KTH ROYAL INSTITUTE OF TECHNOLOGY

Degree Project in Technology

First cycle, 15 credits

A Markovian Approach to Financial Market Forecasting

Examining the predictive ability of a Markov-based trading strategy

WILLIAM BORIN & KEVIN SUN WANG

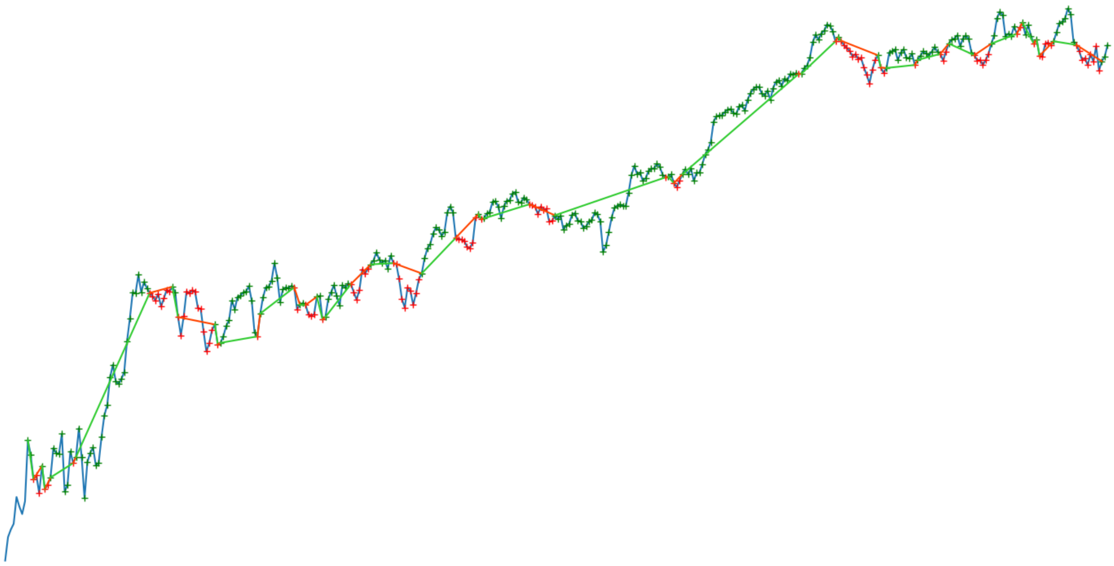

Stockholm, Sweden 2023

Authors

William Borin <wborin@kth.se>
Kevin Sun Wang <kevinsw@kth.se>
Industrial Engineering and Management - Applied Mathematics
KTH Royal Institute of Technology

Supervisor

Björn Wehlin
KTH Royal Institute of Technology

Abstract

This thesis aims to investigate the feasibility of using a Markovian approach to forecast short-term stock market movements. To assist traders in making sound trading decisions, this study proposes a Markovian model using a selection of the latest closing prices. Assuming that each time step in the one-minute time frame of the stock market is stochastically independent, the model eliminates the impact of fundamental analysis and creates a feasible Markov model. The model treats the stock price's movement as entirely randomly generated, which allows for a more simplified model that can be implemented with ease. The model is intended to serve as a starting ground for more advanced technical trading strategies and act as useful guidance for a short-term trader when combined with other resources. The creation of the model involves Laplace smoothing to ensure there are no zero-probabilities and calculating the steady-state probability vector of the smoothed matrix to determine the predicted direction of the next time step. The model will reset daily, reducing the impact of fundamental factors occurring outside trading hours and reducing the risk of carrying over bias from previous trading day. Any open positions will hence be closed at the end of the day. The study's purpose is to research and test if a simple forecasting model based on Markov chains can serve as a useful tool for forecasting stock prices at short time intervals. The result of the study shows that a Markov-based trading strategy is more profitable than a simple buy-and-hold strategy and that the prediction accuracy of the Markov model is relatively high.

Keywords

Markov chain, Markov model, stock market prediction, Laplace smoothing, steady-state, forecasting, trading strategy, stochastic, trading algorithm

Sammanfattning

Denna avhandling syftar till att undersöka möjligheten att använda en markovisk metod för att förutsäga kortsiktiga rörelser på aktiemarknaden. För att hjälpa aktörer på aktiemarknaden att fatta välgrundade handelsbeslut föreslår denna studie en markovisk modell för att förutsäga nästa stängningspris baserat på de senaste stängningspriserna. Modellen antar att varje tidssteg i ett en-minuts intervall på aktiemarknaden är stokastiskt oberoende, vilket eliminerar påverkan från fundamental analys och skapar förutsättningen för en genomförbar markov-modell. Modellen behandlar aktieprisets rörelse som helt slumpmässigt genererat, vilket möjliggör en mer förenklad modell som kan implementeras på marknaden. Modellen är avsedd att tjäna som en utgångspunkt för mer avancerade tekniska handelsalgoritmer och fungera som en användbar vägledning för en akitehandlare med kort tidshorisont i kombination med andra resurser. Skapandet av modellen inkluderar använding av Laplace-jämning för att säkerställa att det inte finns nollsannolikheter samt beräknandet av den stationära sannolikhetsvektorn för den jämnade matrisen i syfte att bestämma den förutsedda riktningen för nästa tidssteg. Modellen kommer att återställas dagligen, vilket minskar påverkan från de fundamentala faktorer som inträffar utanför handelstiderna och ser till att bias inte överförs till nästa börsdag. Detta innebär att alla öppna positioner stängs vid dagens slut. Studiens syfte är att forska och testa om en enkel prognosmodell baserad på Markovkedjor kan vara användbar som ett verktyg för att förutsäga aktiepriser vid korta tidsintervall. Resultatet från studien visar på att en markov-baserad trading strategi är mer lönsam än en enkel köp-och-behåll strategi och att prediktionernas träffsäkerhet från en markov modell är relativt höga.

Nyckelord

Markovkedjor, Markovmodell, prediktion, Laplace-jämning, stationär fördelning, tradingstrategi, stokastisk, trading algoritm

Acknowledgements

We would like to thank our supervisor Björn Wehlin for all the feedback and input which has made this work possible. Special thanks to Lynx Asset Management for providing codes which elevated our trading strategy evaluation. Thanks to Anna Folkesson and Max Aronsson for their peer reviews which significantly improved readability.

Contents

Chapter 1

Introduction

1.1 Background

During the period spanning from 2018 to 2022, the New York Stock Exchange (NYSE), National Association of Securities Dealers Automated Quotations (NASDAQ), and Chicago Board Options Exchange (CBOE), which are recognized as the primary U.S. equities market operators, collectively generated an average monthly turnover of approximately 6.9 trillion U.S. dollars [4]. This figure translates to an estimated daily average turnover of around 330 billion U.S. dollars, considering a typical business month consisting of 21 trading days. It is therefore apparent that there is a high potential upside for anyone with funds (professional and amateur traders alike), to attempt to make a living using stock market movements. However, the high appeal does not come without risks, amongst which is the risk of trading against the market movement which could lead to financial losses. To assist traders in making sound trading decisions, this study attempt to forecast the stock market using a Markovian approach.

In this study each time step in a one-minute time frame of the stock market is assumed to be stochastically independent. With this assumption the influence of fundamental analysis is mitigated and thus allows for the creation of a feasible Markov model. This model aims to predict the closing stock price in the subsequent time step, utilizing a subset of the most recent closing stock prices. The size of the subset will depend on the market in which the strategy is used. To ensure that the model stays relevant to the user, the subset will continuously be updated for each time step, forming a sliding window that encompasses the most recent time steps.

For increased clarity, the term “model” is referring to the Markov model which includes data processing and the generation of predictions, the term “algorithm” refers to the execution of the trades given the predictions from the “model”, and the term “strategy” refers to the both the “model” and “algorithm” combined.

1.2 Problem

The problem is briefly summarised as follows: *To what extent can a discrete time Markov model forecast and predict future stock prices?* With this approach any explainable stock price movements may be considered noise, as it deviates from a purely random behaviour. Although this is a naive assumption, it allows for the creation of a more simplified model which could easily be implemented by professional and amateur alike.

1.3 Purpose and goal

The purpose of this thesis is to research and test if a simple forecasting model based on Markov chains can be useful as a tool for forecasting stock prices at short time intervals. This also includes testing if it is profitable to assume that the stock market behaves stochastically when analysing the stock price at a shorter time interval and whether the model can predict the stock prices at an accuracy higher than 0.5 (i.e. 50%) (0.5 is set as a benchmark since this is the accuracy that would be expected by pure chance). The model will thus act as useful guidance for an intraday trader when combined with other resources.

1.4 Delimitations

In this study it has been assumed that the stock market at a one-minute timeframe behaves stochastically and additionally set 0.5 (i.e. 50%) as the benchmark and the likelihood that a stock trades either up or down.

Additionally, the threshold for defining a significant increase or decrease has been set at 0.03%. This specific value has been selected arbitrarily to classify “high” as an infrequent occurrence, thereby implying that the majority of the data will fall under “increase” or “decrease”. The chosen cutoff value serves as a proof of concept and can be adjusted according to individual preferences.

1.5 Related Work

In a study conducted by Dar, Padi and Rekha, a first-order discrete time Markov Chain model was applied to historical stock prices of a stock traded on the National Stock Exchange of India Limited (NSE). The purpose of the paper was to analyse and estimate the precision of using a Markov chain model to forecast future stock prices, daily closing prices was used and five states (high gain, low gain, no gain, low loss, and high loss) was incorporated in the Markov model. The outcome of the research involved determining the steady-state distribution based on a historical data set spanning 518 days. The conclusion indicated that the obtained results could be beneficial in supporting future investors and shareholders with effective portfolio management [3]. Although the study conducted by Dar et al. differed in terms of goals and objectives from the current paper, it demonstrated the implementation of a Markov model to compare different stocks using a stationary distribution.

In another study more closely related to this paper, Klacksell and Sundberg examined the applicability of a Markov model with various window settings and ordered Markov Chains to predict the stock market. The study yielded less favorable results, with accuracy fluctuating around 0.5 across different settings and often falling below this threshold [8]. The data used in the study encompassed daily and weekly price data from five different stocks. Therefore, previous attempts to utilise Markov chains for stock market prediction have displayed varying levels of complexity with limited success.

This paper differs from other studies by utilising a smaller time step of one-minute frames, in contrast to previous studies that employed daily or weekly time frames. By employing a smaller time step, the hypothesis is that the market will exhibit more stochastic behavior, thereby mitigating the influences of fundamental analysis.

1.6 Outline

In Chapter 2 a brief theoretical background is presented. It is assumed that the reader already has previous knowledge of statistics, Markov processes and the relevant financial indicators, it is thus meant as a brief reminder of the most relevant topics and will not cover the theories in-depth. Chapter 3 presents the method and the implementation of the trading strategy while Chapters 4-6 covers the obtained results, a discussion of said results and finally the conclusion of this study.

Chapter 2

Theoretical Background

2.1 Markov Property

A sequence of random variables X_n is called a Markov chain if it has the Markov property:

$$P\left(X_k = i \mid X_{k-1} = j, X_{k-2}, \cdots, X_1\right) = P\left(X_k = i \mid X_{k-1} = j\right), \tag{2.1}$$

i.e. a stochastic process which is memoryless in that it only depends on the previous outcome.

It is possible to define a Markov chain with finite memory. A Markov chain where the future state is dependent on the previous m states is defined as follows:

$$\begin{aligned} &P\left(X_n = x_n \mid X_{n-1} = x_{n-1}, X_{n-2} = x_{n-2}, \ldots, X_1 = x_1\right) \\ =&P\left(X_n = x_n \mid X_{n-1} = x_{n-1}, X_{n-2} = x_{n-2}, \ldots, X_{n-m} = x_{n-m}\right) \text{ for } n > m \end{aligned} \tag{2.2}$$

2.2 Transition Matrix

The transition matrix is a square matrix ($n \times n$, where n is the number of states) that describes the transitions in the Markov chain. The entries in the transition matrix represent the probability of moving from one state to another in one time step, each of the entries in the matrix is thus non-negative. If the probability of transitioning from state i to j in one step is p_{ij} then the matrix P is given by

$$P = \begin{bmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,n} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n,1} & p_{n,2} & \cdots & p_{n,n} \end{bmatrix} \tag{2.3}$$

The sum of each rows in the matrix is

$$\sum_{j=1}^{n} p_{ij} = 1 \tag{2.4}$$

since the total probability of transitioning from a state i to all other states (including i) must be 1.

2.3 Absorbing state

A state i is absorbing if it is not possible to transition out of this state once it has been entered. The possible absorbing states which could occur in this study is if a stock would reach zero (i.e. bankrupt) or if the stock stops being trading on the market. For simplicity this study will neglect these events, thus the assumption is that the model trades on a market which is not prone to bankruptcy and that the market will be traded.

2.4 Steady state vector

A steady state vector, or the stationary distribution, is the long term probability that the system will be in each state, the steady state vector will hence not change from one time step to the next. The steady state vector π contains entries which are non-zero and sum to 1. Only irreducible and aperiodic Markov chains converge to a unique steady-state probability π.

$$\pi P = \pi \tag{2.5}$$

2.4.1 Irreducibility

A Markov chain is called irreducible if it is possible to transition from one state to any other state in a finite number of steps, hence there are no absorbing states in an irreducible Markov chain.

2.4.2 Periodicity

In this study the Markov chain will be irreducible since it is possible to reach any state from every other state. A state i is periodic with period k if k is the greatest common divisor by which i can be reached from i.

$$k = \gcd\left\{n > 0 : P\left(X_n = i \mid X_0 = i\right) > 0\right\} \tag{2.6}$$

The Markov chain in this paper will be irreducible as well as aperiodic meaning $k = 1$ since it is possible that the price of a stock stays the same in the next transition.

2.5 Laplace smoothing

Laplace smoothing, also known as additive smoothing, is a technique to smooth categorical data and can be useful in cases where zero-probability occurs. Additionally, this smoothing serves to regularise the model, hence reduces the risk of overfitting. For this model an add-1 smoothing is applied as follows:

$$\hat{P}_{ij} = \frac{c_{ij} + 1}{c_i + d} \quad (i = 1, \ldots, d) \tag{2.7}$$

where,
$\hat{P}_{ij} =$ Smoothed probability
$c_{ij} =$ Number of times i has transitioned to j
$c_i =$ Number of i events
$d =$ Number of states

2.6 Financial indicators

Financial indicators are metrics that can be used to measure and analyse the trading strategy from a financial perspective. The relevant indicators for this study are listed below.

2.6.1 Sharpe Ratio

The Sharpe Ratio is a measure of an investment's risk-adjusted performance. The indicator penalises high volatility since this could imply that a high return is based on luck and high risk instead of a stable portfolio. A higher Sharpe ratio is preferred [7].

$$\text{Sharpe Ratio } = \frac{R_p - R_f}{\sigma_p} \tag{2.8}$$

where:
$R_p =$ return of portfolio
$R_f =$ risk-free rate (The theoretical rate of return of an investment with zero risk)
$\sigma_p =$ standard deviation of the portfolio's excess return relative to the risk-free rate

2.6.2 Standard deviation

The standard deviation is a measure of the variation of a data set relative to its mean, the data set is in this study the returns of the trading strategy. In the evaluation of the trading strategy the sample standard deviation normalised by $N - 1$ is used.

$$\hat{\sigma} = \sqrt{\frac{\sum_{i=1}^{N} (x_i - \mu)^2}{N - 1}} \tag{2.9}$$

$\sigma =$ sample standard deviation
$N =$ the size of the sample
$x_i =$ each value from the sample
$\mu =$ the sample mean

2.6.3 Skewness

Skewness is a measurement of the asymmetry observed in a probability distribution. Investors take into account the right-skewedness of a return distribution, as it, similarly to excess kurtosis, provides a more accurate representation of the data set's extreme values [2]. It is desirable for the returns of the strategy to be positively skewed as it indicates that the strategy's smaller and frequent losses can be covered by fewer but larger profits/gains [9]. From a psychological standpoint, the positively skewed returns also eases the loss aversion bias that a trader may experience as frequent smaller losses are more tolerable that infrequent large losses [6]. The skewness value in Chapter 4 is the unbiased skew normalised by $N - 1$ and the skewness value in the strategy evaluation is the skewness of the estimated total returns of the strategy.

$$\tilde{\mu}_3 = \frac{\sum_i^N \left(X_i - \bar{X}\right)^3}{(N-1)\sigma^3} \tag{2.10}$$

$\tilde{\mu}_3 =$ skewness
$N =$ number of variables in the distribution
$X_i =$ random variable

$\bar{X} =$ mean of the distribution
$\sigma =$ standard deviation

2.6.4 Kurtosis

Kurtosis is a measurement that describes how heavily tailed the probability distribution is. A high kurtosis value in the strategy indicates that the outcomes can vary drastically, meaning that a single trade can result in a significant loss or gain [5], thus it is favourable for the strategy to have a lower kurtosis value. In the evaluation, the kurtosis value is calculated on the estimated total returns and a normally distributed return will have a kurtosis value of 3, i.e. the strategy does not apply Fisher's definition of kurtosis. Additionally the kurtosis is the unbiased kurtosis normalised by $N-1$.

$$\operatorname{Kurt}[X] = \operatorname{E}\left[\left(\frac{X-\mu}{\sigma}\right)^4\right] = \frac{\operatorname{E}\left[(X-\mu)^4\right]}{\left(\operatorname{E}\left[(X-\mu)^2\right]\right)^2} = \frac{\mu_4}{\sigma^4} \tag{2.11}$$

$\mu_4 =$ fourth central moment
$\sigma^4 =$ standard deviation to the power of 4

2.6.5 Average holding

The average holding is the average time that a stock is held before it is sold.

2.6.6 Position bias

Position bias in the trading strategy is the tendency for the strategy to hold either a long or a short position. A value of 1 indicates that the model only holds long positions, whereas a value of -1 indicates that the algorithm only shorts the market. Hence, a balanced trading strategy should in theory have a value of 0, meaning that it has no bias/preferred position. In general, position bias is more useful as an indicator for hedge funds which seeks to minimise market exposure. There are a few drawbacks with the strategy, namely that a positively trending market will most likely increase the time that the strategy is in a long position which will be reflected in this indicator.

2.6.7 Slippage

Slippage is the difference between the expected price of a trade and the price at which the trade is executed [1]. In this paper, the slippage also include transaction costs. The slippage is fixed at 2% of the trade size.

Chapter 3

Method and implementation

3.1 Strategy

The generation of the model can be summarised by the following steps:

1. Based on a window size, extract the closing prices of the latest time steps.
2. Create a matrix with rows and columns for each possible state. (In the 2-state model, a 2×2 matrix is generated, where 1 row and 1 column of the matrix corresponded to a particular state.)
3. Using the extracted data, fill the matrix with the number of occurrences of each direction pair.
4. Use Laplace smoothing to regularise the data and eliminate zero-probabilities.
5. Calculate the steady-state probability vector of the smoothed matrix.
6. The option with the highest probability in the steady-state probability vector is chosen as the predicted direction of the next time step. Should there be several occurrences of the highest probability then the model will recommend the same direction as the previous direction one time step earlier. This conservative approach serves to minimise transaction costs related to changing trading position.

The optimal window size as well as the other hyperparameters mentioned in *3.3 Hyperparameters* are first optimised on a training data set before it is validated against the test set.

It is important to note that the model will reset daily, meaning that the latest number of time steps at the start of the day is not last time steps from prior trading day. This implies that the model will not trade on the first couple of time steps at the start of the day. It will start trading when enough time steps (depending on window size) have passed. The idea with this implementation is to reduce the fundamental factors which occur outside the trading hours and additionally to ensure that prior trading day's behaviour is not affecting the current day. Should the behaviour be overlapping over the days, then it can be argued that this model will detect it once enough data is retrieved for the current day. By the reasoning above, the algorithm will have to close any open positions at the end of the day.

3.2 Data collection

The price data of the NASDAQ Composite Index has been extracted from Refinitiv Eikon. Eikon is a platform where different types of financial data can be retrieved, amongst which are historical market prices. It is widely used in the financial sector as well as by academics and thus its data is considered reliable. The motivation for choosing the NASDAQ Composite Index is that a general index would reduce the exposure from each individual company and additionally reduce any market movements that could be caused by fundamental changes.

3.3 Hyperparameters

The following subsection contains a description of the hyperparameters that the model will attempt to optimise during its training. The optimal hyperparameters will then be fixed and used on the test set.

3.3.1 Number of states

The model has two different settings of states that are used; 2-states ("increase", "decrease") and 4-states ("high increase", "increase", "decrease" and "high decrease"). In the 2-state model an "increase/decrease" implies that the next stock price will either increase or decrease in relation to the current price. For the 4-state model the delimiter between "high increase" and "increase" is set as 0.03% meaning that if the forecasted price is higher than (not equal) 0.03% of the current price, the increase is considered to be "high". Similarly, for "high decrease" a change below (not equal) -0.03% is considered "high".

3.3.2 Sliding window

A sliding window will be used in the model. The window uses the latest time steps to predict the next outcome and moves forward when new time steps arises. The optimal size of the window is calculated and presented in Chapter 4.

The model begins with first tuning the window size based on a training data. In Chapter 4, the training data is always a Monday. The idea being that each Monday serves as a new clean slate and also allows for the model to always be updated. Since the model uses Monday as a training set, no actual trading is made during that day which effectively reduces each trading week to 4 days instead of 5. By updating the hyperparameter weekly, the intention is to reduce the risk of overfitting whilst assuring the model stays updated and relevant.

To tune the hyperparameter, the model is run with different window sizes, ranging between 2 and 30. As the window sizes increases the model becomes less sensitive to new information making it less useful at predicting future values. A larger window encompasses a greater number of data points and enables the identification of longer trading patterns that are less likely to reoccur due to overfitting. Additionally, a larger window size will introduce data sparsity (more sparse transition matrix) which results in a worse analysis. Moreover, a larger window places greater significance on historical data points when predicting prices. Conversely, a smaller window generates shorter sequence of patterns, which may repeat multiple times within the model. A shorter window prioritises more recent data, making it more responsive to rapid changes in the market. The window size which returns the highest model accuracy is chosen and set as the hyperparameter for the rest of the week.

3.4 Trading algorithm

To complement the trading model, a trading algorithm is implemented which trades based on the predictions made by the model. Should the model predict that the next price will increase from the current price, the algorithm will enter a long position. The opposite applies for when the model forecast a decrease in price. To maintain a simple algorithm; as soon as the predicted direction changes, the algorithm will exit its current position and enter a new position in the predicted direction. Thus the algorithm is binary; it holds either a long or a short position throughout the trading day.

The only difference in the trading algorithm between 2-state and the 4-state model is that for the 2-state model the algorithm will trade with 1 position size at all times, meaning that when the algorithm enters any trading position, that position size is of size 1. For the 4-state model, a trade with a position size of 2 occurs if the model forecasts a "high increase/decrease" and a position size of 1 if the model predicts an "increase/decrease". The idea is that a higher predicted increase/decrease indicates a stronger conviction which can be maximised by an increased position size.

It is important to note that the trading algorithm does not accumulate its position sizes for each successive time step, which implies that for the 2-state model the algorithm can hold a maximum of 1 position size in either direction whilst the 4-state model has a max limit of 2 in position size.

The trading strategy was compared to a buy-and-hold strategy which serves as a benchmark. The buy-and-hold strategy is a simple strategy which buys a position and holds it forever. It is important to note that this strategy is very naive in that it will hold its position even when the market is trending down. However, it will be very efficient at an upwards trending market.

3.5 Model evaluation

The model was evaluated with regards to accuracy, precision, recall and F1-score (see their definitions below).

$$\text{Accuracy} = \frac{TP+TN}{TP+FP+TN+FN}, \text{Precision} = \frac{TP}{TP+FP}, \text{Recall} = \frac{TP}{TP+FN}$$

where TP = True positives, TN = True Negatives, FP = False Positives, FN = False Negatives.

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The model's accuracy was analysed based on how well it managed to predict the validation data in general. An accuracy above 0.5 was deemed successful since this showed that the model is able to predict better than chance. Precision is the fraction of relevant instances among the retrieved instances and can be seen as the quality of the prediction, recall is the fraction of relevant instances that were retrieved and measures the completeness of positive predictions[10]. The F1-score is the mean of precision and recall, and serves as a useful metric when the value of either precision or recall differs significantly from the other.

Chapter 4

Results

The results below was based on an arbitrarily chosen trading week, which in this case was the week commencing 10/04/2023. For the purpose of this report, two days that manifest divergent market trends were selected for in-depth analysis. These days are Tuesday (11/04/2023) and Thursday (13/04/2023). This selection aimed to assess the performance of the strategy under varying market conditions. The details and findings pertaining to the remaining days was also summarised. For **Figures 4.0.1** and **4.0.2**, the underlying data was the training set (Monday, 10/04/2023). The underlying data for the rest of the week was the validation data set.

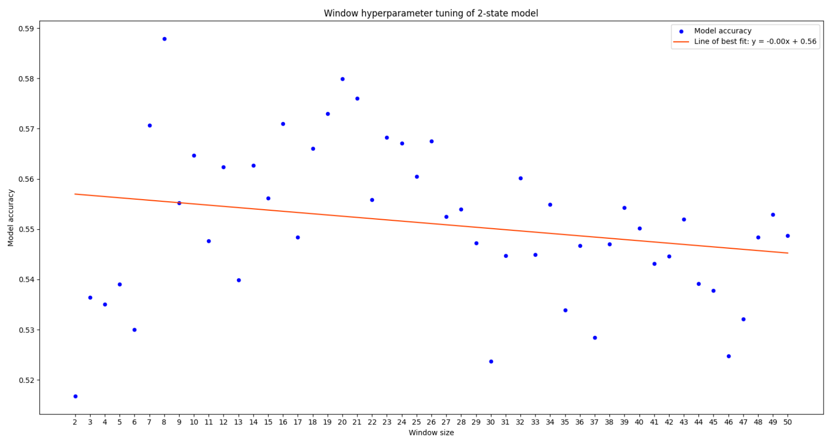

Figure 4.0.1: Model accuracy of varying window sizes for the 2-state model on the training set

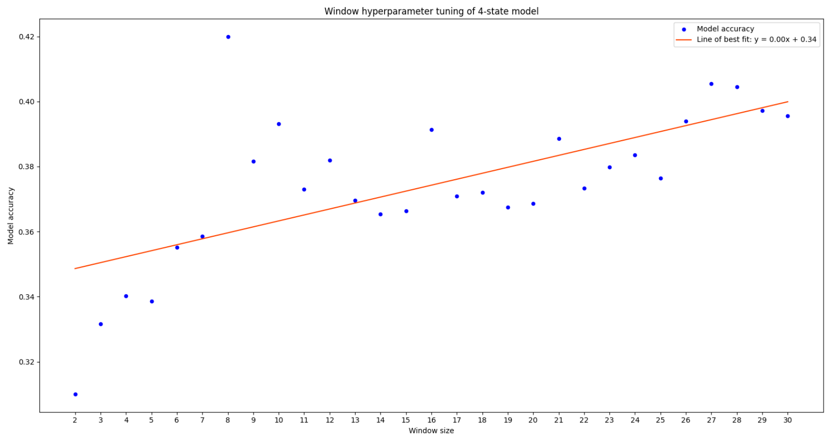

Figure 4.0.2: Model accuracy of varying window sizes for the 4-state model on the training set

2-state model				
	Precision	Recall	F1-score	Accuracy
Tuesday				
Decrease	0.45	0.42	0.43	
Increase	0.55	0.58	0.56	
Accuracy				0.51
Thursday				
Decrease	0.45	0.34	0.39	
Increase	0.60	0.71	0.65	
Accuracy				0.56

Table 4.0.1: 2-state model evaluation

4-state model				
	Precision	Recall	F1-score	Accuracy
Tuesday				
High decrease	0.24	0.17	0.20	
Decrease	0.27	0.25	0.26	
Increase	0.44	0.57	0.50	
High Increase	0.06	0.02	0.03	
Accuracy				0.36
Thursday				
High decrease	0.14	0.08	0.10	
Decrease	0.37	0.30	0.33	
Increase	0.49	0.63	0.55	
High Increase	0.12	0.09	0.10	
Accuracy				0.41

Table 4.0.2: 4-state model evaluation

4.1 Sideways trading day: Tuesday

In **Figures 4.1.1** and **4.1.3**, the green plus signs indicates a buy signal and the red plus signs indicates a sell signal, both of which were the results of the trading model predictions. The starting point of the lines (reading the figures from left to right) indicates that the algorithm enters a new position and the end point of the same line is where the position is closed. If the line is green, it indicates that the entry position was a long position and a red line signals a short position.

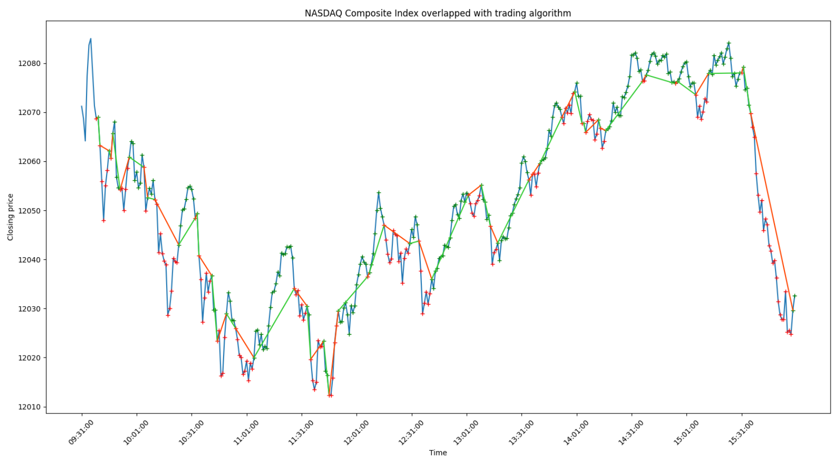

Figure 4.1.1: NASDAQ Composite Index overlapped with 2-state trading algorithm (Tuesday)

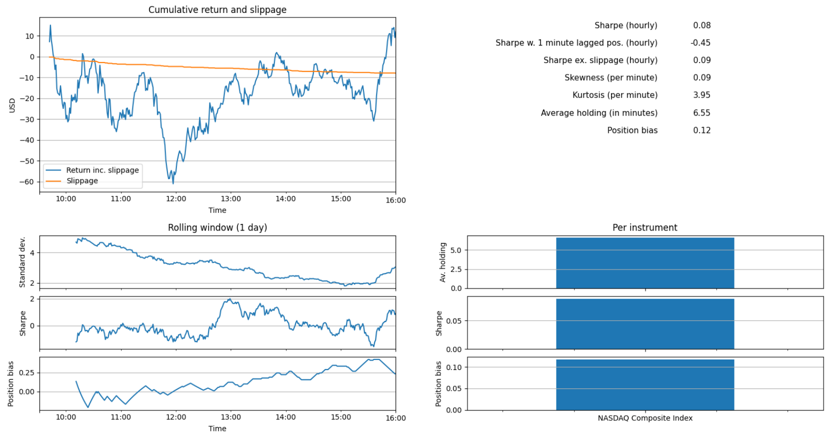

Figure 4.1.2: Evaluation of 2-state trading strategy (Tuesday)

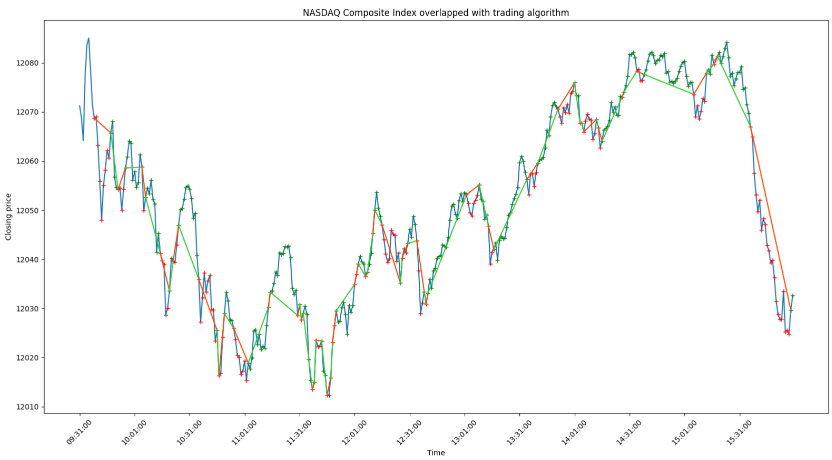

Figure 4.1.3: NASDAQ Composite Index overlapped with 4-state trading algorithm (Tuesday)

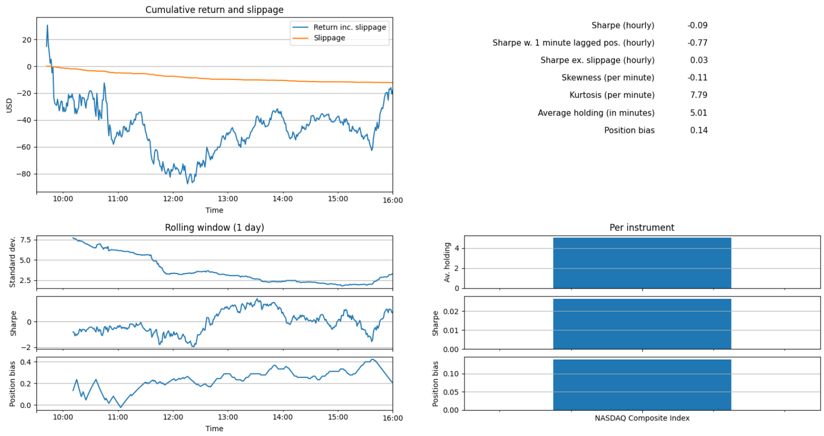

Figure 4.1.4: Evaluation of 4-state trading strategy (Tuesday)

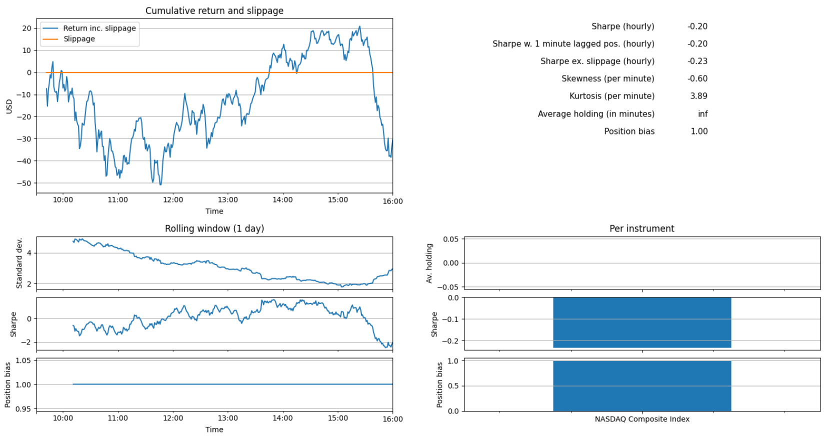

Figure 4.1.5: Evaluation of benchmark (Tuesday)

4.2 Trendy trading day: Thursday

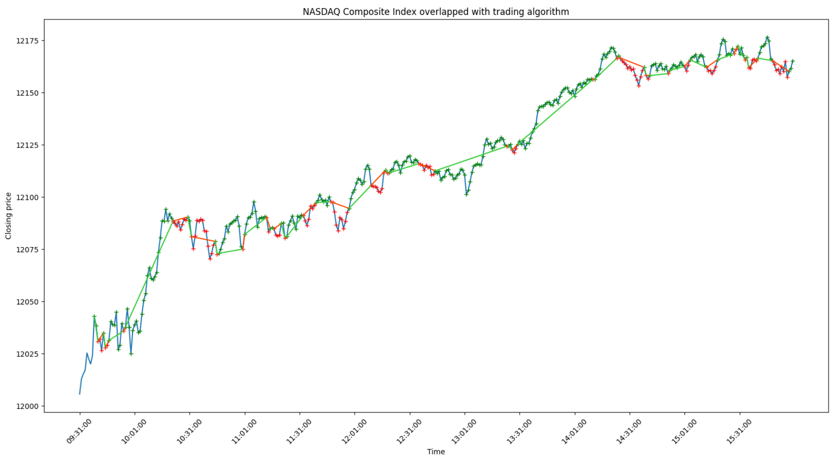

Figure 4.2.1: NASDAQ Composite Index overlapped with 2-state trading algorithm (Thursday)

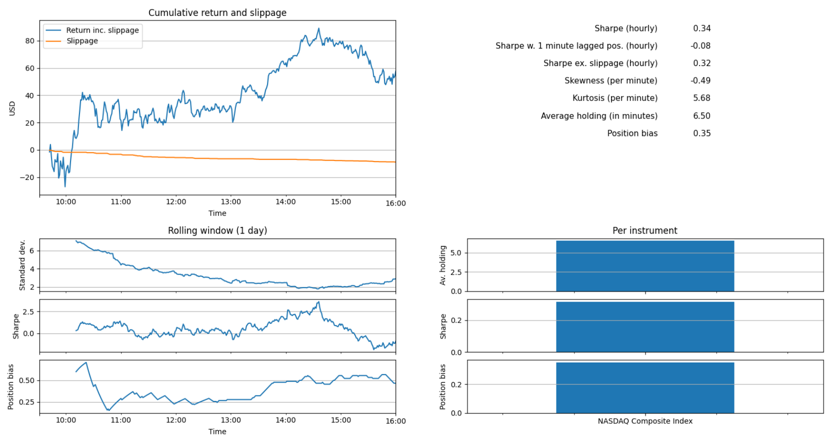

Figure 4.2.2: Evaluation of 2-state trading strategy (Thursday)

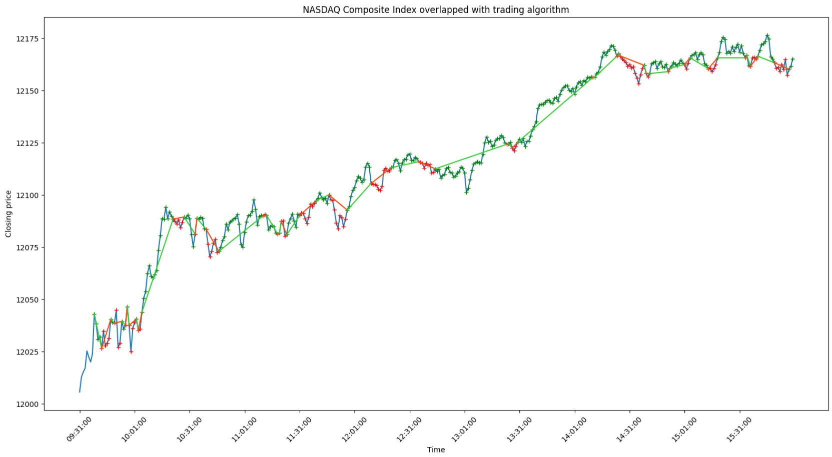

Figure 4.2.3: NASDAQ Composite Index overlapped with 4-state trading algorithm (Thursday)

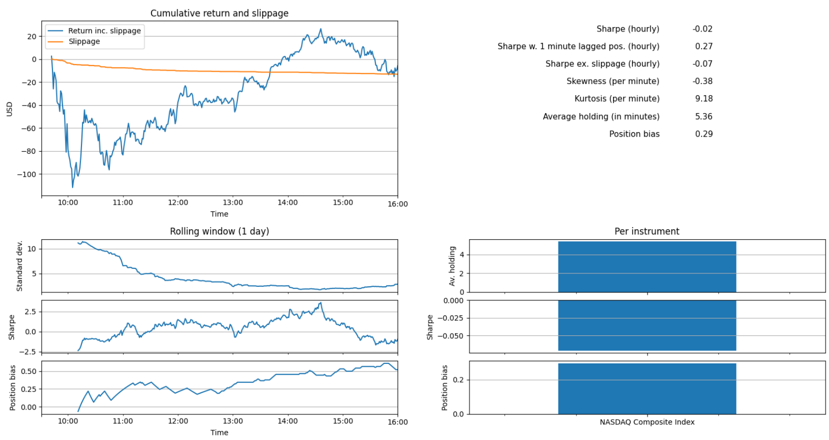

Figure 4.2.4: Evaluation of 4-state trading strategy (Thursday)

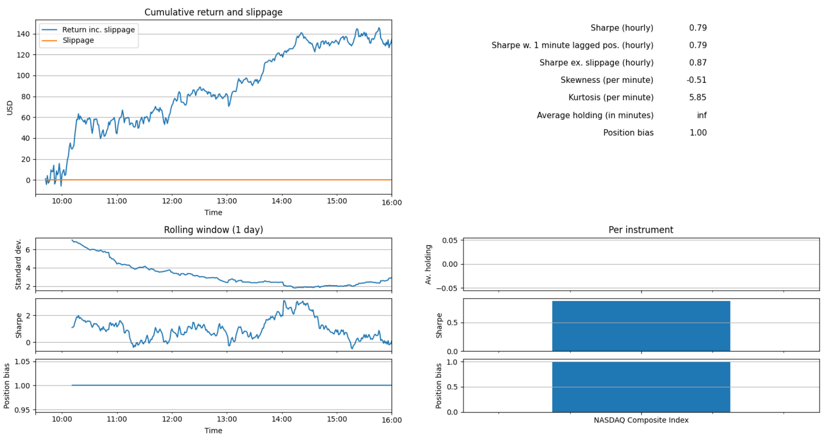

Figure 4.2.5: Evaluation of benchmark (Thursday)

4.3 The rest of the week: Wednesday and Friday

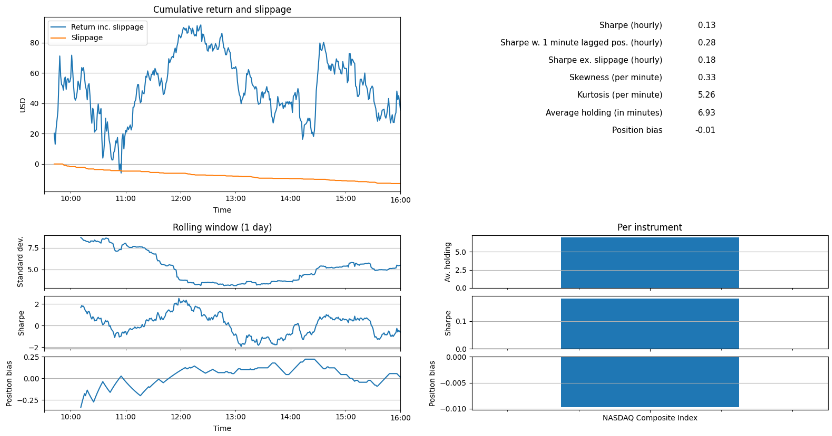

Figure 4.3.1: Evaluation of 2-state trading strategy (Wednesday)

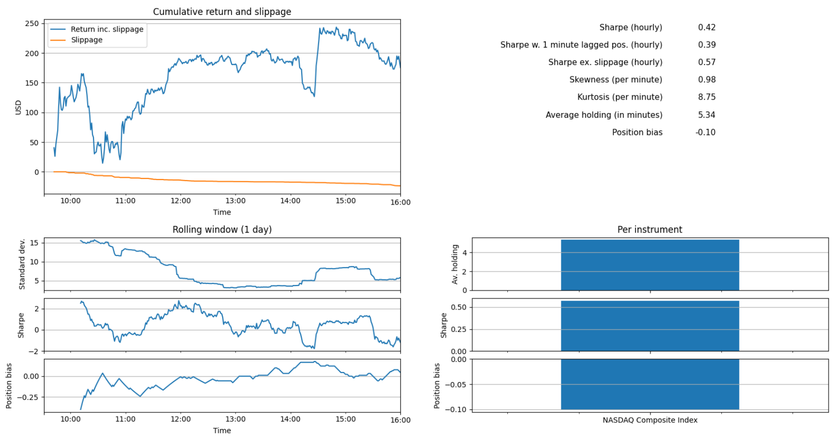

Figure 4.3.2: Evaluation of 4-state trading strategy (Wednesday)

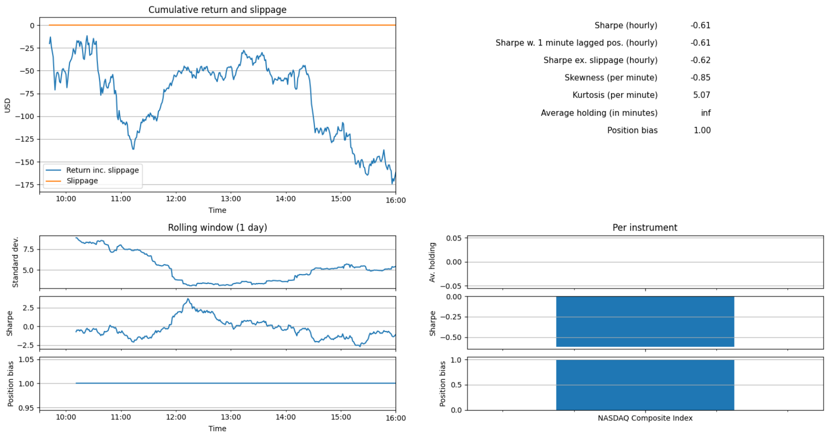

Figure 4.3.3: Evaluation of benchmark (Wednesday)

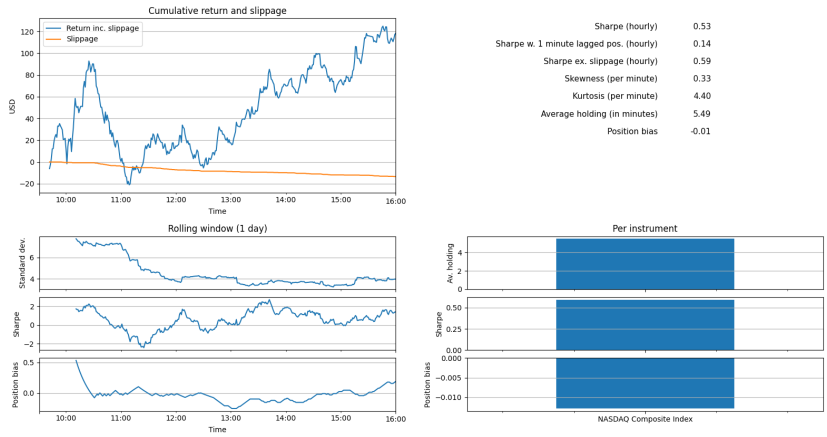

Figure 4.3.4: Evaluation of 2-state trading strategy (Friday)

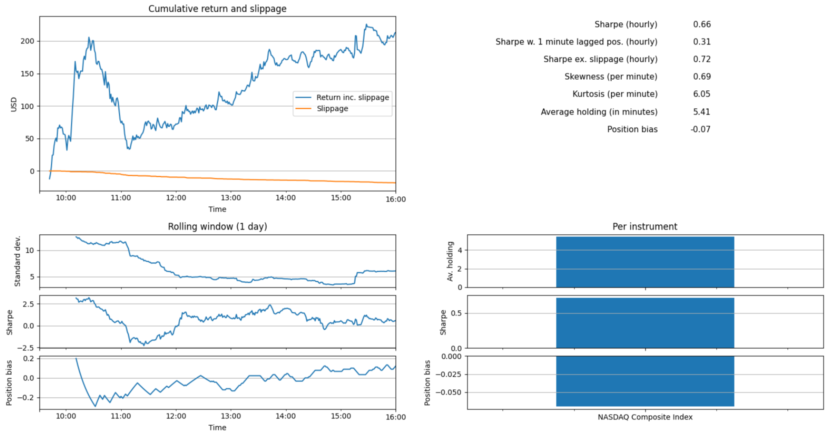

Figure 4.3.5: Evaluation of 4-state trading strategy (Friday)

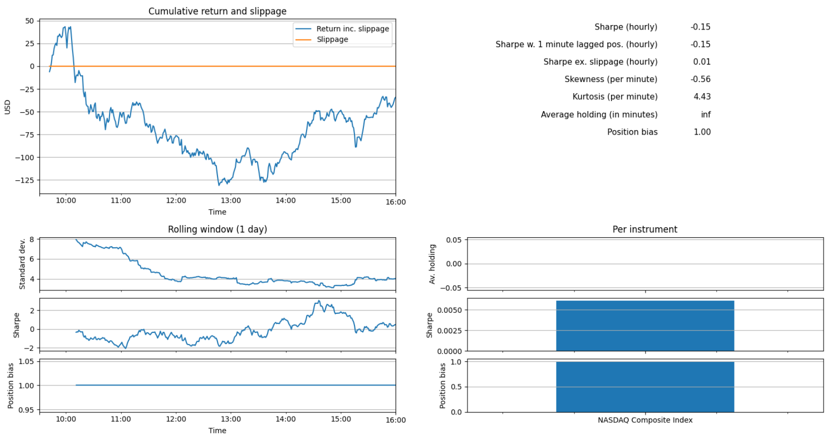

Figure 4.3.6: Evaluation of benchmark (Friday)

Chapter 5

Discussion

5.1 Evaluation using financial indicators

5.1.1 Sharpe

On Tuesday the Sharpe ratio noticeably differed between the different trading strategies (see **Figures 4.1.2** and **4.1.4**). The 2-state strategy had a slightly positive Sharpe of 0.08 while the 4-state had a negative Sharpe of -0.09. Similarily, on Thursday the 2-state strategy had an even higher Sharpe value of 0.34 whilst the 4-state strategy had a Sharpe value of -0.02 (see **Figures 4.2.2** and **4.2.4**). This indicates that the strategy produced less volatile results which can be seen when comparing the Nasdaq Composite **Figures 4.1.1** and **4.1.3** with **Figures 4.2.1** and **4.2.3**. The results suggests that 2-state strategy produces better results than the 4-state strategy during both sideways and trendy trading days. However, for the rest of the week (see **Figures 4.3.1**, **4.3.2**, **4.3.4** and **4.3.5**) the 4-state strategy's Sharpe surpasses the 2-state strategy's Sharpe, hence based on theses results, none of the strategy are consistently outperforming the other. It should be noted, however, that both strategy have a positive Sharpe during the rest of the week which indicate that both strategies are profitable during those days. Additionally, the strategies outperformed the benchmark (see **Figures 4.1.5**, **4.3.3**, **4.2.5** and **4.3.6**) for everyday apart from Thursday when the market was clearly trending upwards.

5.1.2 Skewness

The 2-state strategy had a slightly positively skew of 0.09 on Tuesday (see **Figure 4.1.2**) while the 4-state was negatively skewed with a value of -0.11

(see **Figure 4.1.4**), however on Thursday the 2-state strategy (see **Figure 4.2.2**) was increasingly negatively skewed with a value of -0.49, while the 4-state strategy was less negatively skewed in comparison with a value of -0.38 (see **Figure 4.1.4**). When comparing the strategies' skewness against the benchmark over the entire week; both the 2-state and 4-state strategies appear to be consistently higher than the benchmark. Furthermore, it is only on Tuesday that the 2-state strategy has a higher skewness than the 4-state strategy, on the other trading days, the 4-state strategy has the highest skewness of the three. However, throughout the week the skewness value ranges between 0.33 to -0.49 for the 2-state strategy, 0.98 to -0.38 for the 4-state strategy and -0.51 to -0.85 for the benchmark. Hence, based on skewness; the 4-state strategy is the most preferable. In general the strategies and the benchmark is more negatively skewed which suggest that there is a higher probability of experiencing larger losses which could eliminate the more frequent and smaller gains. On Thursday, when the market traded sideways, the 2-state strategy had the most negative skewness of the two strategies, this observation suggest that the 2-state strategy might be a riskier strategy than the 4-state strategy when the market lacks a clear trend or direction. During Tuesday, when the market was more trendy, the 2-state strategy had a higher skewness value than the 4-state strategy which could suggest that the 2-state strategy is preferable during trendy market conditions.

5.1.3 Kurtosis

The kurtosis value on both strategies and the benchmark ranges between 3.95 to 9.18, where 3 indicates normally distributed returns. Hence it suggests that both gains and losses are relatively manageable. Comparing the strategies against the benchmark; the 2-state strategy and the benchmark appear to have a similar kurtosis value in general whilst the 4-state strategy has a consistently and noticeably higher value throughout the trading week. This could mainly be attributed to the fact that the 4-state algorithm is designed to take larger positions once it foresees a larger movement, hence it is more risk tolerant. For the risk adverse trader, the 2-state strategy would then be more favourable. When comparing the kurtosis value between the sideway (Tuesday) and trendy (Thursday) trading days, the kurtosis value is consistently higher for all strategies/benchmark on the trendy day.

5.1.4 Average holding

The average holding duration consistently exhibits lower values throughout the week for the 4-state strategy when compared to the 2-state strategy. This could be attributed to the fact that the 4-state strategy switches its positions more frequently than the 2-state strategy and that a change from “increase” to “high increase” (or vice versa) corresponds to a change in the positions held. Hence adding additional states increases the likelihood of a lower average holding. Since the benchmark is a buy-and-hold strategy, it will not close its position throughout the week.

5.1.5 Position bias

On the sideways trading day (Tuesday) both the 2-state and the 4-state strategies (see **Figures 4.1.2**, **4.1.4**) have a position bias close to 0. This indicates that the strategies are taking both long and short positions throughout the day, hence has no clear bias toward any type of position. However, both values are slightly positive thus showing a slight tendency towards long positions. This behavior is also prevalent on the trendy trading day (Thursday) (see **Figures 4.2.2**, **4.2.4**) but the bias is even more positive than for Tuesday, this can be explained by Thursday being a more trendy day with more increases than decreases in the market, hence the strategy naturally favouring long positions. The position bias of the strategies are relatively close to one another throughout the week, which is partially explained by the strategies being relatively similar. An increase in the 2-state strategy corresponds to either a high increase or a an increase in the 4-state strategy. The benchmark is a buy-and-hold strategy and thus it will not change its position, hence its position bias having a value of 1 throughout the week.

5.2 Accuracy

The accuracy of the 2-state model is higher than the accuracy of the 4-state model (see **Tables 4.0.1** and **4.0.2**). It should be noted that the benchmark accuracy of the 4-state model should not be 0.5 nor 0.25. This is due to the cutoff value for high increase/decrease being 0.03% which does not separate the results into even buckets. Had the cutoff value separated the results into four even buckets, the benchmark accuracy would be 0.25 instead of 0.5. With the current cutoff of 0.03% the high increase/decrease are deliberately considered to be a rare occurring event hence the bucket sizes is by design uneven, which in turn results in that the benchmark of 0.25 or 0.5 being invalid.

The window is optimised every Monday of the week. **Tables 4.0.1** and **4.0.2** show the model precision, recall, F1-score and accuracy of Tuesday (Sideways trading day) and Thursday (Trendy trading day) for the week. The accuracy of the models are slightly higher on the trendy trading day as opposed to the sideways trading day, this is due to the 2-state and 4-state strategies being a trend following strategy.

5.3 Window

In the case of the 2-state model, there is an observable trend of a slight decline in accuracy as the window size increases (see **Figure 4.0.1**). Conversely, the 4-state model exhibits a slightly increasing trend (see **Figure 4.0.2**), indicating that a larger window size is more advantageous. However, an exceedingly large window recommended by the model renders it overly rigid, resulting in a single prediction of either increase or decrease. Given the assumption that the stock market behaves stochastically in one-minute intervals, with a 50% likelihood of a stock trading up or down, a model that only suggests one direction would yield an average accuracy of 0.5.

This behavior arises when the average accuracy falls below 0.5, prompting the model to suggest larger window sizes until it reaches a plateau at 0.5 accuracy. Furthermore, the 4-state model's accuracy does not adhere to the benchmark of 0.5 due to the uneven distribution of data across different buckets, as previously mentioned. To mitigate this issue and prevent data sparsity and infinitely large windows, the optimal window size has been constrained to a maximum of 30 time steps.

In the results, the highest model accuracy was a window size of 8 for both models (see **Figures 4.0.1** and **4.0.2**) hence this window size was chosen. However, one may argue that this particular window size is a potential outlier for both models as it appears to deviate from the rest of the window sizes. Therefore, it may be more reasonable to choose a window size which is more in line with the cluster of other window sizes. An additional suggestions for future research is tuning the window sizes based on other metrics than accuracy; a window size tuned after returns or a combination of metrics could result in an even more profitable strategy.

5.4 Returns

The returns are shown in the results under the frame “Cumulative return and slippage”. To simulate a real trading run, a trading fee of 2% has been added to all buy and sell orders. The trading fee is included in the slippage and increase as the strategies executes trades on the market.

On Tuesday (sideways trading) the two strategies performed better than the benchmark. This indicates that the strategies can mitigate risk and be profitable even when the market is not trending. The opposite occurred during Thursday (trendy day), where the benchmark outperformed the 2-state and 4-state strategies. Apart from Thursday, the strategies are consistently outperforming the benchmark's final accumulated return. A rough estimate of the total accumulated return at the end of the week can be summarised by the figures in the results, this summary yield the following results: 2-state strategy has a final accumulated return of +215 USD, 4-state strategy a return of +360 USD and the benchmark with a return of -105 USD. Hence, both strategies are more favourable than the benchmark and of the two strategies the 4-state strategy performs better than the 2-state strategy in terms of net return at the end of the week. It should however be noted that the gains of the 2-state strategy is more consistent than the 4-state strategy, in the sense that the 2-state strategy has a positive return at the end of every trading day, whereas the 4-state strategy has a negative return for both Tuesday and Thursday which is compensated for higher returns on the rest of the week. The more volatile returns of the 4-state strategy is a consequence of the higher skewness and kurtosis of the 4-state strategy when compared to the 2-state strategy.

Chapter 6

Conclusions

Based on the results, the 2-state and 4-state strategies are profitable at the end of the week, however as seen in, for instance **Figure 4.1.4**, the returns are not consistent and will occasionally result in losses. When comparing the 2-state and 4-state strategies against the buy-and-hold strategy (benchmark), the Markov-based strategies are outperforming it in metrics such as Sharpe, skewness, position bias and returns. Between the 2-state and 4-state strategies, the 2-state strategy is more consistent in terms of having a positive daily returns. While the 4-state strategy has returns which are more volatile, by the end of the week, its returns surpass the returns of the 2-state strategy. Hence, which strategy is better than the other may depend on the individual's risk preference. The accuracy of the 2-state model appears to surpass 0.5, thus the application of a Markov model for predicting future market movements appear to be viable. By furthering optimising the trading strategy it is not unfeasible that the it can be turned into an even more profitable and applicable trading system.

6.1 Future Work

The existing model serves as a solid foundation upon which further improvements can be built. The addition of more states to the model could prove beneficial and would enable the incorporation of additional parameters, such as volume. By considering factors beyond just price, such as volume, the model's predictive accuracy may be enhanced. This expanded scope holds promise for refining the model's predictions and delivering more accurate insights for traders.

Furthermore, the optimal window size is currently tuned based on accuracy. However future work may instead attempt to tune it by other metrics such as

Sharpe ratio, kurtosis etc. or perhaps a combination of these.

Additionally, the current model is tuning the parameter on data from Monday (see **Figures 4.1.1** and **4.2.1**), which assumes that this day is representative for the rest of the week. Another approach would be to tune based on older historical data in an attempt to find an even more representative data set.

The current 2-state algorithm operates by executing buy and sell orders at the prevailing market price, with a position size of 1 while the 4-state algorithm builds upon this and executes buy and sell orders with different predetermined position sizes depending on how great the increase/decrease is. To further enhance the algorithm's capabilities, a future iteration could be developed to alter the position size based on the probability of a specific event occurring. While this enhancement would introduce greater complexity, it has the potential to generate larger capital gains or losses, which would undoubtedly be of interest to traders. Another interesting area of study is incorporating Kelly criterion to optimise the position sizes.

The incorporation of a stop loss function into the algorithm has the potential to effectively minimise risks and enhance risk management. By integrating this mechanism, the algorithm can automatically trigger predefined actions when certain thresholds or criteria are breached, thereby limiting potential losses. This proactive approach allows for timely intervention and can contribute to maintaining a more controlled and secure trading environment. Consequently, the implementation of a stop loss function serves as a valuable tool for risk mitigation within the algorithm's operations.

Bibliography

[1] CFI. "Slippage". In: (2023). Accessed: 2023-05-04. URL: https://corporatefinanceinstitute.com/resources/capital-markets/slippage/.

[2] Chen, J. "Skewness: Positively and Negatively Skewed Defined with Formula". In: (2023). Accessed: 2023-05-04. URL: https://www.investopedia.com/terms/s/skewness.asp.

[3] Dar, G.F., Padi, T.R., and Rekha, S. "Stock price prediction using a Markov chain model: a study for TCS share values, Advances and Applications in Statistics 80". In: (2022). Accessed: 2023-04-17, pp. 83–101. DOI: http://dx.doi.org/10.17654/0972361722068.

[4] Department, Statista Research. "Monthly turnover value of the equities market in the United States from January 2018 to August 2022, by operator". In: (2022). Accessed: 2023-04-14. URL: https://www.statista.com/statistics/1277225/equities-market-turnover-operator-usa/.

[5] Kenton, W. "Kurtosis Definition, Types, and Importance". In: (2023). Accessed: 2023-05-04. URL: https://www.investopedia.com/terms/k/kurtosis.asp.

[6] Liberto, D. "https://www.investopedia.com/terms/l/loss-psychology.asp". In: (2022). Accessed: 2023-05-04. URL: https://www.investopedia.com/terms/l/loss-psychology.asp.

[7] Sharpe, William F. "The Sharpe Ratio". In: *The Journal of Portfolio Management* (1994). Accessed: 2023-05-04. URL: https://web.stanford.edu/~wfsharpe/art/sr/SR.htm.

[8] Sundberg, J. and Klacksell, G. "Can you describe the stock index with a Markov chain?" In: (2012,). Accessed: 2023-04-17. URL: http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-105228.

[9] Taylor, S. "Positively Skewed Distribution". In: (2023). Accessed: 2023-05-04. URL: https://corporatefinanceinstitute.com/resources/data-science/positively-skewed-distribution/.

[10] Wood, Thomas. "Precision and Recall". In: (2022). Accessed: 2023-05-08. URL: https://deepai.org/machine-learning-glossary-and-terms/precision-and-recall.

Appendix A

Additional figures

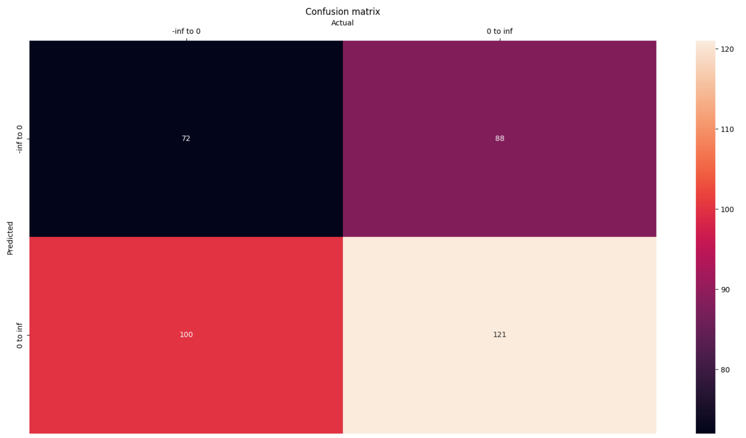

Figure A.0.1: Confusion matrix of 2-state trading model (Tuesday)

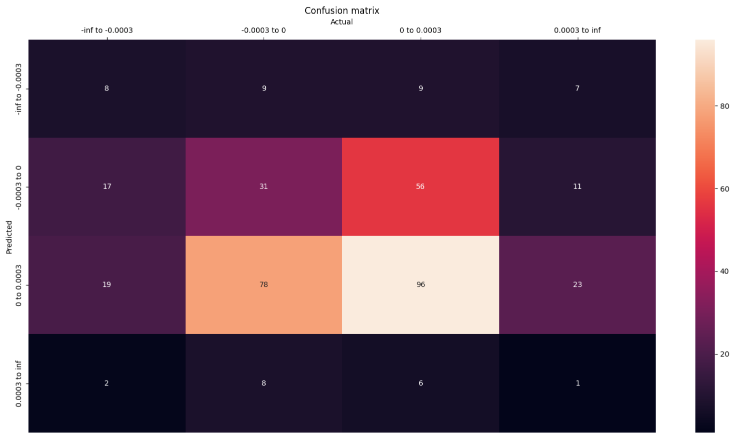

Figure A.0.2: Confusion matrix of 4-state trading model (Tuesday)

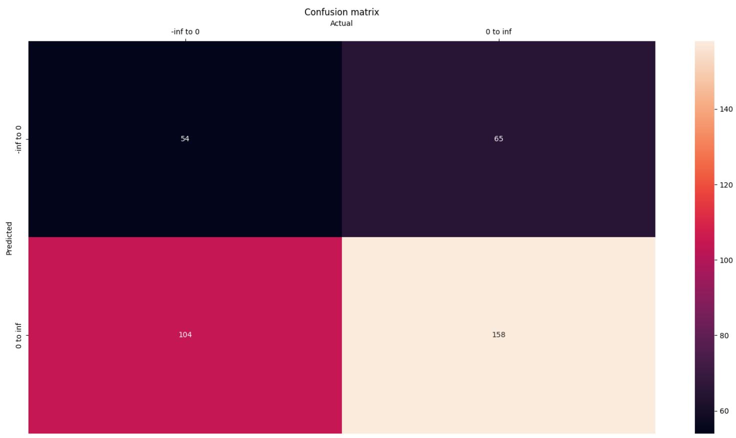

Figure A.0.3: Confusion matrix of 2-state trading model (Thursday)

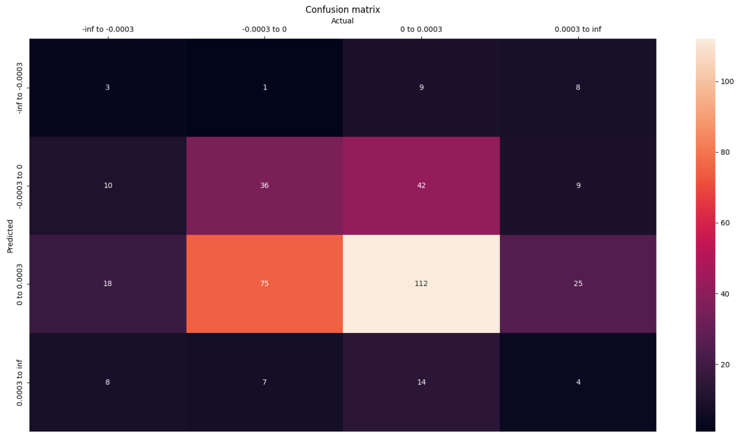

Figure A.0.4: Confusion matrix of 4-state trading model (Thursday)

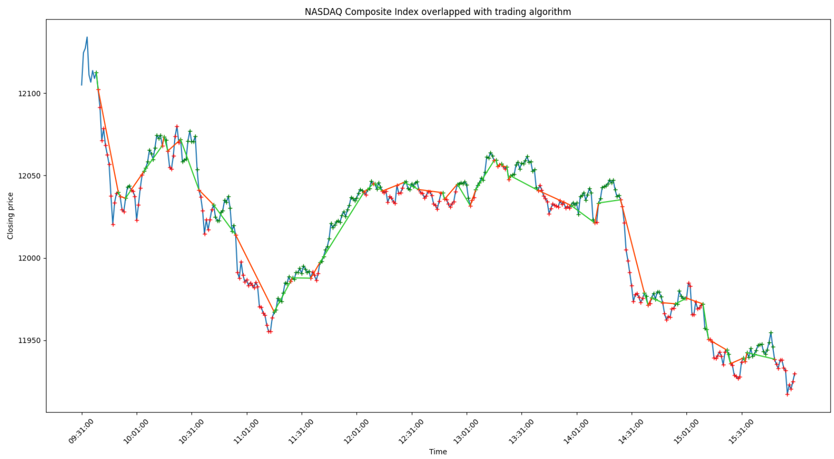

Figure A.0.5: NASDAQ Composite Index overlapped with 2-state trading algorithm (Wednesday)

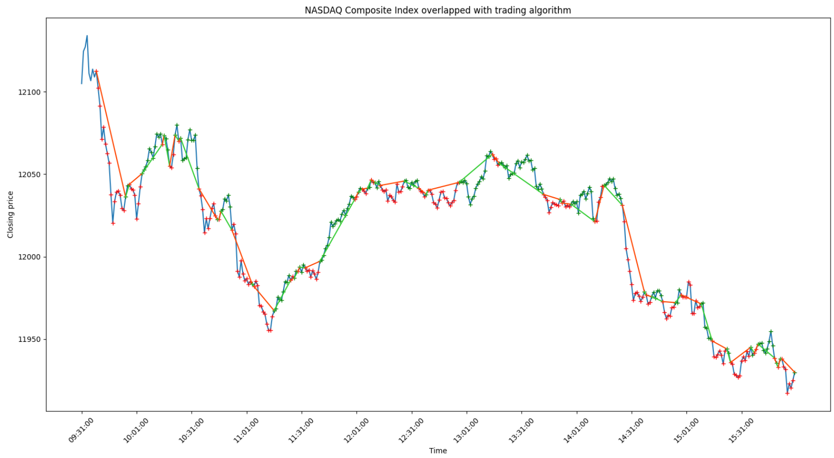

Figure A.0.6: NASDAQ Composite Index overlapped with 4-state trading algorithm (Wednesday)

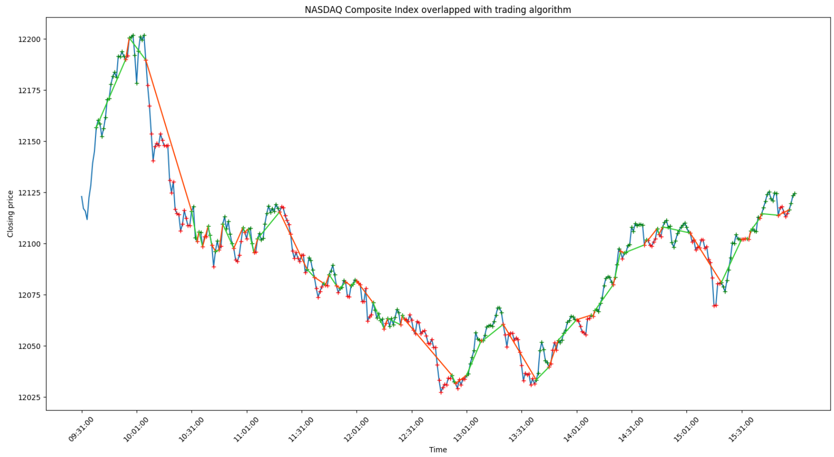

Figure A.0.7: NASDAQ Composite Index overlapped with 2-state trading algorithm (Friday)

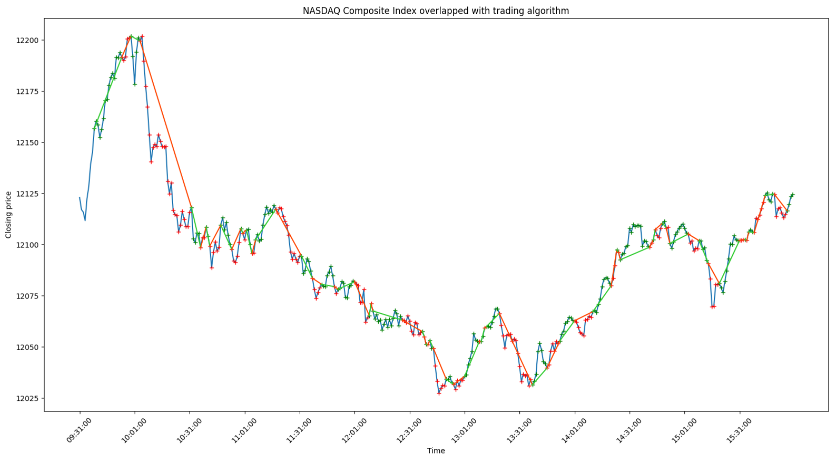

Figure A.0.8: NASDAQ Composite Index overlapped with 4-state trading algorithm (Friday)

TRITA – XXX-XXX 20XX:XX