and real-world applicability of our model. These proposals aim to refine the model's architecture, enrich its input features, and diversify its training data, thereby offering deeper insights into market dynamics and improving its utility for investors and analysts.

An exciting direction for extending this research involves the incorporation of sophisticated neural network architectures such as Bidirectional LSTMs (Bi-LSTMs) and Gated Recurrent Units (GRUs). Bi-LSTMs, which process data in both forward and backward directions, can provide a more comprehensive understanding of temporal dependencies, potentially uncovering patterns missed by traditional LSTMs. Similarly, GRUs, known for their efficiency and performance in handling sequence data, could offer a streamlined alternative for capturing temporal dynamics with fewer parameters, reducing the risk of overfitting [11].

Expanding the model's input features presents another promising avenue for enhancing its predictive performance. Incorporating additional data points such as price volatility, market sentiment indices, and calendar information, including option expiry and ex-dividend dates, could provide the model with a richer context for forecasting volume. These features could help the model account for external factors that influence trading activity, thereby improving its accuracy and relevance for market analysis.

To bolster the model's robustness and applicability to various market conditions, we plan to extend the training dataset beyond the initial focus on BNP.PA data. By training the model on a more diverse set of financial instruments, including stocks from different sectors and regions represented in indices like the S&P 500, we can enhance its generalizability. This expanded dataset will allow the model to learn from a broader spectrum of market behaviors, potentially uncovering universal patterns and idiosyncratic differences across different asset classes and markets.

Acknowledgments

Thank you to Professor Jörg Osterrieder and the Bern University of Applied Sciences for sponsoring this project and hosting me on their campus. I would like to acknowledge Gabin Taibi from BFH for his significant contributions. His insight and support was invaluable. Additionally, I thank Professors Marcel Blais, Stephan Sturm, and Randy Paffenroth for their guidance.

Portions of this research paper were revised with the assistance of generative AI tools, specifically ChatGPT and Gemini. These tools were utilized for editing suggestions, grammar refinement, clarity of expression, brainstorming support, and the refinement of complex ideas.

I would like to extend a special thank you to Ashley for her unwavering support and encouragement throughout the writing process.

References

[1] Robert Almgren and Neil Chriss. "Optimal execution of portfolio transactions". In: *The Journal of Risk* 3.2 (Jan. 2001), pp. 5–39. ISSN: 14651211. DOI: 10.21314/JOR.2001.041. URL: http://www.risk.net/journal-of-risk/technical-paper/2161150/optimal-execution-portfolio-transactions.

[2] J.Scott Armstrong and Fred Collopy. "Error measures for generalizing about forecasting methods: Empirical comparisons". In: *International Journal of Forecasting* 8.1 (June 1992), pp. 69–80. ISSN: 0169-2070. DOI: 10.1016/0169-2070(92)90008-W. URL: https://www.sciencedirect.com/science/article/pii/016920709290008W.

[3] BNP Paribas. *BNP Paribas launches the first tranche of the share buyback programme planned for 2023*. Press release. Mar. 2023. URL: https://invest.bnpparibas/en/document/bnp-paribas-launches-the-first-tranche-of-the-share-buyback-programme-planned-for-2023.

[4] BNP Paribas. *Universal Registration Document and annual financial report 2022*. Tech. rep. May 2023, p. 154. URL: https://invest.bnpparibas/en/groupe-de-document/agm-16-may-2023-documents.

[5] François Chollet et al. *Keras*. 2015. URL: https://keras.io.

[6] François Chollet. *Deep learning with Python*. Second edition. OCLC: on1289290141. Shelter Island: Manning Publications, 2021. ISBN: 978-1-61729-686-4.

[7] Ayman Elhalwagy and Tatiana Kalganova. "Multi-Channel LSTM-Capsule Autoencoder Network for Anomaly Detection on Multivariate Data". en. In: *Applied Sciences* 12.22 (Nov. 2022), p. 11393. ISSN: 2076-3417. DOI: 10.3390/app122211393. URL: https://www.mdpi.com/2076-3417/12/22/11393 (visited on 03/14/2024).

[8] Thomas Fischer and Christopher Krauss. "Deep learning with long short-term memory networks for financial market predictions". In: *European Journal of Operational Research* 270.2 (Oct. 2018), pp. 654–669. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2017.11.054. URL: https://www.sciencedirect.com/science/article/pii/S0377221717310652.

[9] Thierry Foucault, Marco Pagano, and Ailsa Röell. *Market Liquidity: Theory, Evidence, and Policy*. New York: Oxford University Press, Incorporated, 2013. ISBN: 978-0-19-993625-0. URL: http://ebookcentral.proquest.com/lib/wpi/detail.action?docID=5745623.

[10] Paul Glasserman. *Monte Carlo methods in financial engineering*. eng. OCLC: 940691451. New York: Springer, 2010. ISBN: 978-1-4419-1822-2.

[11] Alex Graves. *Generating Sequences With Recurrent Neural Networks*. arXiv:1308.0850 [cs]. June 2014. DOI: 10.48550/arXiv.1308.0850. URL: http://arxiv.org/abs/1308.0850.

[12] Klaus Greff et al. "LSTM: A Search Space Odyssey". In: (2015). Publisher: [object Object] Version Number: 2. DOI: 10.48550/ARXIV.1503.04069. URL: https://arxiv.org/abs/1503.04069.

[13] Olivier Guéant. "Execution and block trade pricing with optimal constant rate of participation". In: (2012). Publisher: [object Object] Version Number: 3. DOI: 10.48550/ARXIV.1210.7608. URL: https://arxiv.org/abs/1210.7608 (visited on 03/16/2024).

[14] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory". en. In: *Neural Computation* 9.8 (Nov. 1997), pp. 1735–1780.

[15] Kif Leswing. *Apple's rise to $3 trillion market cap shows the value of its massive share buybacks*. en. Jan. 2022. URL: https://www.cnbc.com/2022/01/03/apples-3-trillion-market-cap-shows-value-of-share-buybacks-dividend.html (visited on 01/29/2024).

[16] Daniel Libman, Simi Haber, and Mary Schaps. "Volume Prediction With Neural Networks". In: *Frontiers in Artificial Intelligence* 2 (2019). ISSN: 2624-8212. URL: https://www.frontiersin.org/articles/10.3389/frai.2019.00021 (visited on 02/18/2024).

[17] Brooke Masters. "If companies are going to buy back shares, they should pay a fair price". In: *Financial Times* (July 2023). URL: https://www.ft.com/content/5303e9a3-603d-4621-88e3-24f07c87fc69 (visited on 01/29/2024).

[18] John J. Murphy. *Technical analysis of the financial markets: a comprehensive guide to trading methods and applications*. eng. New York: New York Inst. of Finance, 1999. ISBN: 978-0-7352-0066-1.

[19] Tom O'Malley et al. *KerasTuner*. 2019. URL: https://github.com/keras-team/keras-tuner.

[20] Christopher Olah. *Understanding LSTM Networks*. Aug. 2015. URL: http://colah.github.io/posts/2015-08-Understanding-LSTMs/.

[21] Joerg Osterrieder and Michael Seigne. “A Free Lunch for Share Buybacks”. en. In: *SSRN Electronic Journal* (2023). ISSN: 1556-5068. DOI: 10.2139/ssrn.4539767. URL: https://www.ssrn.com/abstract=4539767 (visited on 01/29/2024).

[22] Joerg Osterrieder and Michael Seigne. “Examining share repurchase executions: insights and synthesis from the existing literature”. In: *Frontiers in Applied Mathematics and Statistics* 9 (Nov. 2023), p. 1265254. ISSN: 2297-4687. DOI: 10.3389/fams.2023.1265254. URL: https://www.frontiersin.org/articles/10.3389/fams.2023.1265254/full (visited on 01/29/2024).

[23] Joerg Osterrieder and Michael Seigne. “Temporal Optionality in Share Buyback Execution: An Empirical Anomaly and Value Optimization Approach”. en. In: (July 2023). DOI: 10.2139/ssrn.4525781. URL: https://papers.ssrn.com/abstract=4525781 (visited on 02/09/2024).

[24] Joerg Osterrieder and Michael Seigne. “The Great Deception: A Comprehensive Study of Execution Strategies in Corporate Share Buy-Backs - Decoding the Enigma of Equity Repurchases and their Disproportionate Cost Structures”. en. In: *SSRN Electronic Journal* (2023). ISSN: 1556-5068. DOI: 10.2139/ssrn.4499366. URL: https://www.ssrn.com/abstract=4499366 (visited on 01/29/2024).

[25] Razvan Pascanu et al. *How to Construct Deep Recurrent Neural Networks*. arXiv:1312.6026 [cs, stat]. Apr. 2014. DOI: 10.48550/arXiv.1312.6026. URL: http://arxiv.org/abs/1312.6026.

[26] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

[27] Nitish Srivastava et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: http://jmlr.org/papers/v15/srivastava14a.html.

[28] Cort J. Willmott and Kenji Matsuura. “Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance”. In: *Climate Research* 30.1 (2005). Publisher: Inter-Research Science Center, pp. 79–82. ISSN: 0936577X, 16161572. URL: http://www.jstor.org/stable/24869236 (visited on 03/10/2024).

[29] Yahoo Finance. *BNP Paribas Historical Data*. en-US. Stock Ticker: BNP.PA. URL: https://finance.yahoo.com/quote/BNP.PA/history/ (visited on 01/27/2024).