Pairs Trading Strategy

Independent Work Report

by

NG Tik Sang

Advised by

David Paul ROSSITER

Submitted in partial fulfillment
of the requirements for COMP 4971C
in the Department of Computer Science and Engineering
The Hong Kong University of Science and Technology
2019 − 2020

Date of Submission: May 19, 2020

Contents

Chapter 1

Introduction

1.1 Overview

There is a large number of listed securities in different stock markets all over the world, and most of these companies are financially linked to one another. Thus, when one stock changes, others relevant stocks are also likely to be affected, which poses an obstacle when making a trade. Among all the trading strategies, pairs trading is one of the most common approach which is market neutral. This strategy makes use of this characteristic which reduces the unusual risk in trading. Therefore, this project aims to further study the pairs trading strategy, and optimize the trading returns by tuning some of the parameters in the trading algorithm.

Pairs trading keeps track of two historically correlated securities. It is expected that the difference in stock prices (also known as *spread*) should remain constant. Figure 1.1 shows two stocks of which their spread is mostly the same with some minor fluctuations. The strategy is best deployed when there is a significant divergence in the spread, which can be caused by temporary changes in supply and demand. It is assumed that the stock prices of the two securities revert to their historical trends (i.e. mean-reverting).

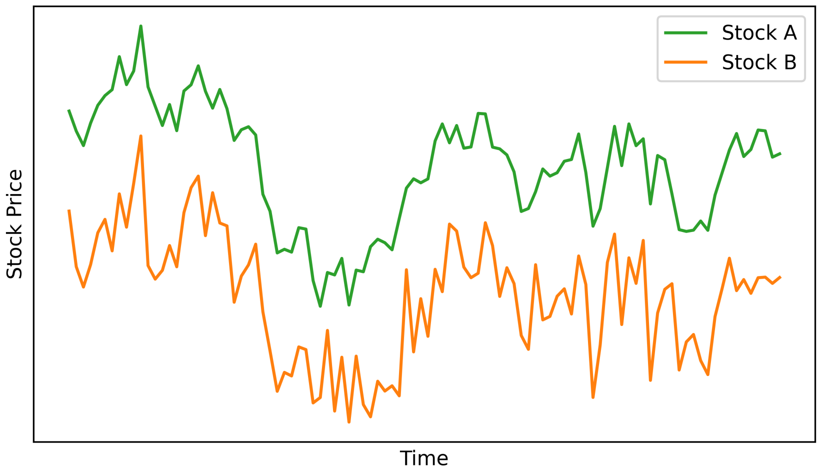

Figure 1.1: Example of Pairs Trading

If two stock prices deviate, the pairs trade suggests to sell the stock that moves up (i.e. to short the outperforming stock) and buy the stock that moves down (i.e. to long the underperforming stock). When the two stock prices converge back to the usual levels, all the positions should be closed, which means buying back the outperforming stock and selling the underperforming stock. In the case where both stocks move up or down together, the spread does not change and no trading is made. Therefore, this strategy not only makes a profit, but also ensures traders to minimize the potential losses.

1.2 Software Framework

In this project, JupyterLab, which is a Python-based IDE, was used for data analysis, optimization and backtesting. The following libraries were also imported to the source code to aid the program development.

- `matplotlib`, for data visualization
- `numpy`, for mathematical computation
- `pandas`, for data management
- `statsmodels`, for cointegration testing

1.3 Source of Data

The stock data was collected from the NASDAQ stock market, since the historical data provided is more complete and there are more types of companies, which facilitates an accurate analysis. The data was first scraped using the Google Finance API and was later exported as several `csv` files for the Python program to read and analyze.

Chapter 2

Methodology

2.1 Preliminary Data Processing

Since pairs trading requires a mutual economical correlation between two securities, it is necessary to restrict the industry where the securities fall in, so that the trading performance can be easily evaluated. After considering the completeness and variety of data between individual sectors, the stocks related to the technical field were chosen as the targets in this study, which include `AAPL` (Apple Inc.), `ADBE` (Adobe Inc.), `AMZN` (Amazon.com, Inc.), `CSCO` (Cisco Systems, Inc.), `GOOGL` (Alphabet Inc.), `INTC` (Intel Corporation), `MSFT` (Microsoft Corporation) and `NVDA` (NVIDIA Corporation). Based on the data availability for all these stocks, the study only focuses on all the trade dates from Jan 1, 2005 to Dec 31, 2019 inclusive, which is a period of 15 consecutive years.

Table 2.1 shows an example of raw stock data which is scraped from Google Finance. For simplicity, only the entries from the Close column were considered, and other columns were dropped throughout this study.

	Open	**High**	**Low**	**Close**	**Volume**
Date					
2001-01-03 16:00:00	3.67	4.02	3.63	4.00	4750700
2001-01-04 16:00:00	3.95	3.95	3.61	3.66	4519300
2001-01-05 16:00:00	3.68	3.95	3.59	3.63	6940600
2001-01-06 16:00:00	3.75	3.82	3.34	3.39	6856500
2001-01-07 16:00:00	3.41	3.61	3.41	3.50	4108800
...	...	...	...	...	...
2020-05-13 16:00:00	312.15	315.95	303.21	307.65	50155639
2020-05-14 16:00:00	304.51	309.79	301.53	309.54	39732269
2020-05-15 16:00:00	300.35	307.9	300.21	307.71	41587094
2020-05-18 16:00:00	313.17	316.5	310.32	314.96	33843125
2020-05-19 16:00:00	315.03	318.52	313.01	314.14	25432385

Table 2.1: Stock Data of Apple Inc.

2.2 Cointegration Testing

Although the domain of data has been confined to the technology sector, their mutual economic relationships is yet to be determined. The parameter *cointegration* is used for determining the statistical connection between two time series. In this study, the Engle-Granger two-step method was utilized and the following hypotheses were tested between the time series, which represent the two stock prices.

$$\begin{cases} H_0: & \text{There is no cointegrating relationship;} \\ H_1: & \text{There is cointegrating relationship,} \end{cases}$$

where H_0 is the null hypothesis and H_1 is the alternative hypothesis.

To show that two stocks are cointegrated, the null hypothesis should be rejected, which means that the p-value has to be lower than a predefined level of significance. Typically, when the p-value is smaller than 0.05, it indicates a strong statistical evidence against the null hypothesis, and the alternative hypothesis should be accepted. Therefore, this study has set 0.05 as a threshold to screen out all the cointegrated stock pairs, and the asymptotic p-values were calculated based on the MacKinnon's approximate used in the Augmented Dickey-Fuller unit root test.

Figure 2.2 shows the comparison of p-values of different stock pairs, in which the greenish grids represent the pairs with a low p-values and the reddish grids represent those with a high p-values. Table 2.3 lists out the top five stock pairs sorted by the ascending order of their p-values. The complete testing result can be found in Appendix A.

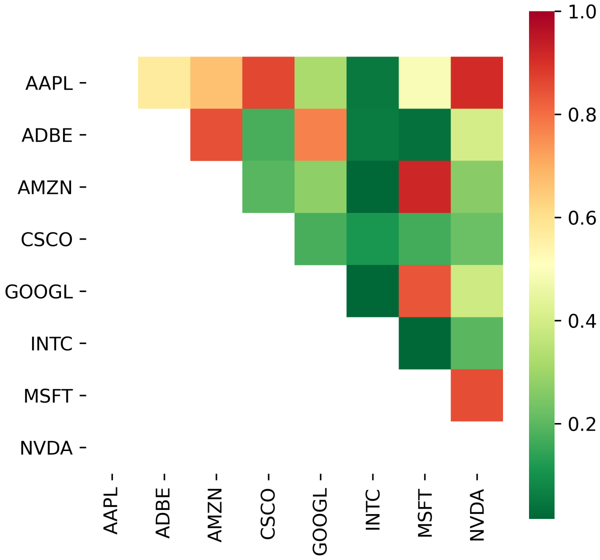

Figure 2.2: Comparison of p-values

	Stock Pairs	p-values
1	(AMZN, INTC)	0.0150584
2	(INTC, MSFT)	0.0152943
3	(GOOGL, INTC)	0.0170111
4	(ADBE, MSFT)	0.0358782
5	(AAPL, INTC)	0.0501868

Table 2.3: Stock Pairs with Low p-values

Based on the cointegration testing results, only the stock pairs with small p-values would be considered for further analysis.

2.3 Normalization

In this section, the stock pair (`AMZN`, `INTC`) was used for illustration. In general, the same adjustment could be applied to all other stock pairs.

When two stocks are cointegrated, their changes in stock prices should be aligned. This scenario is best exemplified in Figure 2.4 between 2008 and 2009, where a series of up-trends and downtrends were captured.

Figure 2.4: Stock Prices of `AMZN` and `INTC` (2005 – 2009)

With such characteristic, trades can be made when the spread between the two stocks varies from the normal range. To determine the usual separation, ratio of the stocks was studied, which is plotted in Figure 2.5 with the red dotted line being the mean ratio across time.

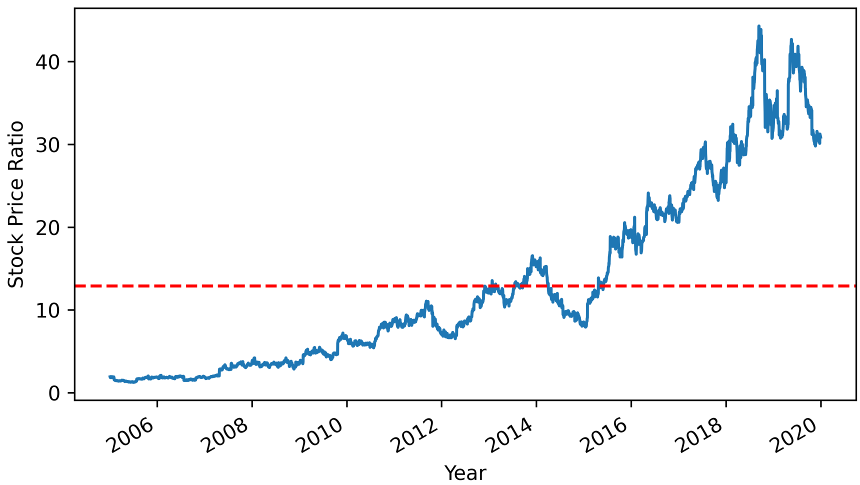

Figure 2.5: Stock Price Ratio of `AMZN` to `INTC`

When making a trade, the actual ratio does not give a precise statistical information. Instead, the relative movement of the ratio should be studied. Hence, normalization was performed on the ratio (x) using the standard score (z), which is defined as

$$z = \frac{x - \mu}{\sigma},$$

where μ denotes the population mean (of all the ratios) and σ denotes the population standard deviation (of all the ratios). The resulting plot is given in Figure 2.6 and the ratio values were normalized.

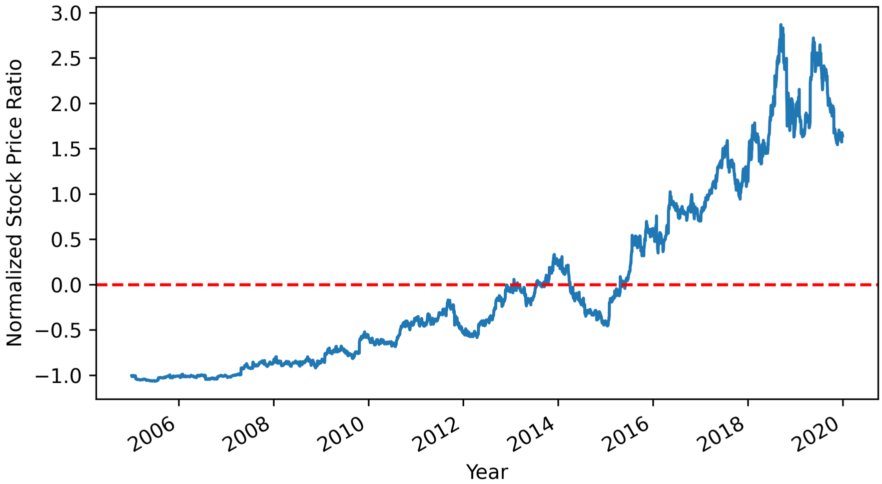

Figure 2.6: Normalized Stock Price Ratio of `AMZN` to `INTC`

However, since the ratio is generally increasing in an exponential scale, the usual range is also expected to grow across time, so the normalization should take the time frame into consideration. Thus, a moving standard score ($\tilde{z}$) can be used instead, which is defined as

$$\tilde{z}(t_{short}, t_{long}) = \frac{\tilde{x}(t_{short}) - \tilde{\mu}(t_{long})}{\tilde{\sigma}(t_{long})}, \quad t_{short} < t_{long},$$

where $\tilde{x}$ denotes a short-term moving average (of all the ratios), $\tilde{\mu}$ denotes a long-term moving average (of all the ratios) and $\tilde{\sigma}$ denotes a long-term moving standard deviation (of all the ratios).

The use of a short-term moving average to replace a single value is to reduce the impact caused by any exceptional recent changes which may not be of interest.

In Figure 2.7, t_{short} was set as 3 and t_{long} was set as 200 as an example.

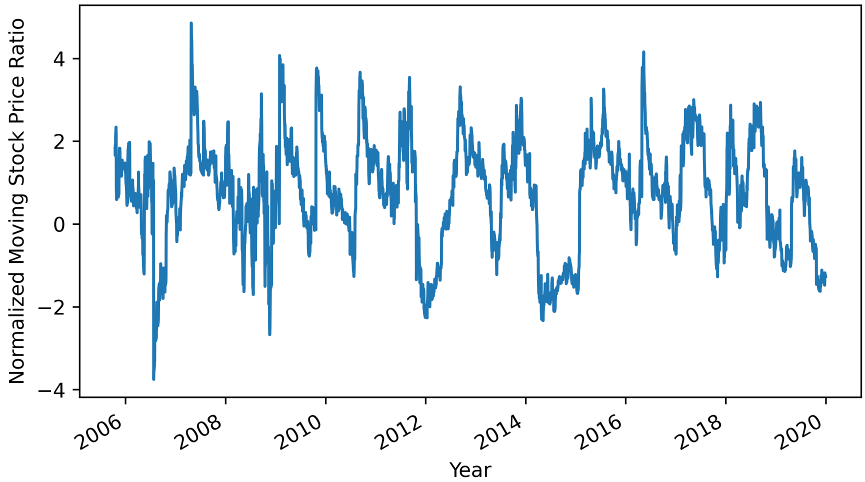

Figure 2.7: Normalized Moving Stock Price Ratio of `AMZN` to `INTC`

After this normalization, the graph is stationary, centered and bounded.

2.4 Trading Strategy

Before any trading, a fixed amount of assets is allocated. The trading signal can be defined with reference to the standard score of the stocks ratio. The variables z_{low} and z_{high} are two benchmarks for performing certain trading decisions, where $z_{low} < z_{high}$.

1. If the standard score is greater than z_{high}, then ratios of the stocks are sold as many as possible.
2. If the standard score is smaller than $-z_{high}$, then ratios of the stocks are bought as many as possible.
3. If the standard score lies between $-z_{low}$ and z_{low}, then all the purchased stocks are cleared, and the profit/loss is recorded.

The third point above aims to prevent unnecessary trades since the transaction fee could be the overheads, especially when the ratio is fluctuating around the mean.

Chapter 3

Optimization

3.1 Parameters

Based on the previous data processing work, there are four variables in which their values could be optimized so that the trading profit is maximized. The following lists the possible ranges of these unknowns.

1. t_{short} (short-term time frame), ranging from 1 to 5 inclusive with a step of 1
2. t_{long} (long-term time frame), ranging from 60 to 180 inclusive with a step of 15
3. z_{low} (low cut-off standard score), ranging from 0.1 to 0.3 with a step of 0.1
4. z_{high} (high cut-off standard score), ranging from 1 to 1.5 with a step of 0.1

These ranges would be used for performing a grid search to figure out the optimal values of the four parameters.

3.2 Indicators

To evaluate the optimal values, the trading performance should be taken into account. Below are three performance indicators extracted from Investopedia.

1. **Compound Annual Growth Rate (CAGR)**
   Compound annual growth rate is the rate of return that would be required for an investment to grow from its beginning balance to its ending balance, assuming the profits were reinvested at the end of each year of the investment's lifespan.

   $$CAGR = \left(\frac{EV}{BV}\right)^{\frac{1}{n}} - 1,$$

   where EV denotes the ending value, BV denotes the beginning value and n denotes the number of years.

2. **Maximum Draw Down (MDD)**
   A maximum drawdown is the maximum observed loss from a peak to a trough of a portfolio, before a new peak is attained. Maximum drawdown is an indicator of downside risk over a specified time period.

   $$MDD = \frac{TV - PV}{PV},$$

   where TV denotes the trough value and PV denotes the peak value.

3. **Sharpe Ratio (SR)**
   The Sharpe ratio is used to help investors understand the return of an investment compared to its risk. The ratio is the average return earned in excess of the risk-free rate per unit of volatility or total risk.

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

   where R_p denotes the return of portfolio, R_f denotes the risk-free rate and σ_p denotes the standard deviation of the portfolio's excess return.

   In this study, R_f was chosen as the 13-week Treasury Bill Yield Index.

3.3 Result

Table 3.1 shows the best 5 combinations of the four parameters based on CAGR. The complete optimization result can be found in Appendix B.

	t_{short}	t_{long}	z_{low}	z_{high}	**CAGR (%)**	**MDD (%)**	**SR**
1	2.0	150.0	0.3	1.0	26.924478	-33.994605	14.196635
2	4.0	120.0	0.2	1.0	25.421691	-32.535939	13.370391
3	4.0	105.0	0.2	1.0	24.327562	-2289.472342	27.362748
4	2.0	105.0	0.3	1.0	24.022689	-40.879782	24.027737
5	4.0	135.0	0.2	1.0	23.553891	-36.517522	12.748263

Table 3.1: Optimization Result

However, since the optimization was done on a subset of all the stock pairs, there may be overfitting on the parameters. To tackle this problem, the rule on choosing the best parameter set was relaxed so that the best 5% of combinations would be considered, which is roughly the top 40 results out of all the 810 possible cases. The targeted combination should maximize both CAGR and SR but minimize MDD.

Table 3.2 shows the finalized values of the four parameters with the performance evaluation in the optimization stage.

t_{short}	t_{long}	z_{low}	z_{high}	**CAGR (%)**	**MDD (%)**	**SR**
4.0	165.0	0.2	1.0	22.041784	-32.559491	24.868090

Table 3.2: Finalized Parameters

Chapter 4

Backtesting

Based on the optimal parameters, mock trading was conducted using the pairs trading strategy. Table 4.1 shows the backtesting result for the stock pairs with p-values smaller than a threshold of 0.05.

	CAGR	**MDD**	**SR**
Benchmark	22.041784	-32.559491	24.868090
(`INTC`, `MSFT`)	22.825154	-9.645808	57.182144
(`GOOGL`, `INTC`)	14.797846	-24.269	28.211776
(`ADBE`, `MSFT`)	21.085135	-2.611080	49.344648

Table 4.1: Backtesting Result

In general, the backtesting shows a positive outcome. All the three testing stock pairs are profitable. Although their CAGR's lie on or below the benchmark, their MDD's are reduced by at least 25%, in the pair (`GOOGL`, `INTC`), and up to over 90%, in the pair (`ADBE`, `MSFT`), when compared to the benchmark. Moreover, all the SR's of the testing pairs are higher than the benchmark. Therefore, the trading strategy is expected to yield a moderate profit with a low risk.

Detailed performance of individual testing stock pair is provided as follows. The cumulative returns will be compared with the investment on a risk-free rate, which is the 13-week Treasury Bill Yield Index.

4.1 Stock Pair (INTC, MSFT)

Figure 4.2 shows that the cumulative returns of the stock pair were higher than that of the risk-free rate, and the level of returns were generally steady across time. Figure 4.3 shows that the draw down near 2018 was immediately stopped right after the occurrence.

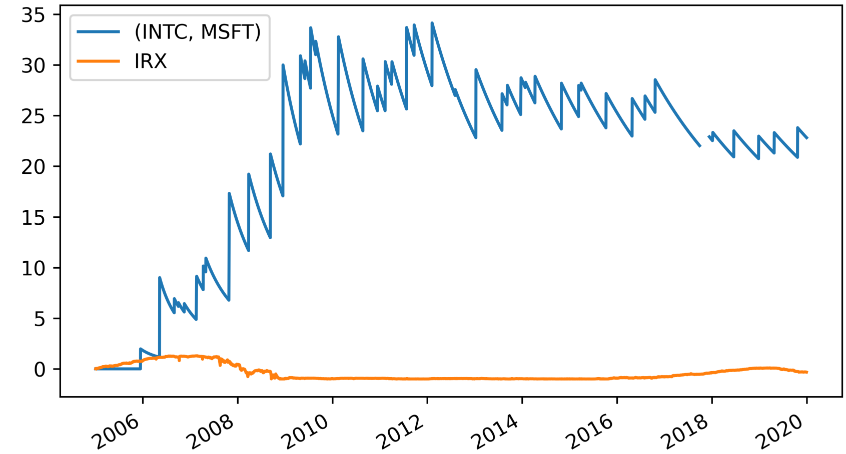

Figure 4.2: Cumulative Returns of Stock Pair (INTC, MSFT)

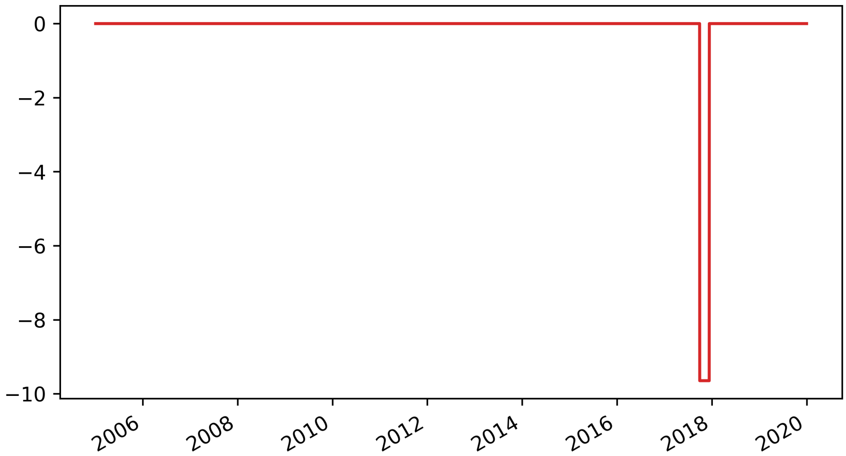

Figure 4.3: Draw Down of Stock Pair (INTC, MSFT)

4.2 Stock Pair (GOOGL, INTC)

Figure 4.4 shows that the cumulative returns of the stock pair were higher than that of the risk-free rate, but the level of returns dropped after 2013. Figure 4.5 shows that longer time was taken to stop the draw down.

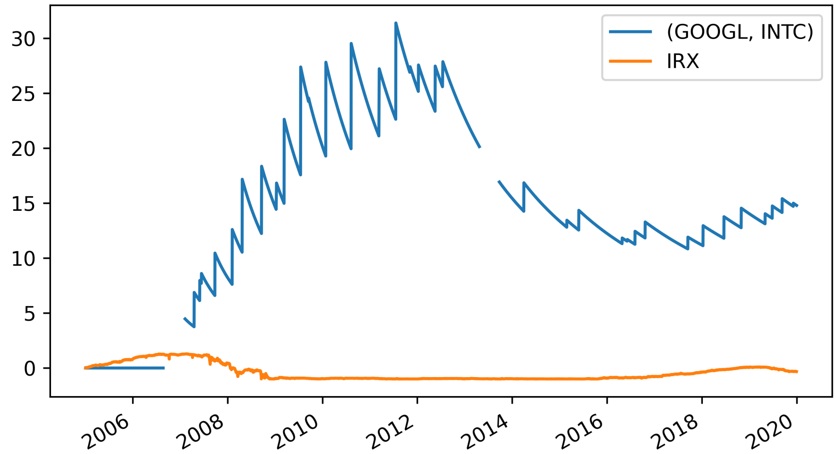

Figure 4.4: Cumulative Returns of Stock Pair (GOOGL, INTC)

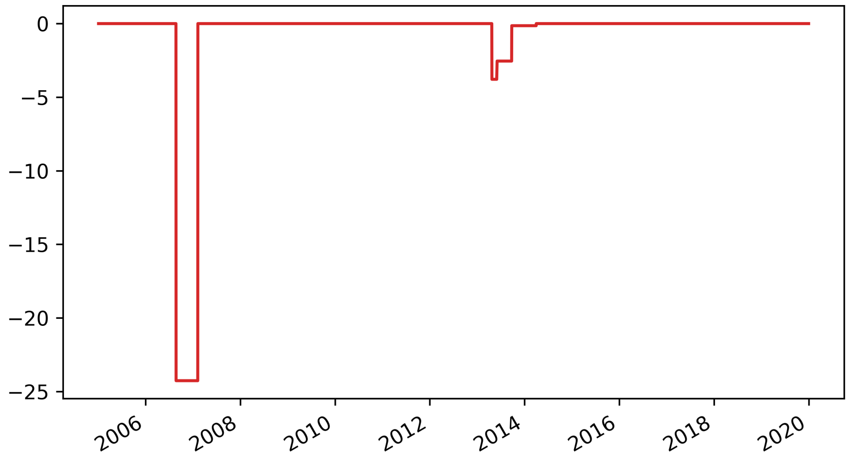

Figure 4.5: Draw Down of Stock Pair (GOOGL, INTC)

4.3 Stock Pair (ADBE, MSFT)

Figure 4.6 shows that the cumulative returns of the stock pair were higher than that of the risk-free rate, but the level of returns dropped after 2012. Figure 4.7 shows that the draw down near 2014 was gradually stopped right after the occurrence.

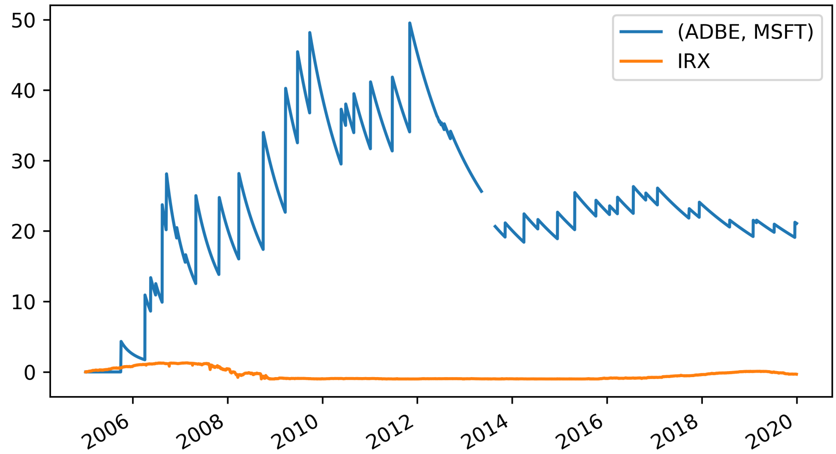

Figure 4.6: Cumulative Returns of Stock Pair (ADBE, MSFT)

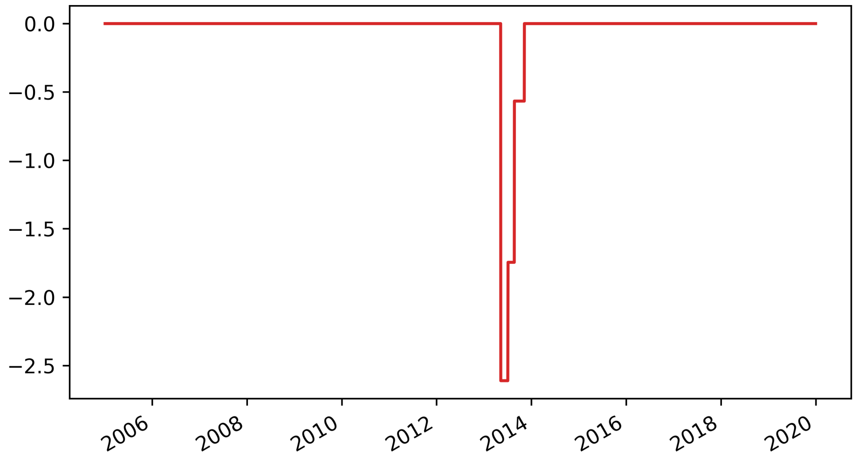

Figure 4.7: Draw Down of Stock Pair (ADBE, MSFT)

Chapter 5

Discussions

5.1 Limitations

Despite of some favorable feedback from the testing result, there are still some constraints in this study.

Firstly, the cumulative returns were not fully optimized. From the testing results, even though the returns were much higher than investing on the risk-free rate, the peak returns may not yet be reached at the end of the trading period, which means that the strategy may lead to some unnecessary losses at the later stage. This may be attributed to the limitation of trading interval. Since stock prices could react quickly on sudden changes, even in an hourly manner, trading on a daily basis may not capture the market information effectively.

Secondly, there may be overfitting or underfitting. During the grid search in optimization, the step sizes and ranges for different parameters were predefined, which could be biased. Hence, the optimization results may not disclose the actual optimized parameters. Besides, it is not guaranteed that the optimized parameters at training stage are the same as those at testing stage.

Thirdly, this study may not be applicable to the entire stock market since the targeted stock source was based on the technology sector. Stocks from other industries may not follow the same trend. Also, there could be stock changes specified to certain sectors, which may not be discovered in this study.

5.2 Future Development

The following suggests some aspects that the strategy can be further improved.

1. **Clustering**
   Stocks in different sectors could also be cointegrated. Clustering algorithms may be applied to determine the set of interested stocks, so that there is a larger sample size for optimization and backtesting.

2. **Normalization**
   Apart from somple moving averages, other calculations such as exponential moving averages could serve as alternatives in normalization, which may yield a better trading profitability.

3. **Optimization**
   The grid search could be intensified by narrowing the step sizes. There could also be weightings applied on different data, so that the recent one deserves more attention. Moreover, advanced optimization methods could be utilized. In fact, some of these approaches have been explored, but an explicit function has to be supplied, which is yet to be figured out.

4. **Trading**
   The trading simulation could be made more robust so that it is capable of handling more conditional trades. Threading could also be done to boost the program efficiency.

Chapter 6
Conclusion

This report has presented the methodology of using a pairs trading strategy in technology-related stocks. Although the findings have shown its profitability, the strategy can be further fine-tuned by introducing more sophisticated techniques.

Through this project, I have broaden my understanding in trading, especially when I do not come from any related backgrounds. Moreover, I was given this opportunity to apply some computing techniques in solving a real-life problem on my own, which was an invaluable experience.

Appendix A

Cointegration Testing Result

	Stock Pairs	p-values
1	(AMZN, INTC)	0.0150584
2	(INTC, MSFT)	0.0152943
3	(GOOGL, INTC)	0.0170111
4	(ADBE, MSFT)	0.0358782
5	(AAPL, INTC)	0.0501868
6	(ADBE, INTC)	0.0553548
7	(CSCO, INTC)	0.113077
8	(CSCO, MSFT)	0.166693
9	(ADBE, CSCO)	0.174653
10	(CSCO, GOOGL)	0.175101
11	(AMZN, CSCO)	0.193142
12	(INTC, NVDA)	0.197746
13	(CSCO, NVDA)	0.221424
14	(AMZN, NVDA)	0.267978
15	(AMZN, GOOGL)	0.275425
16	(AAPL, GOOGL)	0.321601
17	(GOOGL, NVDA)	0.385409
18	(ADBE, NVDA)	0.401503
19	(AAPL, MSFT)	0.488445
20	(AAPL, ADBE)	0.571908
21	(AAPL, AMZN)	0.666002
22	(ADBE, GOOGL)	0.770878
23	(GOOGL, MSFT)	0.840979
24	(ADBE, AMZN)	0.846174
25	(MSFT, NVDA)	0.853373
26	(AAPL, CSCO)	0.864897
27	(AAPL, NVDA)	0.911023
28	(AMZN, MSFT)	0.92193

Appendix B

Optimization Result

(Only the top 50 results are shown as follows.)

	t_{short}	t_{long}	z_{low}	z_{high}	**CAGR (%)**	**MDD (%)**	**SR**
1	2.0	150.0	0.3	1.0	26.924478	-33.994605	14.196635
2	4.0	120.0	0.2	1.0	25.421691	-32.535939	13.370391
3	4.0	105.0	0.2	1.0	24.327562	-2289.472342	27.362748
4	2.0	105.0	0.3	1.0	24.022689	-40.879782	24.027737
5	4.0	135.0	0.2	1.0	23.553891	-36.517522	12.748263
6	5.0	105.0	0.3	1.0	22.949733	-2495.431176	22.476588
7	2.0	90.0	0.2	1.0	22.621927	-69.323238	22.952336
8	2.0	150.0	0.3	1.1	22.537465	-32.422860	13.088574
9	2.0	135.0	0.1	1.0	22.536004	-35.378842	11.878492
10	4.0	135.0	0.3	1.0	22.491172	-33.557224	12.703149
11	2.0	120.0	0.3	1.0	22.424882	-696.334106	21.234039
12	4.0	165.0	0.2	1.0	22.041784	-32.559491	24.868090
13	3.0	150.0	0.3	1.0	21.819821	-38.235692	12.179159
14	2.0	135.0	0.3	1.0	21.602899	-1127.160005	17.385588
15	3.0	90.0	0.3	1.0	21.412859	-53.020578	21.596079
16	2.0	135.0	0.2	1.0	21.380243	-1022.580538	17.385270
17	5.0	120.0	0.3	1.0	21.368677	-31.798950	11.419567
18	5.0	105.0	0.2	1.0	21.183019	-31.103059	12.312617
19	2.0	105.0	0.2	1.0	20.883007	-1236.638299	22.416997
20	1.0	150.0	0.2	1.0	20.858276	-35.681500	11.692496
21	4.0	135.0	0.1	1.0	20.844181	-39.852407	11.213253
22	3.0	180.0	0.3	1.0	20.224668	-32.049719	22.590737
23	1.0	165.0	0.3	1.0	20.161505	-31.568830	11.815353
24	2.0	150.0	0.2	1.0	20.156625	-37.145079	11.441382
25	4.0	165.0	0.3	1.0	20.043387	-43.233331	19.810299

	t_{short}	t_{long}	z_{low}	z_{high}	**CAGR (%)**	**MDD (%)**	**SR**
26	4.0	105.0	0.3	1.0	20.009702	-1854.797035	19.780071
27	3.0	120.0	0.2	1.0	19.894056	-2475.038232	16.923787
28	4.0	135.0	0.3	1.1	19.876498	-32.466427	12.065250
29	3.0	150.0	0.2	1.0	19.826320	-26.079172	11.444591
30	2.0	90.0	0.2	1.1	19.815351	-49.029412	21.408822
31	4.0	120.0	0.2	1.1	19.802211	-31.197130	11.574014
32	5.0	150.0	0.2	1.0	19.799623	-43.529068	19.354006
33	2.0	90.0	0.3	1.0	19.786795	-66.015552	19.983809
34	2.0	165.0	0.3	1.0	19.721520	-25.354600	11.564956
35	4.0	105.0	0.2	1.1	19.705380	-1349.435812	24.715363
36	4.0	165.0	0.2	1.1	19.688487	-32.984908	22.773516
37	3.0	135.0	0.3	1.0	19.477831	-2272.563171	16.143974
38	3.0	165.0	0.2	1.0	19.379142	-32.452233	19.165895
39	1.0	105.0	0.2	1.0	19.298457	-53.587329	22.490190
40	1.0	150.0	0.3	1.0	19.244618	-610.765369	16.415878
41	5.0	165.0	0.3	1.0	19.228636	-31.872792	22.106091
42	5.0	135.0	0.3	1.0	19.174477	-43.579836	16.901537
43	4.0	135.0	0.2	1.1	19.149895	-35.200205	11.171598
44	5.0	135.0	0.2	1.0	19.049495	-35.835240	10.406194
45	2.0	105.0	0.3	1.1	18.920283	-35.019312	21.087833
46	1.0	120.0	0.3	1.0	18.914660	-2071.573754	19.586833
47	2.0	90.0	0.1	1.0	18.905659	-131.320805	22.157410
48	5.0	105.0	0.3	1.1	18.871067	-1399.674687	20.290745
49	4.0	180.0	0.2	1.0	18.702939	-26.403160	22.488163
50	4.0	150.0	0.3	1.0	18.664984	-43.817189	15.107539