# Applying Machine Learning Techniques to Assess Whether a Country's Currency can Predict the Movement of their Respective Stock Market Index

COMP4971C – Independent Work (Fall 2021)

By Tara Relan

Supervised by Dr. David Rossiter

**Abstract**

This project used the machine learning technique Temporal Convolutional Network to determine whether the United States Dollar / Great British Pound exchange rate could be used to predict the closing price of the Financial Times Stock Exchange 100 Index. As there was an accuracy of 89.96%, it can be concluded that the exchange rate could be used to predict the movement of stock index. Two different trading strategies were then implemented on the closing price of the exchange rate using 13,028 data points from 31/12/1985 to 06/10/2021. The two strategies were (1) a pair's trading strategy where trading signals were determined using Bollinger Bands, and (2) a buy and hold strategy. The results were compared and back tested using the Annual Average Return (pair's trading and the buy and hold strategy yielded 3.05% and 1.22% respectively), Maximum Drawdown ($-27.84\%$ from the pair's trading strategy versus $-19.67\%$ from the buy and hold strategy), and Sharpe Ratio (pair's trading and the buy and hold strategy yielded 1.92 and 0.16 respectively). The pair's trading strategy outperformed the buy and hold strategy in both the Annual Average Return and Sharpe Ratio, yet the buy and hold strategy outperformed in terms of the Maximum Drawdown.

**Index Terms**

Temporal Convolutional Network; Bollinger Bands; Buy and Hold; Pair's Trading; Statistical Arbitrage.

# Table of Contents

**Introduction** .... 4
**Method** .... 4
**Prediction Algorithm** .... 5
Augmented Dickey-Fuller and Engle-Granger Tests .... 5
Temporal Convolutional Network .... 6
Mean Absolute Percentage Error .... 8
**Strategy Implementation** .... 9
Bollinger Bands .... 9
Hurst Exponent .... 10
Half-Life Test .... 12
Buy and Hold Strategy .... 13
Pair’s Trading Signals .... 14
**Performance Evaluation** .... 16
Average Annual Return and Compound Annual Growth Rate .... 16
Maximum Drawdown .... 17
Sharpe Ratio .... 18
Summary of Results .... 19
**Conclusion** .... 19
**Further Exploration** .... 20
Kalman Filter .... 20
Clustering Methods .... 20
Basket Trading .... 20
**References** .... 21
**Appendix** .... 24

# Table of Figures

Figure 1: Graph of the Normalised USD/GBP and FTSE100 Closing Prices .......................... 6

Figure 2: A 1D Convolutional Network........................................................................ 7

Figure 3: A 1D TCN Network...................................................................................... 7

Figure 4: The Normalised Closing Prices of the USD/GBP Exchange Rate and FTSE100, Along with the Forecasted FTSE100 Closing Price.................................................... 8

Figure 5: Graph of the Spread of the USD/GBP and FTSE100 Closing Prices...................... 10

Figure 6: Graph with an Example of Bollinger Bands........................................................... 10

Figure 7: Variance of the Hurst Exponent with Lag ............................................................. 12

Figure 8: Output from the Buy and Hold Strategy Implemented on the USD/GBP Exchange Rate Closing Price.............................................................................................. 13

Figure 9:Trading Signals for Selling the FTSE100 from the Pair’s Trading Strategy Between the USD/GBP and FTSE100 Closing Price ............................................................ 14

Figure 10: Trading Signals for Buying the FTSE100 from the Pair’s Trading Strategy Between the USD/GBP and FTSE100 Closing Price ............................................................ 15

Figure 11: Trading Signals for Selling the USD/GBP from the Pair’s Trading Strategy Between the USD/GBP and FTSE100 Closing Price ............................................................ 15

Figure 12: Trading Signals for Buying the USD/GBP from the Pair’s Trading Strategy Between the USD/GBP and FTSE100 Closing Price ............................................................ 16

Figure 13: Cumulative Returns from the Pair's Trading Strategy........................................... 17

Figure 14: Underwater Plot Showing the Drawdown over Time........................................... 18

# Table of Data

Table 1: Parameters Used in Implementing the TCN Model.................................................... 8

Table 2: Variance of the Hurst Exponent with Lag ............................................................... 11

Table 3: Parameters used in Back Testing ............................................................................ 14

Table 4: Summary of Results from Back Testing.................................................................. 19

Table 5: Average Annual Return for the Pair's Trading Strategy ........................................... 24

## Introduction

It is no secret that macroeconomic indicators can be used to analyse stock market movements. More specifically, there is an interesting relationship between a country's exchange rate and their respective stock market index prices. London is one of the world's oldest financial centres and is regarded as one of the world's largest international financial centres. Furthermore, the Pound Sterling is the world's fourth most traded currency. During the Covid-19 global pandemic, there have been changes to the United States Dollar (USD) / Great British Pound (GBP) exchange rate, which can be compared to changes in the Financial Times Stock Exchange 100 (FTSE100) Index's closing price. The FTSE100 is a share index of the 100 companies listed in the London Stock Exchange with the highest market capitalism.

This project focused on the effect of the USD/GBP exchange rate on the FTSE100 closing price from 31/12/1985 to 06/10/2021. This range was chosen as it was the earliest set of available data found on Yahoo! Finance for both the exchange rate and FTSE100 closing price. Both trade by the minute, hence for simplicity, closing prices were used. A supervised machine learning algorithm was created to predict changes in the FTSE100 closing price due to changes in the USD/GBP closing price. Furthermore, a pair's trading algorithm was created to trade both these stocks and back testing was used to evaluate the algorithm's performance.

This project relied on the following assumption. In 2021, foreign investors now own 66% of UK-listed shares [27]. US investors were the biggest shareholders, especially US mutual funds, which now own more UK shares than are held by UK unit trusts. Because of this, the USD/GBP exchange rate has a big impact on the stock index as according to macroeconomic theories, exchange rates have a negative impact on stock index prices. If the exchange rate increases, or the GBP becomes relatively stronger to the USD, by encouragement of excess demand, there is a contraction effect on the stock market. This contraction leads to a decrease in the stock index price, ceteris paribus. The opposite should also hold true when the GBP becomes weaker relative to the USD.

## Method

Python is an incredibly popular programming language and was particularly useful in this project as there is a wide range of libraries available to analyse data. The Pandas library and Yahoo! Finance's API was used to extract the USD/GBP and FTSE100 closing prices. The data was then cleaned and split into two datasets – the training and test sets. A $70:30$ split was used to separate the data into the training and test set respectively [26].

The Statsmodel library was then used to run the Augmented Dickey-Fuller (ADF) and Engle-Granger (EG) tests to determine whether the USD/GBP exchange rate and FTSE100 closing price datasets were cointegrated. Another library, Darts, trained a supervised machine learning algorithm using a Temporal Convolutional Network (TCN). To test the accuracy of the machine learning algorithm, it was run over the test dataset. Mathematically speaking, the Mean Absolute Percentage Error (MAPE) was also calculated to assess the accuracy.

Finally, a pair's trading strategy was implemented using the NumPy, SciPy, Statsmodel and Pandas libraries and the Average Annual Return (AAR), Maximum Drawdown (MDD) and Sharpe Ratio (SR) were also calculated. To compare the pair's trading strategy to a benchmark, Cerebro was used to implement a buy and hold strategy on the USD/GBP closing price and the AAR, MDD and SR were also computed and compared.

**Prediction Algorithm**

Augmented Dickey-Fuller and Engle-Granger Tests

Economic theory suggests that many time-series datasets move together, fluctuating around a long-term equilibrium. In statistics, this long-term equilibrium is tested and measured using cointegration [4]. Cointegration occurs when two non-stationary time-series have a long-term equilibrium, move together in such a way that their linear combination results in a stationary time-series, and share an underlying common stochastic trend. A key consideration prior to testing for cointegration is whether there is theoretical support for a cointegrating relationship [4]. Economic theory implies the strength of the currency is a major factor that affects stock index prices. As such, the USD/GBP exchange rate and FTSE100 are implied to be cointegrated.

The Engle-Granger cointegration test considers the case where there is a single cointegrating vector. The EG test follows the intuition that if the variables are cointegrated, the residual cointegration regression should be stationary [4]. As the cointegration vector is unknown, Ordinary Least Squares (OLS), a linear least-squares method for estimating unknown parameters in a linear regression model, was used to estimate the cointegrating vector from the regression.

$$y_{1,t} = c + \beta y_{2,t} + u_t$$

Where:

| | |
|---|---|
| $Y_t = (y_{1,t}, y_{2,t})$ | The first and secondary time-series respectively. |
| $\beta = \begin{bmatrix} 1 \\ -\beta_2 \end{bmatrix}$ | The cointegrating vector such that $\beta Y_t = y_{1t} - \beta_2 y_{2t}$. |
| $u_t$ | Stationary cointegrating error component, or a short-term deviation from the long-term equilibrium. |

A standard unit root testing should then be used to test the residuals. A unit root is a stochastic trend in a time-series, or a “random walk with drift” [23]. If a time-series has a unit root, the systematic pattern is unpredictable. Unit root tests are also used to test for stationarity in a time-series. A time-series is stationary if a shift in time does not change the shape of the distribution. To test for residuals, the ADF test was used.

In the EG test, the null hypothesis states there is no correlation. The decision rule to reject the null hypothesis is if the ADF test value is less than the critical value (0.05). The results from the ADF test gives a value of 0.0492, which is less than the critical value. Hence, the null hypothesis was rejected, and it was concluded the two datasets were cointegrated.

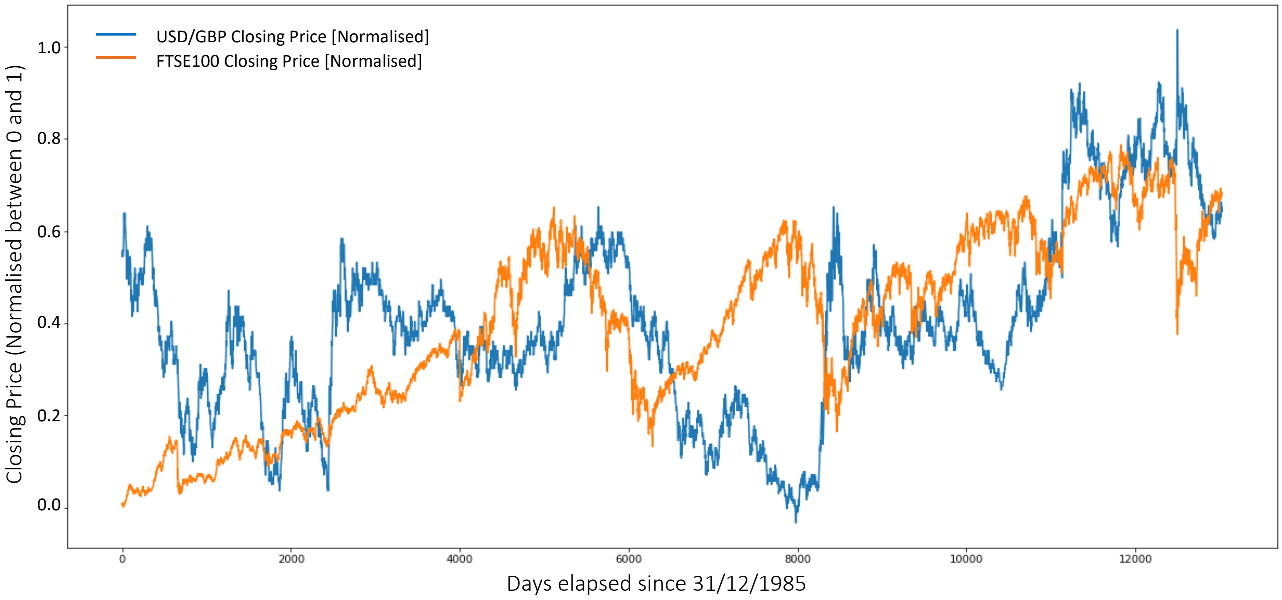


*Figure 1: Graph of the Normalised USD/GBP and FTSE100 Closing Prices*

In $Fig. 1$, since the USD/GBP and FTSE100 closing prices are of different magnitudes of 10, the two datasets were normalised between 0 and 1.

<u>Temporal Convolutional Network</u>

Although commonly associated with image classification tasks, Convolutional Neural Networks (CNNs) have proven to be useful in forecasting time series. Researchers S. Bai et al. have suggested that CNNs should be taken into consideration as a primary candidate when modelling sequential data [7]. They were able to show that CNNs can achieve better performance than Recurrent Neural Networks (RNNs), in many tasks while avoiding common drawbacks of RNNs, such as the exploding or vanishing gradient problem or the lack of memory retention [7].

A TCN consists of dilated, causal 1-Dimentional convolutional layers with the same input and output length [7]. A 1D convolutional network takes a 3D input and output tensor, whose shapes are $\begin{bmatrix} batch\_size \\ input\_length \\ input\_size \end{bmatrix}$ and $\begin{bmatrix} batch\_size \\ output\_length \\ output\_size \end{bmatrix}$ respectively.

To understand how a single input layer converts its input to output, consider a 1D input and output tensor, i.e., $num\_input\_channels = 1$ and $num\_output\_channels = 1$ respectively. To compute one element of the output, the dot product of the input and kernel vectors of learned weights are taken. To calculate the rest of the output elements, the same procedure is applied, but the $kernel\_size$ window of the input sequence is shifted to the right by a single element [7].

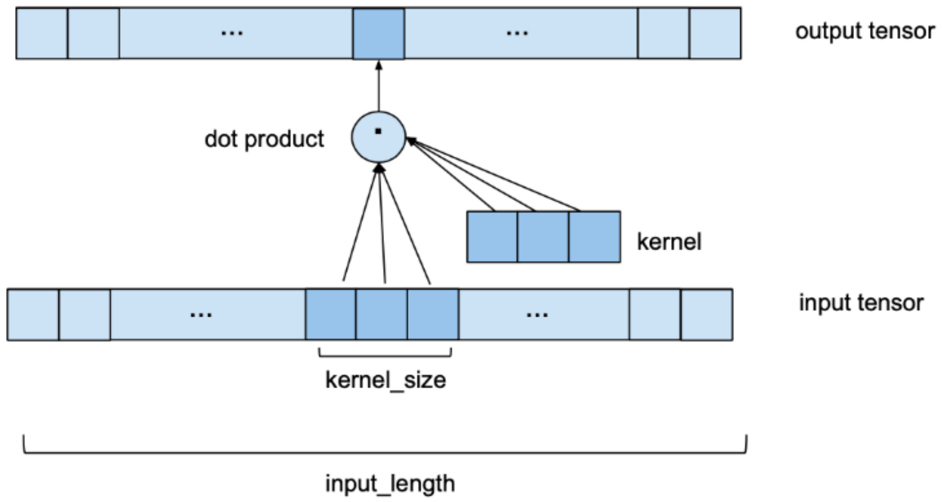


*Figure 2: A 1D Convolutional Network*

If $num\_input\_channels > 1$ and $num\_output\_channels > 1$, the process is repeated for every output channel with a different kernel matrix. The output vectors are stacked on top of each other, resulting in an output tensor of shape $\begin{bmatrix} input\_channel \\ num\_output\_channel \end{bmatrix}$ [7]. The number of kernel weights is then $kernel\_size * num\_output\_channels * num\_input\_channels$ [7].

Dilation refers to the distance between elements of the input sequence that are used to compute one entry of the output sequence [7]. A conventional convolutional layer is a 1D diluted layer since the input elements to compute an output element are adjacent.

Given the $input\_length$, $kernel\_size$, $dilation\_base$ and the minimum number of layers required for full coverage, the basic TCN network would look like $Fig. 3$.

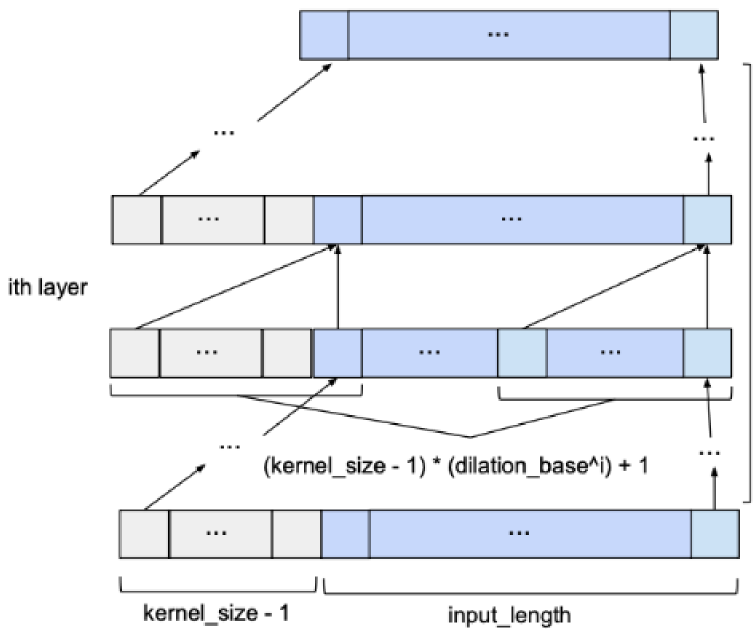


*Figure 3: A 1D TCN Network*

The parameters used in this project were as follows:

| Parameter | Value |
|---|---|
| Datapoints | 13028 |
| Input chunk length | 2889 |
| Output chunk length | 1444 |
| Epochs | 115 |

*Table 1: Parameters Used in Implementing the TCN Model*

Out of the 13,028 datapoints from 31/12/1985 to 06/10/2021, a 70:30 split was imposed to separate the data into a training and testing dataset. The training dataset had 8,695 datapoints, and testing 4,333. A 2/3 chunk of the 4,333 datapoints were taken as the input length (2,889) and the rest (1,444) as output length. To reduce the MAPE, the number of epochs was chosen to be 115.

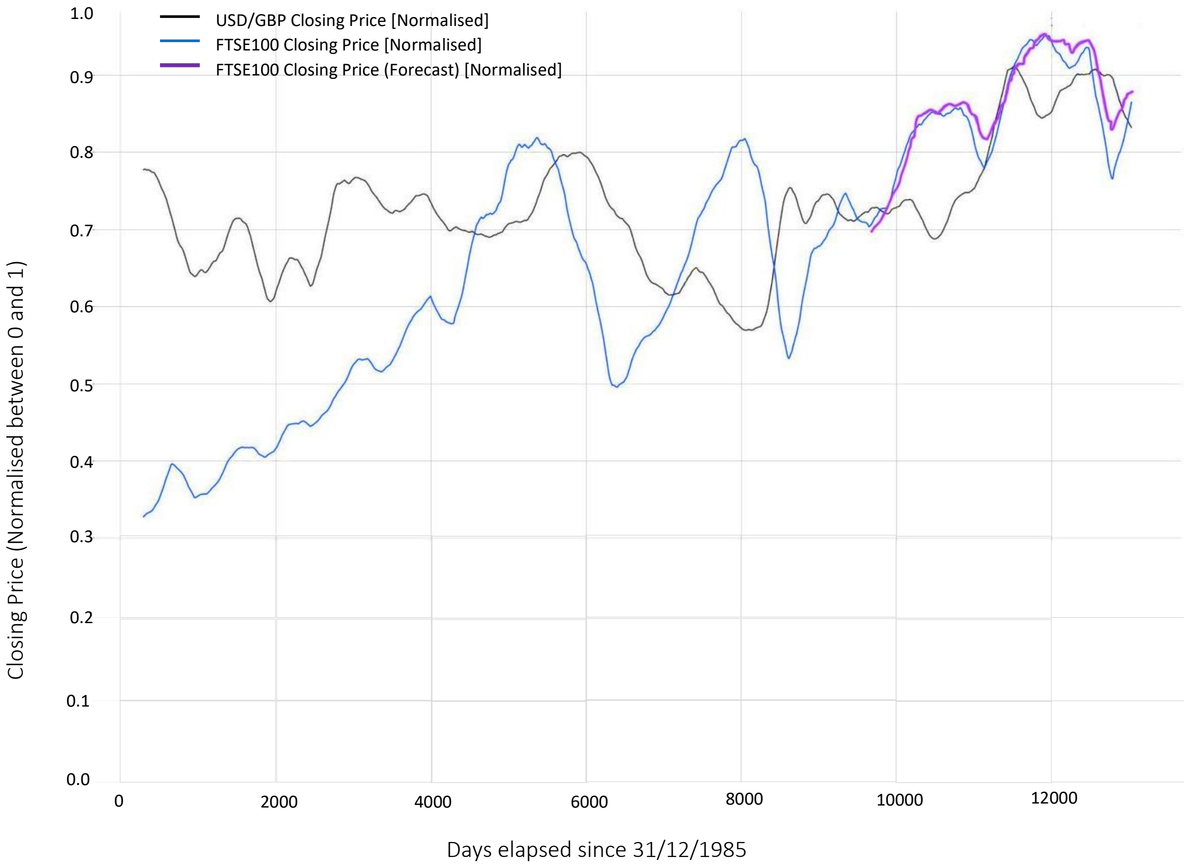


*Figure 4: The Normalised Closing Prices of the USD/GBP Exchange Rate and FTSE100, along with the Forecasted FTSE100 Closing Price*

<u>Mean Absolute Percentage Error</u>

The MAPE is a measure of prediction accuracy of a forecasting method [24], defined as:

$$MAPE = \frac{100}{n}\sum_{t=1}^{n}\left|\frac{A_t - F_t}{A_t}\right|$$

Where:

| | |
|---|---|
| $A$ | The actual value at time $t$. |
| $F_t$ | The forecasted value at time $t$. |

The value of the MAPE after running the machine learning algorithm was $10.04\%$, meaning the accuracy of the algorithm was $100 - 10.04 = 89.96\%$.

**Strategy Implementation**

Pair's trading has long been a popular statistical arbitrage strategy. The strategy involves matching a long and short position of two stocks with high cointegration [23]. The fundamental idea is to short-sell the relatively overvalued stock and buy the relatively undervalued stock based on the spread and its standard deviation [24]. To identify when these trades should be made, Bollinger Bands were used.

If the mean-reverting relationship breaks down over time, the pair's trading strategy risks resulting in heavy short-term losses [23]. Several tests were used to evaluate the behaviour of this pair, including the ADF test as mentioned before, the Hurst Exponent (H) which helps identify whether a series is trending, mean-reverting, or undergoing a random walk, and a half-life test from the Ornstein-Uhlenbeck (OU) process which was used to test for the speed of mean-reversion.

<u>Bollinger Bands</u>

A Bollinger Band is a technical analysis tool defined by a set of trendlines often plotted $\pm 2$ standard deviations away from a 20-day simple moving average of the security's price [2].

$$Upper\ Bollinger\ Band\ (UB) = \mu(z_t, n) + \sigma * \sigma[z_t, n]$$

$$Lower\ Bollinger\ Band\ (LB) = \mu(z_t, n) - \sigma * \sigma[z_t, n]$$

Where:

| | |
|---|---|
| $\mu$ | The moving average. |
| $z_t$ | The spread. |
| $n(= 20)$ | The number of days in the smoothing period. |
| $\sigma(= 2)$ | The number of standard deviations. |
| $\sigma[S, n]$ | The rolling standard deviation over the last $n$ periods of $S$. |

The mean-reverting spread is defined by [23]:

$$z_t = y_{1,t} - \lambda y_{2,t}$$

Where:

| | |
|---|---|
| $y_{1,t}$ | The closing price of the first stock, i.e., the USD/GBP exchange rate. |
| $y_{1,t}$ | The closing price of the second stock, i.e., the FTSE100. |
| $\lambda(= 2)$ | The hedge ratio. It indicates how many stocks of $y_{2,t}$ need to be sold to buy one stock of $y_{1,t}$, and vice versa to keep the spread mean-reverting. |

If the spread is low, i.e.,

$$z_t < LB$$

The first stock is undervalued and the second overvalued. This indicates to traders they should buy the spread, meaning they should short-sell the first stock and buy the second. The opposite is also true when the spread is high, i.e.,

$$z_t > UB$$

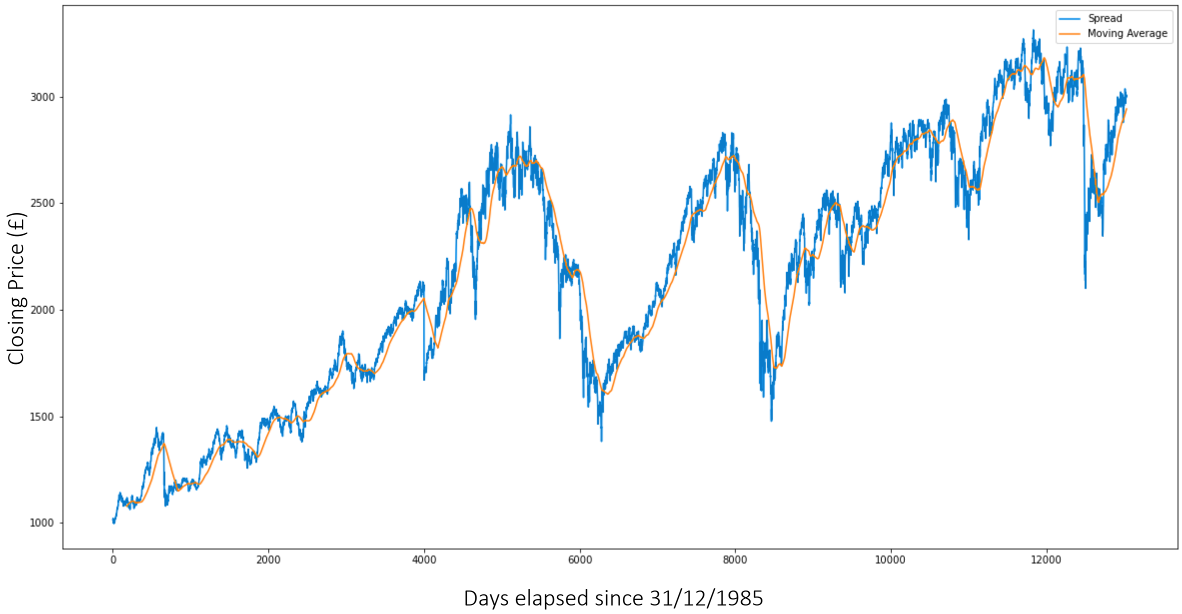


*Figure 5: Graph of the Spread of the USD/GBP and FTSE100 Closing Prices from the Equation* $z_t = y_{1,t} - 2y_{2,t}$


*Figure 6: Graph with an Example of Bollinger Bands*

Hurst Exponent

Investors use the Hurst Exponent (H) as a measure for the long-term memory of a time-series, that is, to measure the degree to which a time-series deviates from a random walk [24]. The scalar value represents the relative tendency of a time-series to either regress strongly to the mean (mean-reverting pattern), or to move in a certain direction (trending pattern) [24].

Based on the value of H, a time-series can be classified as one of the following:

| | |
|---|---|
| $0 < H < 0.5$ | A mean reverting series. The closer H is to $0$, the stronger the mean-reversion process. |
| $H = 0.5$ | A geometric random walk. |
| $0.5 < H < 1$ | A trending series. The closer H is to 1, the stronger the trend. |

There are multiple ways to calculate H, but here H is based on estimating the rate of the diffusive behaviour based on the variance of the logarithm of the spread.

First, define $x_t$ as the logarithm of stock prices $S_t$ of the spread:

$$x_t = \log(S_t)$$

The variance for an arbitrary lag (denoted by $\tau$) can be expressed as:

$$Var(\tau) = \langle |x_{t+\tau} - x_t|^2 \rangle$$

Since stock prices usually do not follow a Geometric Brownian Motion (GBM), if there is some deviation from a random walk, the variance is not proportional to the lag, but it instead gains an exponent [24]. The new relationship is therefore:

$$Var(\tau) \sim \tau^{2H}$$

After running the algorithm over the following lags:

| **Lag ($\tau$)** | **Hurst Exponent (H)** |
|---:|---:|
| 20 | 0.4817 |
| 100 | 0.4564 |
| 300 | 0.4485 |
| 500 | 0.4471 |
| 1000 | 0.4401 |
| 1200 | 0.4336 |
| 1500 | 0.4166 |

*Table 2: Variance of the Hurst Exponent with Lag*

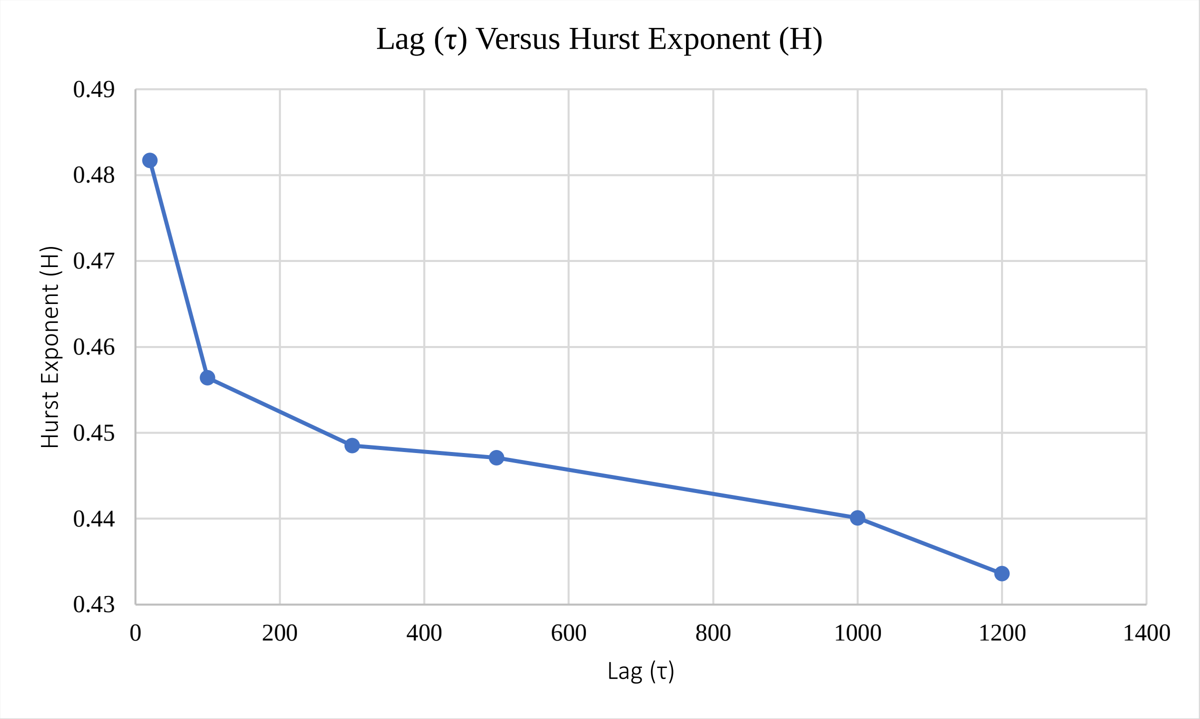


*Figure 7: Variance of the Hurst Exponent with Lag*

The value of H for these lags are all less than 0.5, indicating the spread is a mean-reverting time-series.

<u>Half-Life Test</u>

The half-life test measures how long it takes for the time-series to revert to half its initial deviation from the mean [1]. Investors use this as a selection criterion to only trade a time-series whose half-life is no longer than some specified period or it could be used as a time limit after entering a trade [1]. The Ornstein-Uhlenbeck process $x_t$ is a mean-reverting process defined by the stochastic differential equation [1] [24]:

$$dx_t = -\theta(\mu - x_t)dt + \sigma^2 dW_t$$

Where:

| | |
|---|---|
| $\theta$ | The speed of the mean-reversion. |
| $\mu$ | The mean value of the process. |
| $x_t$ | Brownian motion with unit variance parameter. |
| $\sigma^2$ | Variance of the process. |
| $W_t$ | A Wiener process. |

$\theta$ can be used to calculate the half-life of the mean-reverting process:

$$t_{1/2} = \frac{\log(2)}{\theta}$$

In this project, the half-life of the spread was 1,277 days. This half-life was used as an initial value of the lookback period to compute moving averages and standard deviations to determine trading entries and exits [24].

<u>Buy and Hold Strategy</u>

Buy and hold strategy is a long-term passive strategy where investors keep a relatively stable portfolio over time, regardless of short-term fluctuations [5] [6]. This strategy involves riding out any peaks or troughs in the stock closing price, rather than trying to swing trade the price movement [6]. A buy and hold strategy entails buying stocks or other securities and not selling them for long periods of time [23].

As comparison to the pair's trading strategy, the buy and hold strategy was implemented on the USD/GBP closing prices from 31/12/1985 to 06/10/2021.

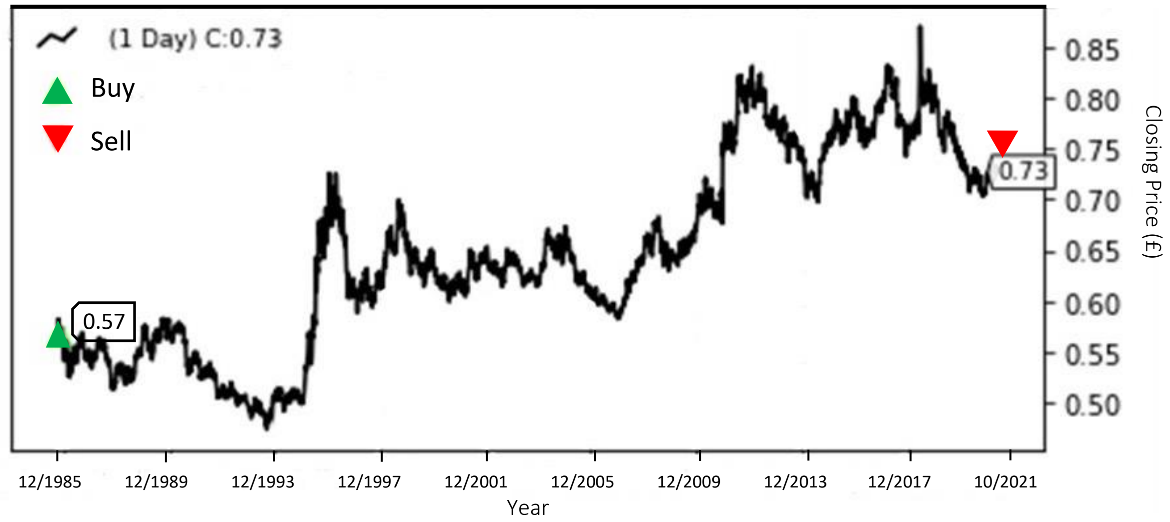


*Figure 8: Output from the Buy and Hold Strategy Implemented on the USD/GBP Exchange Rate Closing Price*

Pair's Trading Signals

For a stationary and cointegrated series, a spread between the pairs of assets is assumed to fluctuate around an equilibrium value. A long entry position was generated when the spread was lower than the LB. This meant the spread was bought. Likewise, a short entry position was generated when the spread was higher than the UB. This means, the spread was sold.

| **Controls** | |
|---|---|
| Time Frame | Daily |
| Initial Capital (GBP) | 1,000,000.00 |
| Portfolio after making first purchase (20 $*$USD/GBP and 10 $*$FTSE100) | 975,860.21 |
| Benchmark (Buy and Hold Strategy) | USD/GBP Closing Price |
| **Trading Signals** | |
| Open z-score | 1.0 |
| Close z-score | 0.0 |
| **Hurst Exponent** | |
| Window Length (Days) | 180 |
| Upper Threshold | 0.4166 |
| Lower Threshold | 0 |
| Lag (Days) | 1500 |
| **Half-Life Test** | |
| Duration (Days) | 1277 |
| **Augmented Dickey-Fuller Test** | |
| Window (Days) | 180 |
| Maximum $\rho$-value | 0.05 |

*Table 3: Parameters used in Back Testing*

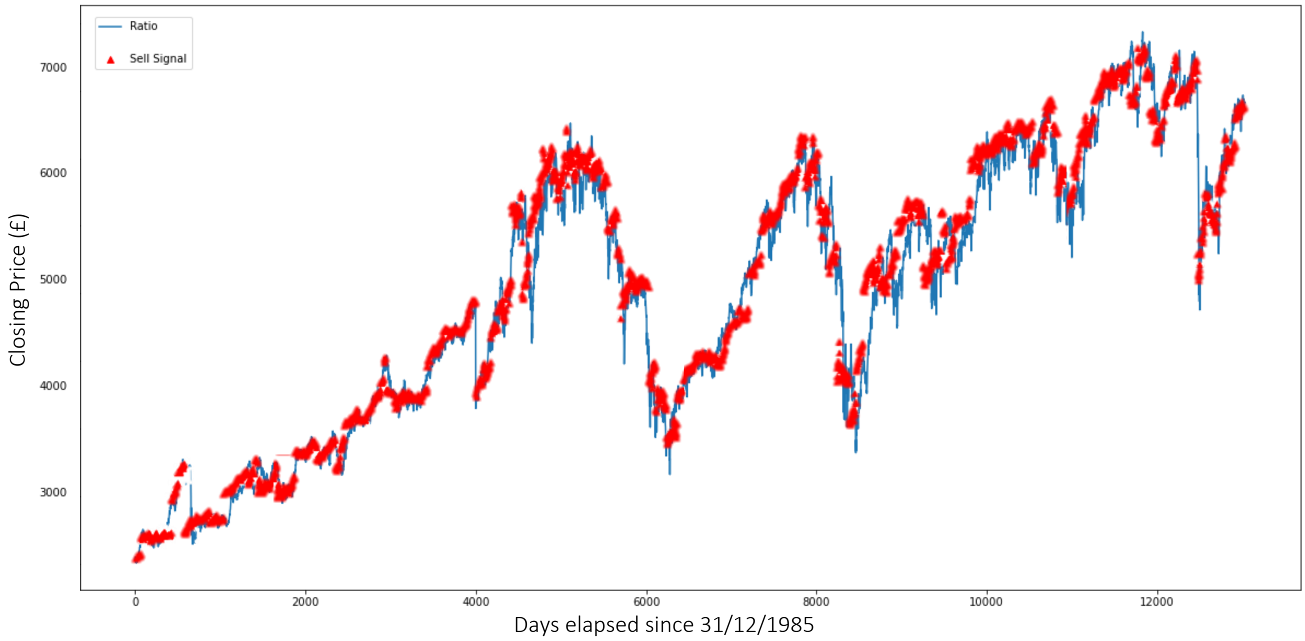


*Figure 9:Trading Signals for Selling the FTSE100 from the Pair's Trading Strategy Between the USD/GBP and FTSE100 Closing Price*

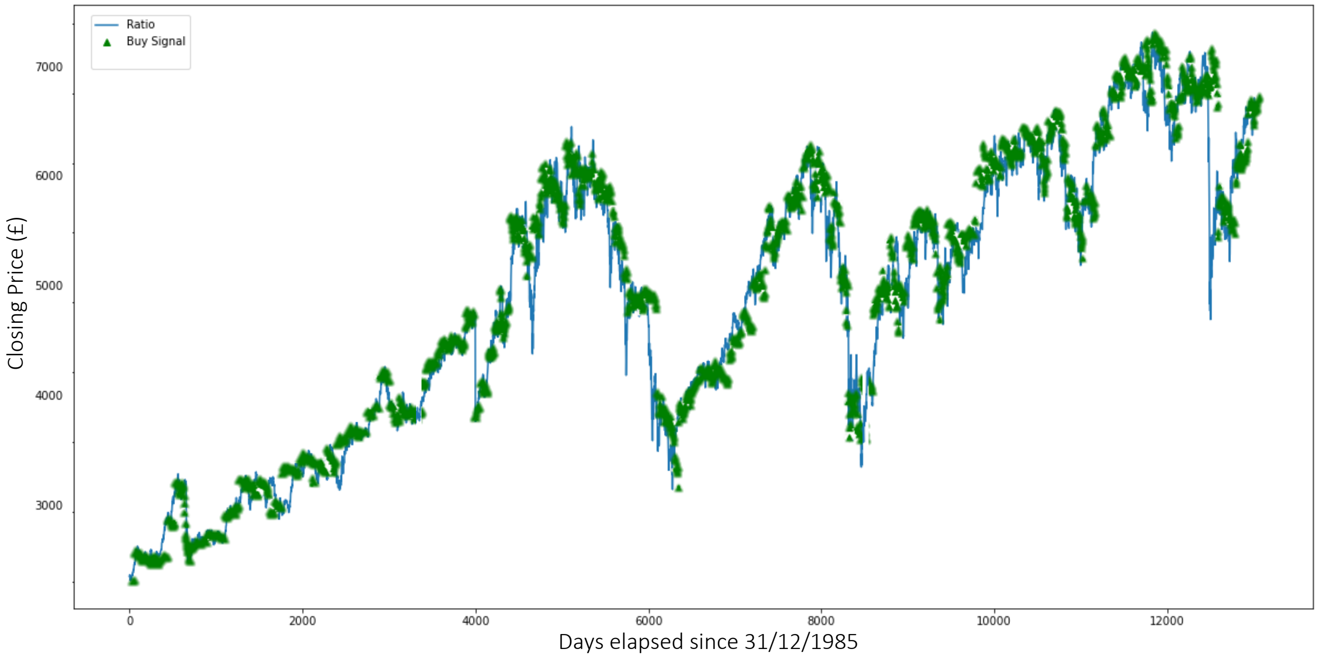


*Figure 10: Trading Signals for Buying the FTSE100 from the Pair's Trading Strategy Between the USD/GBP and FTSE100 Closing Price*

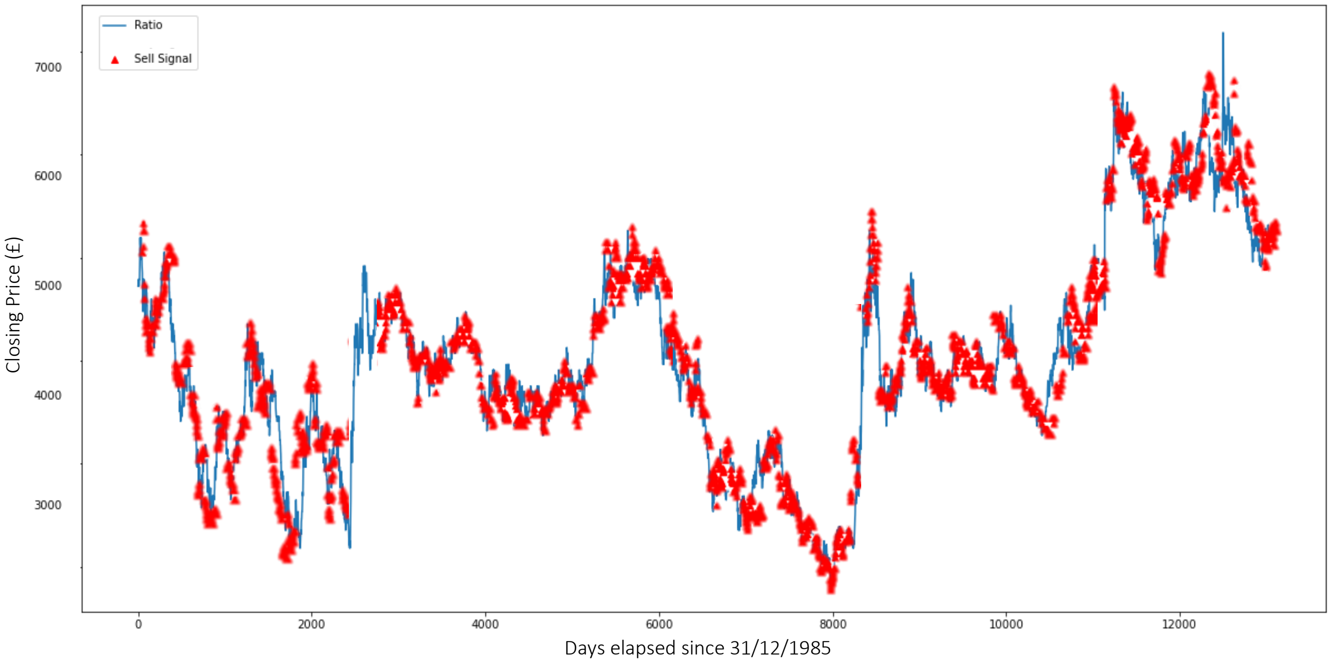


*Figure 11: Trading Signals for Selling the USD/GBP from the Pair's Trading Strategy Between the USD/GBP and FTSE100 Closing Price*

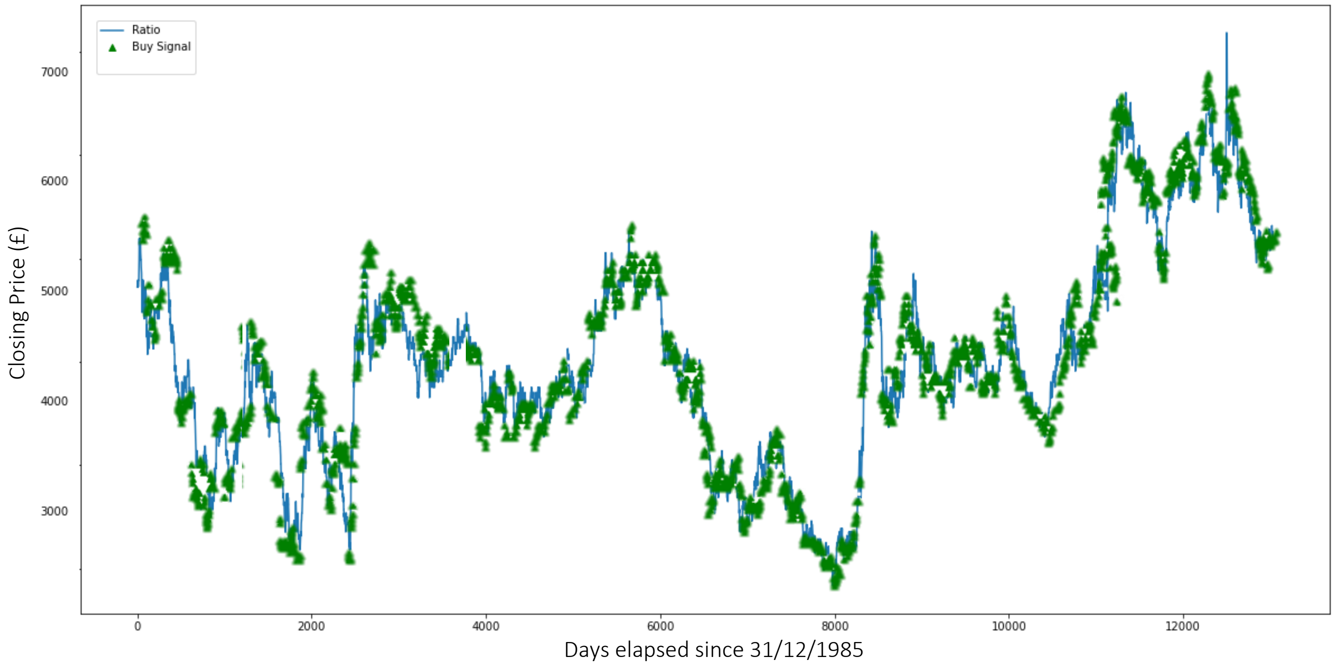


*Figure 12: Trading Signals for Buying the USD/GBP from the Pair's Trading Strategy Between the USD/GBP and FTSE100 Closing Price*

The graphs above illustrate the trading signals after implementing the pair's trading strategy on the USD/GBP closing price and FTSE100 closing price pair. There were 10,441 trading days, 5,942 of which were dedicated to buying the spread (long), and 4,499 were days where the spread was sold (short).

## Performance Evaluation

Average Annual Return and Compound Annual Growth Rate

The annual return (AR) on an investment is the return an investment provides over a certain time period, expressed as a time-weighted percentage [24]. The rate of annual return is measured against the initial amount put into the investment and represents a geometric mean rather than a simple arithmetic mean [24]. The Compound Annual Growth Rate (CAGR) shows how much the value of the portfolio grows with time, and is defined as:

$$CAGR = \left( \left[ \frac{P_f}{P_i} \right]^{\frac{1}{y}} - 1 \right) * 100$$

Where:

$P_f$ The final value of the portfolio.
$P_i$ The initial value of the portfolio.
$y$ The trading period.

To calculate the Annual Average Return (AAR), the formula then becomes:

$$\frac{\sum_{t=0}^{34} \left( \left[ \frac{P_t}{P_{t-1}} \right]^1 - 1 \right)}{35} * 100$$

Hence, $P_f$ became the value of the portfolio at time $t$, $P_i$ became the value at time $t - 1$, and $y$ became 1. The AAR was then calculated by summing the CAGR values from years 0 to 34, then divided by 35 since there were 35 years' worth of datapoints to analyse.

The AAR in this project was 3.05%. This is a relatively low percentage, but it may be since the project analysed data from the past 35 years. Had the project analysed data from a shorter period, or a shorter time frame, this yield could have been higher. For example, the AAR from 1985 to 2000 was 6.74%. The AR values from 1985 (year 0) to 2021 (year 34) are in $Table$ 5 in the appendix.

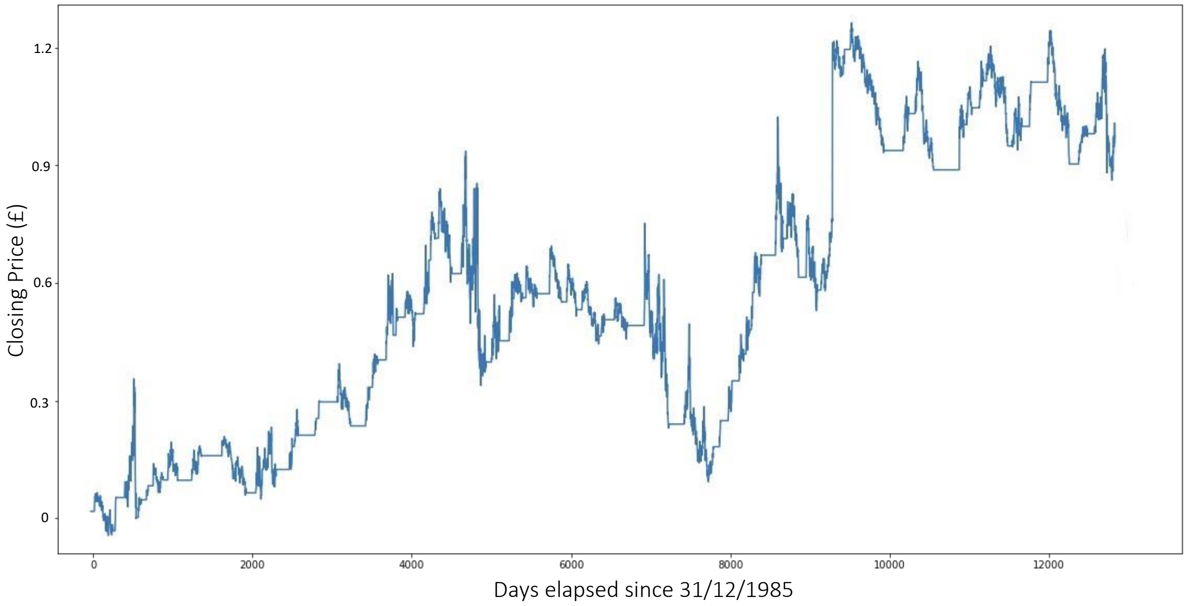


*Figure 13: Cumulative Returns from the Pair's Trading Strategy*

<u>Maximum Drawdown</u>

A Maximum Drawdown (MDD) is the maximum observed loss from a peak to trough of a portfolio before a new peak is attained [3]. The MDD is a specific measurement of drawdown that looks for the greatest movement from a high to low point before a new peak is reached [3]. However, it is important to note that it only measures the magnitude of the greatest loss, without considering the frequency of large losses. Since it measures the greatest drawdown, it does not indicate how long it took for an investor to recover from the loss [23]. A trading strategy with a large drawdown indicates a high-risk and high-volatile trading system [23].

$$MDD = \frac{V_T - V_P}{V_P} * 100$$

Where:

$V_T$ The trough value.
$V_P$ The peak value.

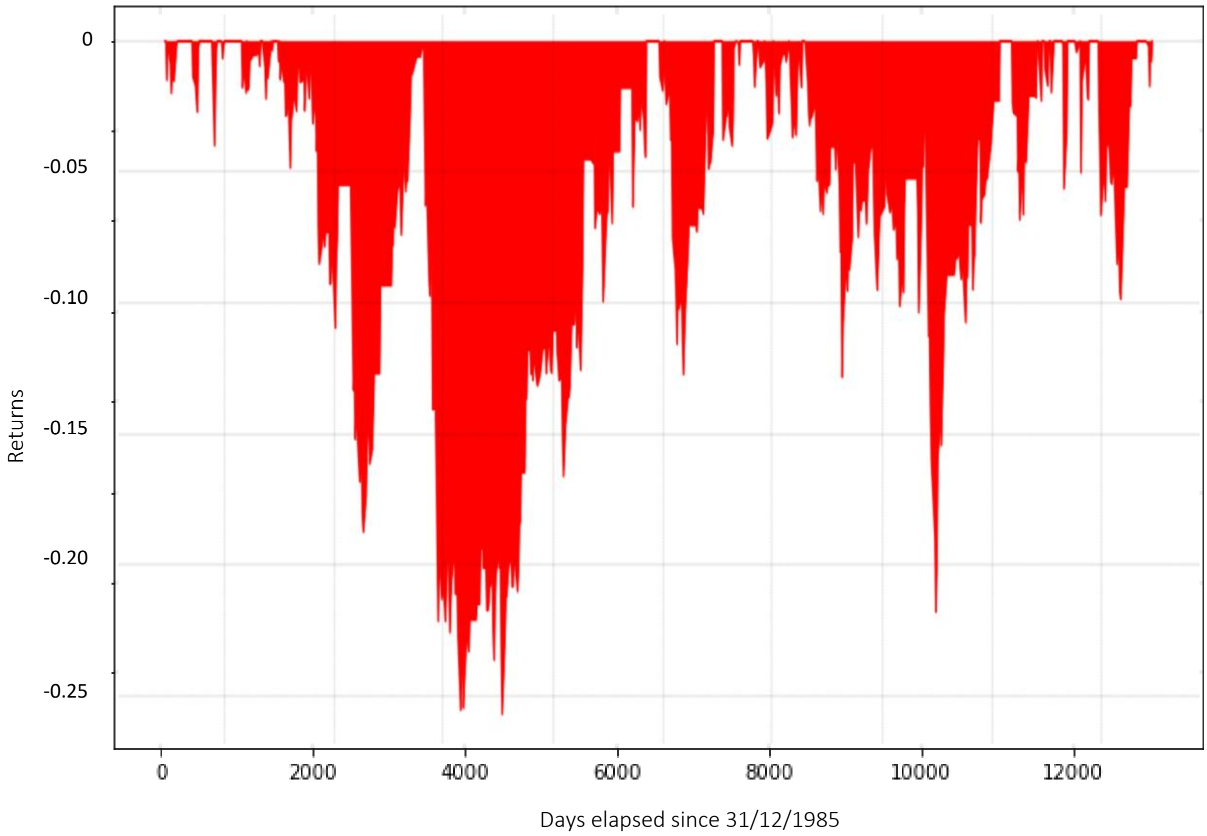


*Figure 14: Underwater Plot Showing the Drawdown over Time*

$Fig.$ 14 shows the drawdown of the strategy over time. The maximum drawdown of the pair's trading strategy would be the lowest point in the graph, or $-27.84\%$, which is acceptable.

<u>Sharpe Ratio</u>

The Sharpe Ratio is used by investors to help them understand the return of an investment compared to its risk [24]. The ratio is the average return earned in excess of the risk-free rate per unit of volatility [24]. The risk-free rate in this project was the UK 40-year treasury yield. The risk-free rate is the rate of return of an investment with no risk of loss.

$$SR = \frac{R_p - R_f}{\sigma_p}$$

Where:

$R_p$ The return of the portfolio.
$R_f$ The risk-free rate.
$\sigma_p$ The standard deviation of the portfolio's excess return.

The SR for the pair's trading strategy is 1.92, which is often deemed as "acceptable" by investors. However, if it were greater than 2.0, it could be rated as very good.

Summary of Results

| | **Average Annual Return (%)** | **Cumulative Return (%)** | **Maximum Drawdown (%)** | **Sharpe Ratio** |
|---|---|---|---|---|
| Benchmark | $1.22$ | | $-19.67$ | $0.16$ |
| FTSE100/GBP | $3.05$ | $109.88$ | $-27.84$ | $1.92$ |

*Table 4: Summary of Results from Back Testing*

On average, the pair's trading strategy yielded a greater Annual Average Return (3.05%) versus the benchmark (1.22%). Furthermore, the pair's trading strategy had a higher Sharpe Ratio than the buy and hold strategy ($1.92$ versus $0.16$). However, the pair's trading strategy suffered from a higher maximum drawdown than the benchmark ($-27.84\%$ versus $-19.67\%$). It could be argued that the pair's trading strategy is more successful than the buy and hold strategy.

**Conclusion**

This project explored the feasibility of using the Temporal Convolutional Network algorithm to use the USD/GBP exchange rate to predict the closing price of the FTSE100 from $31/12/1985$ to $06/10/2021$. With $13{,}208$ datapoints, the algorithm yielded an MAPE of $10.04\%$, implying an accuracy of $89.96\%$. This is an acceptable value but could have been higher given more datapoints. This, however, may have been difficult to do since Yahoo! Finance was the only source which had data from 1985 for both the USD/GBP and FTSE100 closing prices.

After running the machine learning algorithm, a pair's trading strategy was created using statistical arbitrage. This pair's trading strategy was then compared to a buy and hold strategy with the USD/GBP closing prices. Besides the maximum drawdown value (the pair's trading strategy had a higher value of $-27.84\%$ as opposed to $-19.67\%$ for buy and hold), the pair's trading strategy generally performed better than the buy and hold strategy as it had a higher Sharpe Ratio ($1.92$ for the pair's trading strategy compared to $0.16$ for the buy and hold strategy) and higher average returns ($3.05\%$ for the pair's trading strategy versus $1.22\%$ for the buy and hold strategy).

An average return of $3.05\%$ for the pair's trading strategy is low, especially if compared to the Standard & Poor 500 (S&P500), a popular standard benchmark, whose AAR is 8% adjusted for inflation. This may be because the value $3.05\%$ is obtained by averaging 36 years of average returns (see $Table$ 5 in the appendix). However, because the nature of the Forex market is much different to the stock market, it may not be feasible to compare the two average annual returns. If this project focused on data by the minute and the number of days were reduced, the perhaps the pair's trading strategy would have a higher average annual return.

In $Table$ 5, there are several years to note, which have been highlighted as they are either abnormally low or abnormally high:

1998 Since 1998, the UK has been running a trade deficit. Imports could be needed to produce the country's exports or sales overseas. An increase in exports contributes positively to economic growth as there is an increase in foreign sales by domestic companies. The growing economy could have led to higher stock market prices [8].

2001/2 The performance of the UK economy in 2001 was satisfactory as it had managed to survive a major international crisis. The downturn of economic activity in the USA led to problems globally, and although the Bank of England responded by relaxing monetary policy, it was not enough to offset recessionary forces which became evident in 2002 [22] [23].

2009 The biggest annual fall in GDP prior to 2020 occurred in 2009, when the UK economy contracted by 4.1% at the height of the global financial crisis of the late 2000s [8].

2017 The UK decided to leave the European Union in 2017. The fall in the Pound Sterling boosted earnings since many of the FTSE100 companies had large operations abroad [9].

2020/1 The UK economy shrank by a record 9.8% in 2020, due to the economic fallout due to the Covid-19 pandemic.

## Further Exploration

Kalman Filter

A Kalman Filter is an algorithm that uses a series of measurements observed over time and produces an estimate of unknown variables [24]. This algorithm could be used to estimate the value of the Hedge Ratio instead of using Ordinary Least Squares. Kalman Filters are ideal for systems that are continuously changing [24]. They have the advantage that they are light on memory (they don't need to keep any history other than the previous state).

Clustering Methods

The project could also be further extended to multiple sectors and use clustering techniques such as K-Means Clustering or Density-Based Spatial Clustering of Applications with Noise (DBSCAN), to identify suitable clusters of stock markets and currency exchange rates, and then implementing the pair's trading strategy on those stock / currency pairs.

Basket Trading

More stock exchanges could be implemented in the project, and a basket trading strategy could be implemented to trade multiple stock exchanges alongside a single currency exchange rate. For example, the USA has three major stock exchanges – the NASDAQ composite, the Dow Jones Industrial Average, and S&P 500. These three stock exchanges could be traded alongside the US Dollar Index.

**References**

[1] A. Constandinou, “Quant Post 3.1: A guided path into mean reversion,” *Medium*, 25-Jul-2018. [Online]. Available: https://medium.com/@constandinou.antonio/quant-post-3-1-a-guided-path-into-mean-reversion-8b33b3c279e4. [Accessed: 23-Nov-2021].

[2] A. Hayes, “Bollinger Band,” *Investopedia*, 28-Jul-2021. [Online]. Available: https://www.investopedia.com/terms/b/bollingerbands.asp. [Accessed: 23-Nov-2021].

[3] A. Hayes, “Maximum Drawdown (MDD),” *Investopedia*, 19-May-2021. [Online]. Available: https://www.investopedia.com/terms/m/maximum-drawdown-mdd.asp. [Accessed: 23-Nov-2021].

[4] A. Hussain, “Market Neutral Definition,” *Investopedia*, 08-Sep-2021. [Online]. Available: https://www.investopedia.com/terms/m/marketneutral.asp. [Accessed: 23-Nov-2021].

[5] A. Levy, “What is a Buy-and-Hold Strategy in Investing?,” *The Motley Fool*, 25-Jun-2009. [Online]. Available: https://www.fool.com/investing/how-to-invest/stocks/buy-and-hold-strategy/. [Accessed: 23-Nov-2021].

[6] B. Beers, “Buy and Hold Definition,” *Investopedia*, 19-May-2021. [Online]. Available: https://www.investopedia.com/terms/b/buyandhold.asp. [Accessed: 23-Nov-2021].

[7] B. Or, “Temporal convolutional networks, the next Revolution for Time-series?,” *Medium*, 15-Aug-2020. [Online]. Available: https://towardsdatascience.com/temporal-convolutional-networks-the-next-revolution-for-time-series-8990af826567?gi=d392fa43a60a. [Accessed: 23-Nov-2021].

[8] D. Clark, “UK: GDP growth 2000-2018,” Statista, 12-Aug-2021. [Online]. Available: https://www.statista.com/statistics/281734/gdp-growth-in-the-united-kingdom-uk/. [Accessed: 09-Dec-2021].

[9] D. Jordan, “UK stock markets close at a record high,” BBC News, 29-Dec-2017. [Online]. Available: https://www.bbc.com/news/business-42512023. [Accessed: 09-Dec-2021].

[10] D. P. Palomar, “Pair’s Trading,” in *Portfolio Optimization with R*.

[11] “Drawdown Forex,” *Compare Forex Brokers*, 01-Nov-2021. [Online]. Available: https://www.compareforexbrokers.com/forex-trading/drawdown/. [Accessed: 23-Nov-2021].

[12] E. Lewinson, “Introduction to the Hurst exponent,” *Medium*, 26-May-2021. [Online]. Available: https://towardsdatascience.com/introduction-to-the-hurst-exponent-with-code-in-python-4da0414ca52e. [Accessed: 23-Nov-2021].

[13] Erica, “A Guide to Conducting Cointegration Tests,” *Aptech*, 12-May-2021. [Online]. Available: https://www.aptech.com/blog/a-guide-to-conducting-cointegration-tests/. [Accessed: 23-Nov-2021].

[14] J. Chen, “Annual Return,” *Investopedia*, 19-May-2021. [Online]. Available: https://www.investopedia.com/terms/a/annual-return.asp. [Accessed: 23-Nov-2021].

[15] J. Chen, “Pair’s Trade,” *Investopedia*, 19-May-2021. [Online]. Available: https://www.investopedia.com/terms/p/pair’strade.asp. [Accessed: 23-Nov-2021].

[16] J. Corrius, “Simple Stationarity Tests on Time Series,” *Medium*, 09-Oct-2018. [Online]. Available: https://medium.com/bluekiri/simple-stationarity-tests-on-time-series-ad227e2e6d48. [Accessed: 23-Nov-2021].

[17] J. Fernando, “Sharpe Ratio,” *Investopedia*, 20-Nov-2021. [Online]. Available: https://www.investopedia.com/terms/s/sharperatio.asp. [Accessed: 23-Nov-2021].

[18] J. Steffen, P. Held, and R. Kruse, Otto von Guericke University Magdeburg , Magdeburg, tech.

[19] K. Thune, “What Is the Buy and Hold Strategy?,” *The Balance*, 09-Jul-2021. [Online]. Available: https://www.thebalance.com/what-is-buy-and-hold-2466543. [Accessed: 23-Nov-2021].

[20] “Mean Absolute Percentage Error,” *Wikipedia*, 21-Nov-2021. [Online]. Available: https://en.wikipedia.org/wiki/Mean_absolute_percentage_error. [Accessed: 23-Nov-2021].

[21] P. Turner, “Review of the UK economy in 2003.,” Loughborough University, 11-Aug-2019. [Online]. Available: https://repository.lboro.ac.uk/articles/journal_contribution/Review_of_the_UK_economy_in_2003_/9493970. [Accessed: 09-Dec-2021].

[22] P. Turner, “The UK economy in 2001.,” Loughborough University, 11-Aug-2019. [Online]. Available: https://repository.lboro.ac.uk/articles/journal_contribution/The_UK_economy_in_2001_/9494066/1. [Accessed: 09-Dec-2021].

[23] Stephanie, “Unit Root: Simple Definition, Unit Root Tests,” *Statistics How To*, 01-Jan-2021. [Online]. Available: https://www.statisticshowto.com/unit-root/. [Accessed: 23-Nov-2021].

[24] T. Babb, “How a Kalman Filter Works, in Pictures,” *Bzarg*, 11-Aug-2015. [Online]. Available: https://www.bzarg.com/p/how-a-kalman-filter-works-in-pictures/#:~:text=Kalman%20filters%20are%20ideal%20for,time%20problems%20and%20embedded%20systems. [Accessed: 23-Nov-2021].

[25] "Tensor," *Wikipedia*, 16-Nov-2021. [Online]. Available: https://en.wikipedia.org/wiki/Tensor. [Accessed: 23-Nov-2021].

[26] P. Ryser-Welch, "Is there an ideal ratio between a training set and validation set? which trade-off would you suggest?," ResearchGate, 22-Feb-2021. [Online]. Available: https://www.researchgate.net/post/Is-there-an-ideal-ratio-between-a-training-set-and-validation-set-Which-trade-off-would-you-suggest. [Accessed: 16-Dec-2021].

[27] P. Inman, "Foreign investors own 66% of UK-listed shares, analysis shows." The Guardian, 07-Jun-2021. [Online]. Available: https://www.theguardian.com/business/2021/jun/07/foreign-investors-own-66-of-uk-listed-shares-analysis-shows. [Accessed: 16-Dec-2021].

## Appendix

| Year | Number of Trading Days | | Annual Average Return (%) |
|---|---|---|---|
| | Buy | Sell | |
| 1985/6* | 176 | 124 | 7.3282 |
| 1987 | 209 | 90 | 8.4331 |
| 1988 | 163 | 125 | 1.68693 |
| 1989 | 223 | 86 | 1.63874 |
| 1990 | 115 | 170 | 12.8882 |
| 1991 | 160 | 147 | −5.6256 |
| 1992 | 163 | 138 | 5.5158 |
| 1993 | 191 | 127 | 8.7067 |
| 1994 | 139 | 162 | 12.4891 |
| 1995 | 206 | 87 | −8.0505 |
| 1996 | 171 | 130 | 12.6773 |
| 1997 | 212 | 92 | −11.9782 |
| 1998 | 181 | 119 | 20.1912 |
| 1999 | 175 | 106 | 16.9164 |
| 2000 | 127 | 164 | 18.2910 |
| 2001 | 108 | 171 | −12.3672 |
| 2002 | 110 | 173 | −18.3370 |
| 2003 | 188 | 87 | −8.0622 |
| 2004 | 164 | 119 | 12.6690 |
| 2005 | 224 | 97 | 7.1973 |
| 2006 | 192 | 96 | 16.3393 |
| 2007 | 166 | 111 | 12.6101 |
| 2008 | 115 | 175 | 5.2103 |
| 2009 | 187 | 95 | −35.5990 |
| 2010 | 206 | 120 | 25.6762 |
| 2011 | 128 | 153 | 12.2943 |
| 2012 | 181 | 113 | −10.0460 |
| 2013 | 172 | 132 | 9.9472 |
| 2014 | 162 | 130 | 14.9833 |
| 2015 | 139 | 140 | −1.7031 |
| 2016 | 166 | 116 | −7.1843 |
| 2017 | 160 | 123 | 17.4621 |
| 2018 | 128 | 171 | 11.5963 |
| 2019 | 190 | 121 | −15.0614 |
| 2020 | 135 | 127 | −18.7180 |
| 2021 | 110 | 62 | −10.1321 |
| | | | |
| **Sum** | 5942 | 4499 | 109.88 |
| **Average** | 165 | 125 | 3.12 |

*Table 5: Average Annual Return for the Pair's Trading Strategy*

*There is only one datapoint in 1985