International Journal of Economics and Financial Issues

ISSN: 2146-4138

available at http: www.econjournals.com

International Journal of Economics and Financial Issues, 2018, 8(4), 78-88.

High-frequency Pairs Trading on a Small Stock Exchange[#]

Andreas Mikkelsen[1]*, Frode Kjærland[2]

[1]Nord University Business School, 8049 Bodø, Norway, [2]NTNU Business School, Norwegian University of Science and Techcology, 7491 Trondheim, Norway. *Email: andreas.mikkelsen@nord.no

ABSTRACT

We study the performance of a high-frequency pairs trading (PT) strategy on the 100 most liquid stocks, in 15-min intervals, on a small commodity dominated stock exchange (Oslo Stock Exchange) using a comprehensive dataset from January 2012 to March 2016. We use both the distance and cointegration approach. Moreover, we let the formation (trading) period vary between 2 (1), 4 (2), and 6 (3) weeks, in order to test the impact on the PT profit. We find that the distance and cointegration approaches both have their strengths and weaknesses. In addition, we find that a shorter formation (and trading) period yields better results. As a further contribution to the literature, our findings imply that a simple static PT strategy, still is profitable using high-frequency data. Further, our results show better performance in a bull market than in a sideways-moving, volatile market.

Keywords: Finance, Pairs Trading, Statistical Arbitrage, High-Frequency Trading, Distance Method, Cointegration
JEL Classifications: G11, G15

[#] We would like to thank Francisco Santos for helpful comments

1. INTRODUCTION

We study high-frequency pairs trading (HFPT) on a small, and relatively illiquid stock exchange to examine whether the reported profits from larger markets hold in this type of market as well.

The empirical setting is the Oslo stock exchange (OSE), representing a small commodity driven stock exchange. This is complementary to the majority of pairs trading (PT) studies performed on S&P500 data or in other, larger stock exchange contexts. In addition, this study adds to the focus on HFPT linked to oil companies (Liu et al., 2017) as well as other recent HFPT studies (Fallahpour et al., 2016; Stübinger and Bredthauer, 2017). We use a comprehensive dataset from the OSE with data binned in 15-min intervals from January 2012 to March 2016. These data provide the opportunity to study HFPT performance both in a bull market, from January 2012 to June 2014, and in a more volatile, stable-trend situation, from July 2014 to March 2016.

The choice of formation period is crucial in any PT implementation because of the sensitivity of returns when formation periods are changed (Huck, 2013). Previous research is inconclusive in this regard. Hence, we compute three formation periods (2, 4, and 6 weeks) followed by a trading period that is half the size of the formation period, to unfold more information regarding this important feature of PT performance. Moreover, we perform a comparison of the distance and cointegration approaches on high frequency data. The body of literature on HFPT is scarce, making any thorough study a contribution to the knowledge on this subject.

We contribute in the following ways. We add to a small body of literature on HFPT, including that on the impact of transaction costs. We compare the two most conventional approaches in PT, the distance and cointegration approaches, in the arena of HFPT. Finally, we compare the impact of varying the formation and trading period in a HFPT setting.

The remainder of the paper is structured as follows. Section 2 introduces the PT strategy with a special focus on the formation and trading periods as well as the idiosyncrasy of high-frequency trading (HFT), and reviews the relevant literature. Section 3 presents the data and methodology, while Section 4 provides the results. Finally, Section 5 concludes and we discuss some implications of the findings.

2. LITERATURE SECTION

HFT has evolved as a part of financial markets during the last decade. Some pundits refer to it as merely faster trading, while others claim that it represents a new paradigm. Moreover, as a relatively new phenomenon, much of the discussion on HFT is "not backed by solid academic research" (Chordia et al., 2013. p. 637).

Within HFT, Hagströmer and Norden (2013), when studying Swedish data, subcategorize HFTs into market-making and opportunistic strategies. Market makers are defined as "proprietary-only firms that use algorithms in their order submission" (p. 742), while opportunistic refers to a more classical arbitrage strategy. They find more players of the first category than of the latter. HFPT obviously represents an opportunistic approach, being a sophisticated development from the PT strategy invented on Wall Street in the 1980s.

The PT concept is intuitively easy to follow: Identify two traded securities that co-move over a certain time period. If they diverge from an equilibrium, buy the security that is relatively underpriced (long) and sell that which is relatively overpriced (short). The concept is, then, to speculate that the securities will reverse to their equilibrium, at which the trades are reversed. The trader gains a profit from the transactions by configuration. This implies that the spread of the securities is bought and sold. PT is referred to as a market-neutral strategy, since one gains profits regardless of whether the market goes up or down (for a more comprehensive review, Krauss, 2017).

The most known and common method is the distance approach, proposed by Gatev et al. (2006), which dominates the empirical work on PT (Krauss, 2017). This approach involves normalizing the price series to start at 1. Then, one subtracts one series from the other, and the goal is to find pairs with the smallest possible sum of squared deviations, meaning pairs with a small distance between the normalized price series.

Cointegration is also an approach that has been frequently used in previous studies. This technique allows us to find two (or more) time series with common factors that drive their evolution. The common factors ensure that a linear combination of the two with a stationary long-term equilibrium and finite variance can be found. Vidyamurthy (2004) provides the most cited work regarding this approach. Further, cointegration represents a more formal procedure in comparison to the distance approach.

Common to any approach is that the data sample is "split" into two parts. The first part consists of a formation, or training, period, in which one evaluates the pairs that show the best characteristics for trading. In the second part, which does not overlap with the first, the pairs are actually traded. Once a selection approach is chosen and pairs are found to be promising for trading, a trading algorithm is operationalized. However, certain choices need to be made. In the literature, a trigger based on an SD metric is commonly used to decide when to enter a trade. The threshold is not obvious, although one has to bear in mind that a smaller threshold usually results in a more sensitive trigger and, thus, more trades, which increases transaction costs and can eliminate possible profits earned from the strategy. A high trigger results in fewer trades; yet, if too strictly defined, hardly any trading occurs and no profits are generated. When the spread crosses back over the historical equilibrium, the trades are reversed to exit the longshort positions. Regarding the profitability of PT, relevant studies find ambiguous results Gatev et al. (2006); Do and Faff (2010; 2012). This subject has also been debated by Jacobs and Weber (2015); Engelberg et al. (2009), Krauss (2017). Several studies find that profits has decreased in recent years. Moreover, most studies on PT are on U.S. data or from other large stock exchanges[1]. Studies in a Nordic context is done by Broussard and Vaihekoski (2012) and Mikkelsen (2018).

In addition to equity markets, PT has been studied on the US treasury securities market (Nath, 2003), bitcoin exchanges (Lintilhac and Tourin, 2017) and in commodity markets (e.g. Dunis et al., 2015).

Several studies point out the significance of the formation period as a key element for excess returns (Bowen et al., 2010). (Huck 2013) chooses four different formation periods (6 months, 1 year, 1½ years, and 2 years) and finds that excess returns are highly sensitive to this input parameter. Although the results are somewhat mixed, he finds that when an 18- or 24-month formation period is applied, the excess returns are 0.60% higher per month and are described as "likely to be economically significant" (p. 1302).

Since several papers report of declining PT profits in recent years, recently a number of papers are investigating new methods for identifying tradable securities. Examples are copula (Liew and Wu, 2013, Krauss and Stubinger, 2017, Stübinger x 2, et al., 2016), machine learning (Krauss et al., 2017), time series with regime switching (Yang et al., 2016) and partial cointegration (Clegg and Krauss, 2018) approaches.

2.1. HFPT

The vast majority of PT studies are on daily data (Huck, 2013; Rad et al., 2016). Concerning HFPT, there are far less studies, despite it being more closely related to real world trading. However, the access to and processing of huge amounts of data have, so far, limited the number of studies in the literature. Consequently, there is a need for these types of realistic approaches in order to better understand the performance of these types of strategies in different markets.

With regard to our HFPT approach, there are a few relevant studies. Nath (2003) studies PT in U.S. treasury securities (from 1994 to 2000) using the distance approach, outperforming the relevant

1 Perlin (2009) uses Brazilian data, Bolgün et al. (2010) Turkish data, Bogomolov (2011) Australian data, Mashele et al. (2013) South African data, and Li et al. (2014) Chinese and Hong Kong data. However, all these studies are on daily data.

benchmarks. However, an examination of exposure to systematic risk factors was not performed.

Bowen et al. (2010) use the distance approach on FTSE100 equities between January 2007 and December 2009 using data in 60-min intervals. The formation period is 264 h and the following trading period is 132 h. They find that the strategys returns are sensitive to both transaction costs and the speed of execution. When waiting one period for the trade and adding transaction costs of 15 basis points, profits are eliminated.

Dunis et al. (2010) study the Eurostoxx50 index between July 3, 2009 and November 17, 2009 (5½ months). They use 5-, 10-, 20-, 30-, and 60-min data, in addition to daily data. They also use a cointegration approach, in addition to employing the Kalman filter for time-varying coeffcient estimation. Moreover, they seem to use a 1-week period for both formation and trading periods. On average, they find that the results are not attractive; however, when using the top five pairs with the most attractive in-sample indicators, they obtain positive results.

Kim (2011) studies the equities listed on the KOSPI100 index (Korea), using the cointegration approach with the Kalman filter used to estimate the time-varying coeffcients. The data are in 30-min intervals, with a 2-week formation period and 1-week trading period. He finds positive excess returns, after transaction costs are taken into account, when trading takes place one period after the trading signal. Such a strategy performs better during bear markets, that is, in this case, during the financial crisis.

Further, Fallahpour et al. (2016) provide another HPFT study based on S&P500 data from June 2015 to January 2016. Yet, their scope is somewhat different, focusing on how reinforcement learning outperforms other methods in obtaining the best parameters for performing a cointegration HFPT strategy.

Liu et al. (2017), studies the oil companies listed on the NYSE. Their approach concerns a novel way of modeling spreads between pairs of stocks in order to search for temporary market mispricing ineffciencies in a more dynamic way, when compared to the distance and cointegration approaches. The data are on 5-min intervals in 2008 and June 2013 to April 2015. They find clearly positive results in back-testing the strategy, also in 2008. Hence, they confirm the results of other studies finding PT especially profitable in bear markets (like Gatev et al., 2006; Kim, 2011; Do and Faff, 2010; Rad et al., 2016). The results are of interest to our study on the data from a commodity-driven (especially oil and gas) stock exchange.

Miao (2014) employs data of 177 oil and gas stocks from the US market, in a study from May 2012 onwards. The data is in 15-min intervals and he uses a two-stage correlation and cointegration approach for choosing pairs suitable for trading. He find that the strategy performs well, with a cumulative return of 56.58% over a 12 month period.

Mikkelsen (2018) studies PT on a small sample of seafood companies listed on OSE. The study is comparing the use of high frequency data and daily data for both the distance and cointegration approach, where he find that none of them give significant profits. The comparison of daily and high frequency data is nevertheless not quite apples to apples, since they have the same length in formation and trading periods.

Stübinger and Bredthauer (2017) provides a thorough examination of high HFPT on the S&P500 constituents from 1998 to 2015 using transaction data binned to 1-min intervals. They use three related approaches for pairs selection, the distance approach, correlation and fluctuation behaviour, in the spirit of Liu et al. (2017). In the trading period, they use three different approaches. A static threshold approach, which is the standard in the literature, a benchmark. Dynamic thresholds, using running mean and SD, following Bollinger (1992). Finally, reverting thresholds, which is similar to the dynamic threshold approach. Albeit, that a position is opened on the second crossing of the threshold, instead of the first, to ensure mean-reversion in the spread. Thus, they have in total nine strategies, with a formation period of 2 weeks, which they trade for 1 week. They find that using the distance approach with a dynamic threshold for trading, yields the best performance, with a return of 50% p.a. The static method, which is closely related to our study, also performs well at 21.5% p.a. Further, they confirm the findings that PT returns are decreasing over time, thus more advanced approaches should be considered.

In this spirit, Stübinger and Endres (2018) models the spread of pairs as a mean-reverting jump diffusion model. Their empirical setting is oil and gas companies on the S&P500 from January 1998 to December 2015, using 1-min data intervals. Their model far out-performs the classic approaches with returns of 60% p.a, and Sharpe Ratio of 5.3 after transactions costs.

The abovementioned studies represent the few that have been conducted on HFPT. Moreover, these studies focus on different aspects, making them more or less relevant to our study. Clearly, there is a need for more studies unfolding the impact that HFT/HFPT features have on returns. Thus, our study provides an extension within a small body of literature on this topic.

More specifically, we intend to extend the literature on HFPT, to test whether the findings from a highly liquid market (S&P500), translates to a more illiquid smaller stock exchange in a high frequency setting. Further, we are testing a basic model with static thresholds, which one would expect to be the least profitable. Thus, if we can show that the basic static model is profitable, practioners, which is usually considered smart money, should be able to come up with more advanced trading models that are more profitable.

3. DATA AND METHODOLOGY

The dataset consists of the high-frequency tick-by-tick transaction data of the OBX and OB Match segments of the OSE; these segments consist of roughly the 100 most liquid stocks at the exchange at any time. The data span the period from January 2012 to March 2016, and the entire sample consists of 226 unique stocks. As noted by Liu et al. (2017), using high-frequency data, one might get the impression that the sample period is rather short in comparison to studies utilizing daily data. In our case, the stocks have 30 daily ticks over 4 years and a quarter, which translates to 30 × 252 × 4.25 ÷ 252 = 127.5 years of daily data.

In the literature, it is well known that high-frequency data usually need to be filtered, as the raw data consist of outliers and errors, not representing the usual market conditions, nor being suitable for statistical analysis. In order to account for these errors, we first transform the tick-by-tick data to 5-s intervals and then remove all trades that have not occurred between 09:01 and 16:20, since these trades would not have occurred in the continuous trading session at the exchange. To remove outliers and errors, we follow the procedure suggested by Verousis and ap Gwilym (2010), shown in Figure 1. Having removed the outliers, we transform the data to 15-min intervals in open-high-low-close format; if a price is missing in an interval, the previous price is used. To account for dividend payments and corporate events, we calculate the adjusted prices using adjustment factors from the TITLON (nd) database. Finally, we wish to remove the least liquid stocks from our sample, and we do so in a backward-looking fashion in 2-week windows.

We test in the window at time t-1 and remove stocks in the window at time t; this can be compared to a training and trading period. First, we remove all stocks that have unadjusted prices less than NOK 10 on the last tick of the training period, to remove the effects from small-cap stocks. Second, for the remaining stocks, we calculate the average share volume per tick during the training period window, and remove those in the lowest volume quartile. The effect of the sorting on market cap and volume is that, on average, we have 60 stocks in our sample, of which we are forming pairs.

The distance approach of Gatev et al. (2006) has become the standard procedure for pair selection in the PT literature. We normalize prices to start at NOK 1 in the beginning of every formation period. The aim is to find pairs with the smallest average squared spread between two return series, as shown in Equation 1. We have a requirement that there must exist observations for at least 50% of the possible observations in a formation period window. We may obtain missing values if two time series are not completely overlapping, for instance, if a stock is listed in the middle of the training period.

$$\mathrm{SSD} = \frac{1}{n}\sum_{t=1}^{n}\left(P_t^1 - P_t^2\right)^2 \qquad (1)$$

We want the best possible pairs for our trading strategy; therefore, the top 1% pairs in every training period are selected. This is similar to choosing approximately the top 20 pairs in every period, which is the standard in the literature. Further, this choice makes the number of pairs approximately the same using both approaches, thus making them comparable.

We also employ a cointegration approach. The basic intuition is that if we have two price series that are integrated of order one – I(1), there might exist a linear combination of the two series that is stationary; if so, the two series are said to be cointegrated. This is shown in Equation 2. We are using the Johansen (1988) procedure, and the test is the max eigenvalue test, without a trend or constant specification. The number of lags is determined by the AIC criterion. All pairs that are found to be cointegrated at a 5% significance level are chosen for trading. Since the cointegration vector β determines the cointegration relationship, as well as the long-run equilibrium, we also use this for order sizing in the trading periods. The β is normalized with respect to stock A, so we enter positions with 1 NOK in stock A and β NOK in stock B.

$$\mu = P_t^1 - \beta P_t^2 + \epsilon_t \qquad (2)$$

The strategies differ in the length of the formation period. In our study, therefore, we have six unique strategies. The strategies differ in the length of the formation (trading) periods, with 2 (1), 4 (2), and 6 (3) weeks. In the following, we denote the strategies D2–D6 (e.g. distance method, 2 week formation period) and C2–C6 for the distance and cointegration approaches respectively. Since the first 2 weeks of the sample are used for sorting on liquid stocks, the first possible trading period is in the start of February 2012.

Figure 1: Algorithm for removing outliers. The chart is found in Verousis and ap Gwilym (2010)

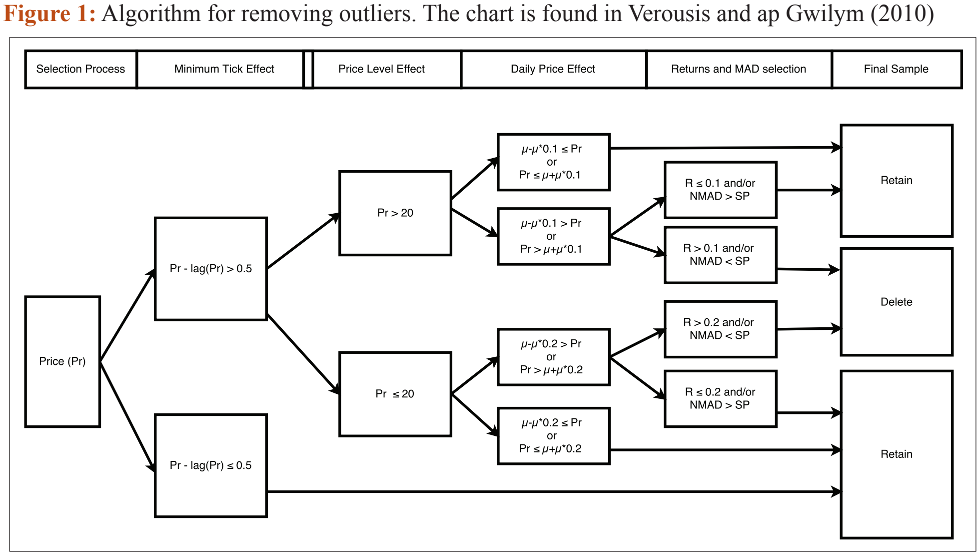

We roll the training period forward by the length of the trading period; thus, we have new pairs for trade when the previous trading period is finished. This also ensures that we do not have overlapping trading periods, which, according to Broussard and Vaihekoski (2012), is unreasonable for an illiquid market. Further, for the trading signal, we are using the adjusted close price, while the trade is occurring at the next timestamp on the open price. A position is opened if the spread has diverged more than two SD, both for the distance and cointegration periods, as measured in the formation period. We close the position, or reverse the trades, if the indicator crosses the mean (equilibrium) from either side. If positions do not converge, they will be closed at the end of the period; the same applies if a stock gets delisted. We have not employed any stop-loss mechanisms; although, in the case of the shortest trading period, one might argue that this is implicitly in place because of the short trading period.

For every position, transaction costs are included — 0.0195% for every trade in the position — which equals 7.8 basis points for the entire position. The rate is chosen based on the rates offered by internet brokers, where the most attractive rate is 2.9 bp per transaction for private investors. It is natural to believe that an institutional investor would have lower transaction costs than this. In addition, it is well known in the literature that HFTs often take the role as liquidity providers, meaning they can get rebates from the exchange for this service. Hagströmer and Norden (2013) find that opportunistic HFT are on the supply side of the transaction in over 30% of their trades. We have not accounted for price impacts of trading or trading slippage, which would increase the trading costs for an investor. However, these impacts are hard to estimate (Do and Faff, 2012). As a comparison to other studies, Stübinger and Bredthauer (2017) use transaction costs of 10 bp per roundtrip in their study.

We calculate returns on a daily basis. For the individual pair, we measure the return by the change in net equity (profit or loss), labeled ΔPL on both the long and short positions from time t-1 to time t. The change is divided by the initial exposure on each position, in addition to transaction costs, as shown in Equation 3. If a pair has multiple round trips in a trading period, the cash sitting in the account does not accrue interest. The average daily return is the mean return on all pairs available at time t, as shown in Equation 4. Finally, we compound the returns to weekly, as shown in Equation 5. Since the returns stem from a long and a short position, it is natural to view them as excess returns (Gatev et al., 2006). Thus, we will address the returns from PT strategies as excess returns.

$$R_{pair,t} = \frac{PL_{long}}{InitialExposureLong + \frac{Txn}{2}} + \frac{PL_{short}}{InitialExposureShort + \frac{Txn}{2}} \tag{3}$$

$$R_t = \frac{\sum_{pair=1}^{n} R_{pair,t}}{n_t} \tag{4}$$

$$R_w = \prod_{1}^{n} \left(1 + R_t\right) - 1 \tag{5}$$

4. RESULTS

The results regarding the cumulative excess returns for all six strategies are presented in Figure 2. The figure shows that if NOK 1 was invested at the start, the D2 strategy would yield approximately NOK 7 at the end (600% increase!), 4 years and 3 months later. Moreover, we see that five out of the six strategies yield (highly) positive returns. The three strategies using the distance approach all end up with nice profits. Two out of the three strategies using cointegration also provide positive excess returns. It is only the C6 strategy that is unprofitable in the studied period. These findings provide support to HFPT as a feasible approach to excess returns. These findings are in contradiction with those in Huck and Afawubo (2015) and, more interesting, with those in Bowen et al. (2010). Bowen et al. (2010) may be the closest study

Figure 2: Performance of the six HFPT strategies, cumulative returns January 2012 to March 2016

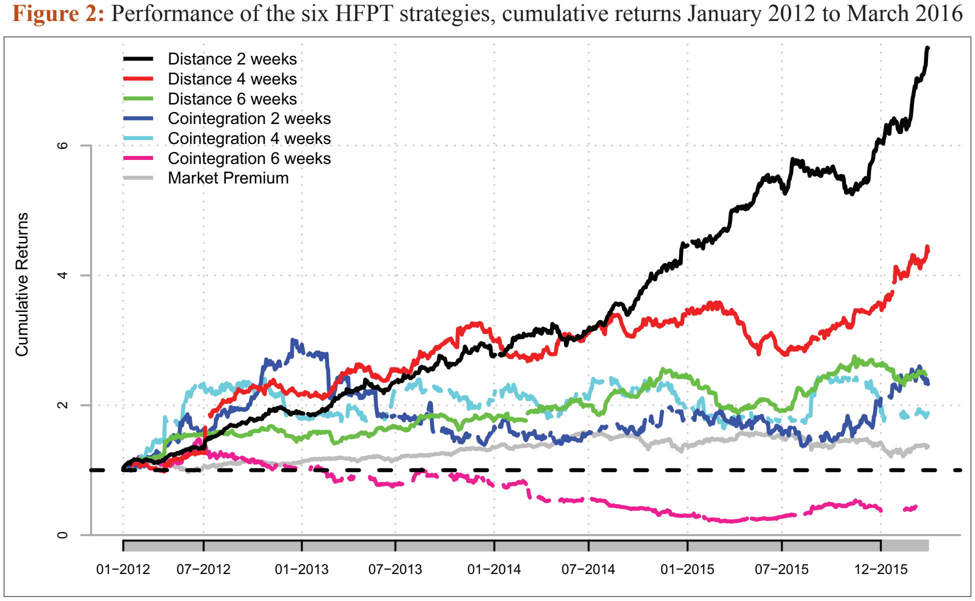

to ours geographically, given that they use European data. They do not find an HFPT approach to be profitable when adjusting for transaction costs, when studying the FTSE100 in 2007– 2009 - a period that should be especially suitable for PT (since the literature often refers to bear markets as especially suitable for a PT strategy). On the other hand, our results are similar to those of Stübinger and Bredthauer (2017), which has use a large dataset on the S&P500.

We find that the distance approach yields higher cumulative returns than the cointegration approach. All three distance approaches perform better than the best (C2) strategy using cointegration. Moreover, the strategies with shorter formation (trading) periods are more profitable than those with longer formation (trading) periods. This is consistent for both approaches, although the 4-week cointegration approach performs better than the 2-week strategy for a lot of the period. More comprehensive information regarding the distribution of excess returns is presented in Table 1. These descriptive statistics provide more insight into the profitability of the strategies. From Table 1, we observe that annualized excess returns provide fine results for the first five strategies (all five provide more than 13% annualized excess returns). To formally test whether the trading strategy has produced significant profits, we perform a Newey-West t-test using four lags on the weekly return series. The estimate for the mean weekly excess return for the C2 strategy is 0.29% (t-statistic: 3.96), when including all weeks in the sample, even those without any trading. This is significantly different from zero, implying that the profit is statistically significant, a weekly return of 0.29% must also be said to be economically significant. From the table we see that the D4 and C2 strategies have significant returns as well, where the latter has the strongest performance of the two. We also observe that the shorter strategies (2- and 4-week) clearly outperform the 6-week approach, since neither the D6, nor the C6 strategy has significant returns.

The distance approach is also less risky as measured by the SD, than the cointegration approach, when comparing strategies of similar formation periods. Moreover, we see positive skewness in all cases, a good characteristic for any trading strategy. Although, this is less pronounced for cointegration. The distribution of all strategies are leptokurtic, more so for the distance strategies. The annualized Sharpe ratios provide some information concerning the risk adjusted excess returns. We see, again, the best performance for the two shortest (2-week) strategies, and we do note the low risk associated with the D2 strategy. The market annualized Sharpe ratio for the same period is 0.245. Hence, the first five strategies by far outperform the market using this metric. These observations contribute to the understanding of why we observe positive overall excess returns. We further test whether the D2 outperforms the C2 stategy on a risk adjusted basis using the probabilistic Sharpe ratio measure proposed by Bailey and Lopez de Prado (2012), shown in Equation 6[2]. This measure also accounts for skewness and kurtosis in the return distribution, which is suitable for our purpose with non-normal returns and the track record of the strategy. The test value is 0.77, thus well below the critical value of 0.95. Hence, we cannot say whether the performance of the D2 approach is significantly different from the C2 approach.

$$\widehat{PSR}\left(SR^*\right) = Z\frac{(\widehat{SR} - SR^*)\sqrt{n-1}}{\sqrt{1 - \widehat{\gamma_3}\widehat{SR} + \frac{\widehat{\gamma_4} - 1}{4}\widehat{SR}^2}} \tag{6}$$

The study by Stübinger and Bredthauer (2017) report daily returns of 0.08% for their D2 strategy, where the corresponding result in our study is 0.2%, thus a difference of 0.12% favouring our study. This is not an unreasonable result, since we would expect the liquid market to be more effcient than the illiquid market. In addition, they archieve twice the return with a more dynamic approach. Some words of caution though, as the numbers are however not quite comparable; 1) they use a slightly higher transaction cost than we do and 2) one would expect a higher degree of trading slippage and price impact on the OSE compared to the stocks listed in the S&P500, which would impact the

2 $\widehat{SR}$ is the D2 Sharpe ratio, SR* is the C2 Sharpe ratio,n denotes the number of weekly returns, γ_3 and γ_4 denotes skewness and kurtosis of the D2 empirical distribution respectively and Z is the cdf of the standard normal distribution.

Table 1: Weekly return distribution, N is the number of weeks where trading is occuring for each strategy

	Distance 2 weeks	Distance 4 weeks	Distance 6 weeks	Coint. 2 weeks	Coint. 4 weeks	Coint. 6 weeks
N	218	216	214	218	216	214
Average excess return	0.0029	0.00356	0.00138	0.00502	0.00288	−0.00119
Standard error (Newey-West)	0.00073	0.00205	0.00097	0.00159	0.00187	0.00144
t-statistic	3.959***	1.735*	1.426	3.15***	1.543	−0.827
Excess return distribution						
Median	0.0014	−0.0001	0.0012	0.0029	0.0007	−0.002
Standard deviation	0.0114	0.0268	0.0155	0.0235	0.0272	0.025
Skewness	2.189	8.064	5.81	1.104	3.406	0.254
Excess Kurtosis	11.87	85.38	61.28	4.65	27.28	2.61
Minimum	−0.023	−0.039	−0.036	−0.066	−0.058	−0.084
Maximum	0.08	0.31	0.17	0.12	0.24	0.11
Obs. with excess return <0	0.44	0.5	0.44	0.43	0.48	0.56
Annualized return	0.145	0.178	0.069	0.251	0.144	−0.06
Annualized SD	0.081	0.19	0.11	0.166	0.192	0.177
Annualized Sharpe ratio	1.8	0.94	0.63	1.51	0.75	−0.34

Whether the average excess return is di erent from zero, is tested using Newey-West standard errors with 4 lags on the weekly return series, the corresponding t-statistic is shown in the next row, with significance stars. The annualized excess returns, SD and is acquired by scaling the corresponding weekly estimates. The annualized Sharpe ratio is calculated from the annualized excess return and SD. Significance codes: ***0.01, **0.05, *0.1

Table 2: Summary statistics

	Distance 2 weeks	Distance 4 weeks	Distance 6 weeks	Coint. 2 weeks	Coint. 4 weeks	Coint. 6 weeks
a. Trading statistics. The panel shows descriptive statistics for the different strategies. A roundtrip is defined as the opening and closing of a position, and corresponds to 4 trades. The holding period is measured in days. Whether a position is win, loss or neutral is measured before transactions costs.						
Avg. pairs selected per period	18.17	5.03	5.36	14.02	5.27	3.66
Avg. pairs trading	15.49	4.41	4.69	6.10	2.20	1.50
per period Avg. trades per	103.38	31.85	33.53	32.15	12.11	7.70
Period	25.84	7.96	8.38	8.04	3.03	1.93
Avg. roundtrips per Period	1.87	3.76	5.21	1.75	3.71	4.82
Holding period						
Biggest winner	0.53	0.47	0.49	0.34	0.51	0.29
Biggest loser	−0.19	−0.18	−0.40	−0.36	−0.22	−0.39
Win %	0.48	0.50	0.52	0.46	0.50	0.44
Neutral %	0.08	0.06	0.05	0.10	0.07	0.06
Loss %	0.44	0.44	0.43	0.44	0.43	0.50
W/L ratio	1.10	1.16	1.21	1.05	1.17	0.88
b. Pairs composition. The first two rows shows the percentage of convergence or non-convergence of the positions. The last three rows shows the percentage of pairs that have single or multiple convergence after opening a position, in addition to percentage of non-convergence pairs.						
% Position closed at roundtrip	52.18	55.29	54.18	49.94	48.93	48.30
% Position closed at end of period	48.38	45.00	45.93	50.51	51.38	51.70
% Non-convergence pairs	47.38	46.80	46.31	50.41	53.57	54.52
% Multiple roundtrip pairs	19.57	22.77	22.41	11.74	12.39	11.21
% Single roundtrip pairs	33.05	30.43	31.27	37.85	34.03	34.27

performance of the trading strategy. This may be reflected in the riskyness of the strategies as well, where the D2 strategy in the US market is 0.69% less risky as measured by daily SD, compared to our study.

The overall trading statistics are reported in Table 2. We observe that the distance approach is traded far more often than that of cointegration, while the holding periods are quite similar.

The average holding periods range from 1.75 days to 5.21 days, depending on the strategy. We see that the variation is not too sensitive when comparing the two methods and, logically, it increases with increased formation (trading) period. The holding period for the D2 strategy is similar to that of Stübinger and Bredthauer (2017), who reports a holding period of 1.62 days. Moreover, we observe that the distance approach is far more trading-intensive when compared to that of cointegration. Thus, a more formal approach, such as cointegration, is less sensitive to trading frictions and transaction costs, which might be of great importance in real world situations. One reason for this is that a much smaller percentage of the pairs selected for the cointegration approach is actually trading, compared to the distance approach. For the former, roughly 40–43% of the pairs are trading, while for the latter the number is much higher at roughly 85–88%. The win–loss ratios may not be overly impressive in support of an HFPT strategy, but nevertheless, they are (mainly) above 1 and comparable to studies like Nath (2003); Kim (2011) and Liu et al. (2017).

From the material, we do see quite high numbers for non-convergence, even in such short time periods as those studied. Evidently, this represents high risk when so many positions are not converging (Do and Faff, 2010). We also observe that the distance approach has less non-converging pairs than does that of cointegration. Moreover, we see that the distance approach has approximately twice as many multiple roundtrip pairs as does that of cointegration. Regarding the D2 strategy, we select on average ca. 18 pairs every trading period, thus it is comparable to the top 20 pairs from the US market. The non-convergence risk is much higher of the D2 strategy of Stübinger and Bredthauer (2017), compared to the Norwegian case, at 74% for the former and 47% for the latter. They are able to reduce the non-convergence risk by more than 50% by using more dynamic trading signals, which one can expect would translate to the Norwegian market as well.

In summary, although the Sharpe ratios between the D2 strategy and the C2 strategy are not statistically different, each the strategies has some distinct advantages. In the case of the D2 strategy, we observe a lower risk and a more positively skewed distribution. In addition, a larger degree of multiple roundtrips pairs and a slightly lower proportion of non-converging pairs. For the C2 strategy, we observe a higher average excess return, while the risk is also higher. Hence, this approach would benefit the most from stop-loss rules or dynamic trading signals. Further, the C2 strategy is a lot less trading intensive than the D2 strategy. The trading intensity is about a third of that of the D2 strategy, which is naturally beneficial with regards to trading costs.

To further investigate the risk factors that affect HFPT, we regress the weekly excess return series against the Fama and French (1993) three-factor model, and add the momentum factor proposed by Carhart (1997). We also include a liquidity factor described in Næs et al. (2009), since our data stem from a small, relatively less liquid stock exchange. All factors, including market return and the risk free rate, are retrieved from Ødegaard (nd), and the daily series is compounded to weekly series. The results are presented in Table 3.

First, we notice that the shorter strategies has significant alpha values, and this is more evident for the cointegration method, where only the shortest strategy has a significant alpha. Moreover, as previously mentioned, the PT strategy is market-neutral, and our regression confirms this, as shown by the numbers in Table 3.

One might expect that trading on a relatively small stock exchange would call for some liquidity premium. However, according to the last row in Table 3, we do not find strong evidence for a liquidity premium (except for the D4 strategy). This implies that there is suffcient liquidity in the stocks constituting our sample. For robustness tests, we also regress the contemporaneous PT excess returns on lagged market and liquidity factors (up to four lags, not reported), however, the lagged variables have no explanatory power.

The five factors explain only a small portion of the returns (R^2 values are below 3%), which is similar to other relevant studies. The model fails to explain the HFPT excess returns.

We also split the sample into two periods: January 2012 to June 2014 and July 2014 to March 2016. The reason for these actual

Table 3: Fama-French 3 Factor+Momentum+Liquidity. The asset pricing model is tested on the weekly strategy excess return series. The daily factors from Ødegaard (nd) are compounded to weekly factors. The t-statistics are HAC consistent, using the Newey-West method with 4 lags

	Dependent variable					
	Strategies					
	D2	D4	D6	C2	C4	C6
Alpha	0.003***	0.004*	0.001	0.005***	0.002	−0.002
	t=3.309	t=1.820	t=1.049	t=2.901	t=1.259	t=−1.279
MP	−0.039	−0.081	0.148	0.082	0.111	0.211
	t=−0.724	t=−1.508	t=1.517	t=0.799	t=0.834	t=1.531
SMB	−0.057	0.357	0.114	−0.051	0.249	0.313**
	t=−0.889	t=1.609	t=1.033	t=−0.361	t=1.528	t=2.031
HML	−0.018	−0.077	0.009	0.052	0.246**	−0.057
	t=−0.307	t=−0.777	t=0.169	t=0.457	t=2.078	t=−0.513
PR1YR	0.016	0.017	−0.093	−0.085	0.046	−0.008
	t=0.400	t=0.328	t=−1.303	t=−0.830	t=0.553	t=−0.097
LIQ	0.015	−0.374**	0.104	0.078	−0.120	0.149
	t=0.222	t=−2.027	t=0.936	t=0.703	t=−0.943	t=1.272
Observations	218	216	214	218	216	214
Adjusted R^2	−0.018	0.014	0.019	−0.012	−0.002	0.023
Residual Std. Error	0.011	0.027	0.015	0.024	0.027	0.025

Significance codes: ***0.01, **0.05, *0.1

Figure 3: Oslo all share index from January 2012 to April 2016

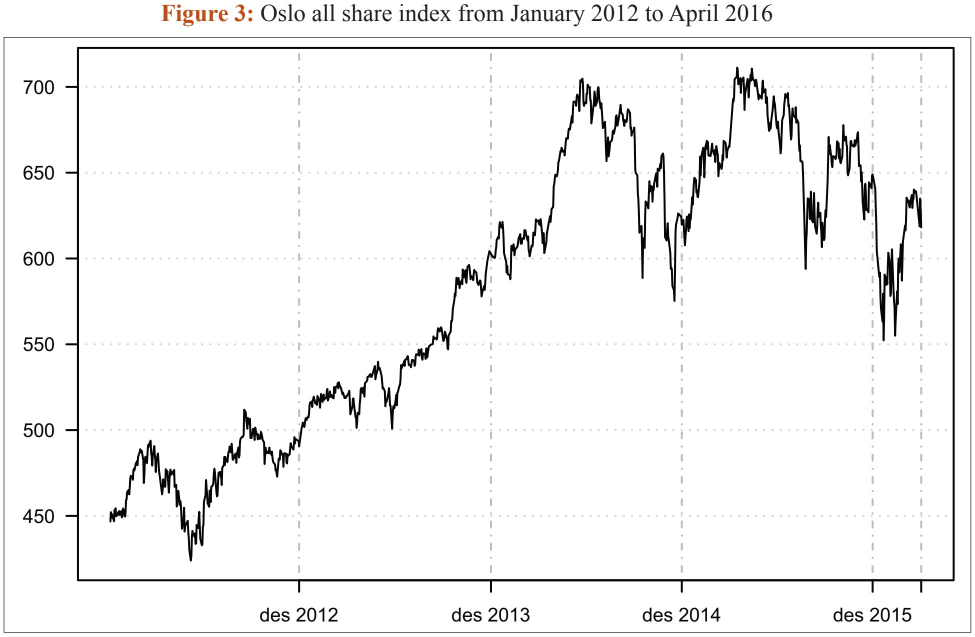

Table 4: Fama-French 3 Factor+Momentum+Liquidity from January 2012 through June 2014, this period is considered the case of a bull market [Figure 3]

	Dependent variable					
	Strategies					
	D2	D4	D6	C2	C4	C6
Alpha	0.004	0.006	0.003	0.007	0.005	−0.001
	t=2.857***	t=1.725*	t=2.074**	t=2.143**	t=1.889*	t=−0.419
MP	−0.168	−0.138	0.020	−0.095	0.031	0.299
	t=−1.199	t= −0.982	t=0.209	t = −0.389	t=0.098	t=1.286
SMB	−0.013	0.657	0.212	−0.413	0.318	0.183
	t=−0.133	t=1.665*	t=1.953*	t = −1.475	t=0.982	t=0.675
HML	0.101	−0.024	0.066	−0.175	0.294	−0.093
	t=0.974	t= −0.134	t=0.905	t = −0.831	t=1.429	t=−0.469
PR1YR	−0.022	−0.027	−0.262	−0.174	0.315	0.170
	t=−0.216	t= −0.173	t=−1.841*	t=−0.924	t=1.333	t=0.916
LIQ	−0.116	−0.640	0.057	0.231	−0.097	0.404
	t=−0.894	t=−2.086**	t=0.388	t=1.246	t=−0.410	t=2.222**
Observations	126	124	122	126	124	122
Adjusted R^2	−0.014	0.022	0.053	−0.008	−0.005	0.005
Residual Std. Error	0.013	0.033	0.018	0.027	0.032	0.029

The asset pricing model is tested on the weekly strategy excess return series. The daily factors from Ødegaard (nd) are compounded to weekly factors. The t-statistics are HAC consistent, using the Newey-West method with 4 lags. Significance codes: ***0.01, **0.05, *0.1

Table 5: Fama-French 3 Factor+Momentum+Liquidity from July 2014 through March 2016, this period is considered the case of a volatile sideways moving market [Figure 3]

	Dependent variable					
	Strategies					
	D2	D4	D6	C2	C4	C6
Alpha	0.002	0.001	−0.001	0.004	−0.001	−0.004
	t=1.801*	t=0.551	t=−0.645	t=1.915*	t=−0.670	t=−2.312**
MP	0.007	−0.056	0.171	0.150	0.105	0.124
	t=0.186	t=−1.267	t=1.528	t=1.450	t=1.106	t=0.768
SMB	−0.105	0.037	0.094	0.251	0.140	0.451
	t=−1.358	t=0.504	t=0.586	t=1.830*	t=1.099	t=2.326**
HML	−0.112	−0.077	−0.029	0.212	0.183	−0.084
	t=−2.174**	t=−1.641	t=−0.369	t=2.264**	t=1.523	t=−0.773
PR1YR	0.025	0.011	0.011	0.009	−0.043	−0.029
	t=0.630	t=0.329	t=0.248	t=0.084	t=−0.810	t=−0.400
LIQ	0.094	−0.078	0.020	−0.147	−0.067	−0.080
	t=1.768*	t=−1.005	t=0.231	t=−1.284	t=−0.543	t=−0.537
Observations	92	92	92	92	92	92
Adjusted R^2	0.030	−0.037	0.040	0.025	−0.003	0.057
Residual Std. Error	0.008	0.012	0.010	0.018	0.018	0.019

The asset pricing model is tested on the weekly strategy excess return series. The daily factors from Ødegaard (nd) are compounded to weekly factors. The t-statistics are HAC consistent, using the Newey-West method with 4 lags. Significance codes: ***0.01, **0.05, *0.1

subsets can be seen from Figure 3, where the OSE All Share Index shows a bull period in the first 2½ years, while, the second period, highlighted in grey, was quite volatile and sideways-moving. We acknowledge that splitting the sample in this way may seem a bit arbitrary, and that this information would not have been available for a trader when trading. However, it may give us some insights on how the PT strategy works under varying market conditions. Given these insights, traders can implement forward looking measures/proxies for the market conditions and adjust the strategy accordingly. Our findings of the regressions of the subsamples are found in Tables 4 and 5. We observe that during the first period, in a bull market, the D6, D6 and C2 strategies have significant alphas on a 5% significance level, while the D4 and C4 strategies alphas are significant on a 10% level. However, in the case of a sideways moving market, only the D2 and C2 strategies has significant alphas, albeit on a 10% significance level. Hence, a weaker performace in the sideways market, with a much lower and less significant alpha estimates. Splitting the sample in two, in general, does not change our conclusion regarding the exposure to the other risk factors.

As mentioned, the alphas of the D4-6 and C4-6 strategies are no longer significant in the second subsample. This can be associated with the excess return contribution from the long and short legs, respectively. As seen in Table 6, the long position provides a higher return than does the short position of the strategy in five out of the six cases. Hence, this strategy performs better in a bull market period. This finding contradicts the findings of previous studies, which showed that PT performs better in bear market periods (Gatev et al., 2006; Do and Faff,

Table 6: Individual leg return contribution

	Distance 2 weeks	Distance 4 weeks	Distance 6 weeks	Coint. 2 weeks	Coint. 4 weeks	Coint. 6 weeks
Long leg return contribution						
Estimate	0.002	0.004	0.002	0.004	0.004	−0.0004
	t=2.765*	t=1.833	t=1.726	t=3.440**	t=2.354*	t=−0.306
Short leg return contribution						
Estimate	0.001	−0.001	−0.0005	0.001	−0.001	−0.001
	t=0.980	t=−0.762	t=−0.584	t=0.457	t=−0.913	t=−0.790

This table show the average excess return for the long and short position respectively, thus how each leg of the transaction contribute towards the strategy's profits. In the top panel the excess returns of the long position is shown, while the bottom shows the short position excess returns. The t-statistics are HAC consisting using the Newey-West method with 4 lags. Significance codes: ***0.01, **0.05, *0.1

2010). The same studies find that the excess return is mostly generated from the short leg, explaining the good performance in bear markets. Jacobs (2015), does however find that the long leg of the transaction has the strongest performance, which is in line with our study. We find the long positions to be significantly greater than the short positions, except for the C6 strategy, when performing a t-test on the long and short excess return series. This can imply that the long positions (losers) have mean-reverting attributes to a larger extent than do the short positions (winners).

5. DISCUSSION AND CONCLUSION

Our study investigates the performance of HFPT using a large dataset from the OSE. We find several points that provide better insight into the impact of the formation period on profit, while comparing distance versus cointegration as a superior approach and profitability dependent on market conditions.

The formation period indeed matters. Huck (2013) focuses on the length of the formation period in PT strategies. However, there is no clear finding on what is preferable and, as mentioned, his study is on daily data. Our data, in a high-frequency setting, show that the shorter the formation (and trading) period, the better the profit. These results are clearly in favor of the shortest time strategy, regardless of which approach (distance or cointegration) is used.

Why do the data provide such evidence? One possibility might be that a more sensitive indicator leads to more profitable trades. This enables the exploitation of smaller, but still profitable, opportunities, even when transaction costs are taken into account. Including more data can therefore results in a more stale indicator, not up to speed with current market conditions.

This study also provides input on the comparison between the two most popular methods for the PT strategy. We find that each approach has their strengths and weaknesses, i.e. the distance approach is less risky, while the cointegration approach is less trading intensive. Still, we are unable to separate the two on a risk-adjusted basis. Using daily data, the distance approach is better according to Rad et al. (2016), while cointegration is superior according Huck and Afawubo (2015). Using HF data, we are unable to conclude on this point. Still, such a comparison of the approaches has not been previously conducted in the HFPT literature.

Previous PT studies claim that the HFPT strategy performs especially well in bear market periods. However, our findings are different, in that we find it to perform better in a bull market period, when compared to a sideways-moving, volatile market.

We add to a small body of literature on HFPT and acknowledge the need for more research in order to better understand how HFT works in different markets and contexts. Our study may be most closely related to Bowen et al. (2010) and Stübinger and Bredthauer (2017); however, we encourage more PT studies on more data, especially concerning HFPT.

6. ACKNOWLEDGMENTS

Frode Kjærland acknowledges the support from the regional research fund of Mid-Norway for the project "Styring og samarbeid i offentlig sektor."

REFERENCES

Bailey, D.H., Lopez de Prado, M. (2012), The sharpe ratio effcient frontier. Journal of Risk, 15(2), Available from: https://www.dx.doi.org/10.2139/ssrn.1821643.

Bogomolov, T. (2011), Pairs Trading in the Land Down Under. In: Finance and Corporate Governance Conference.

Bolgün, K.E., Kurun, E., Güven, S. (2010), Dynamic pairs trading strategy for the companies listed in the Istanbul stock exchange. International Review of Applied Financial Issues and Economics, 1, 37-57.

Bollinger, J. (1992), Using Bollinger bands. Stocks and Commodities, 10(2), 47-51.

Bowen, D., Hutchinson, M., O'Sullivan, N. (2010), High frequency equity pairs trading: Transaction costs, speed of execution and patterns in returns. Journal of Trading, 5(3), 31-38.

Broussard, J. P., Vaihekoski, M. (2012), Profitability of pairs trading strategy in an illiquid market with multiple share classes. Journal of International Financial Markets, Institutions and Money 22(5), 1188-1201.

Carhart, M.M. (1997), On persistence in mutual fund performance. The Journal of Finance, 52(1), 57-82.

Chordia, T., Goyal, A., Lehmann, B.N., Saar, G. (2013), High-Frequency Trading. Johnson School Research Paper Series (20-2013).

Clegg, M., Krauss, C. (2018), Pairs trading with partial cointegration. Quantitative Finance, 18(1), 121-138.

Do, B., Faff, R. (2012), Are pairs trading profits robust to trading costs? Journal of Financial Research, 35(2), 261-287.

Do, B., Faff, R. (2010), Does simple pairs trading still work? Financial Analysts Journal, 66(4), 83-95.

Dunis, C.L., Giorgioni, G., Laws, J., Rudy, J. (2010), Statistical Arbitrage And High-Frequency Data with an Application to Eurostoxx 50 Equities. Liverpool Business School, Working Paper.

Dunis, C.L., Laws, J., Middleton, P.W., Karathanasopoulos, A. (2015), Trading and hedging the corn/ethanol crush spread using time-varying leverage and nonlinear models. The European Journal of Finance, 21(4), 352-375.

Engelberg, J., Gao, P., Jagannathan, R. (2009), An Anatomy of Pairs Trading: The Role of Idiosyncratic News, Common Information and Liquidity. In: Third Singapore International Conference on Finance.

Fallahpour, S., Hakimian, H., Taheri, K., Ramezanifar, E. (2016), Pairs trading strategy optimization using the reinforcement learning method: A cointegration approach. Soft Computing, 20(12), 5051-5066.

Fama, E.F., French, K.R. (1993), Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3-56.

Gatev, E., Goetzmann, W.N., Rouwenhorst, K.G. (2006), Pairs trading: Performance of a relative-value arbitrage rule. Review of Financial Studies, 19(3), 797-827.

Hagströmer, B., Norden, L. (2013), The diversity of high-frequency traders. Journal of Financial Markets, 16(4), 741-770.

Huck, N. (2013), The high sensitivity of pairs trading returns. Applied Economics Letters, 20(14), 1301-1304.

Huck, N., Afawubo, K. (2015), Pairs trading and selection methods: Is cointegration superior? Applied Economics, 47(6), 599-613.

Jacobs, H. (2015), What explains the dynamics of 100 anomalies? Journal of Banking and Finance, 57, 65-85.

Jacobs, H., Weber, M. (2015), On the determinants of pairs trading profitability. Journal of Financial Markets, 23, 75-97.

Johansen, S. (1988), Statistical analysis of cointegration vectors. Journal of Economic Dynamics and Control, 12(2), 231-254.

Kim, K. (2011), Performance Analysis of Pairs Trading Strategy Utilizing High Frequency Data with an Application to Kospi 100 Equities. Available from: http://www.SSRN1913707.

Krauss, C. (2017), Statistical arbitrage pairs trading strategies: Review and outlook. Journal of Economic Surveys, 31(2), 513-545. Available from: https://www.onlinelibrary.wiley.com/doi/abs/10.1111/joes.12153.

Krauss, C., Do, X.A., Huck, N. (2017), Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the s&p 500. European Journal of Operational Research, 259(2), 689-702.

Krauss, C., Stu¨binger, J. (2017), Non-linear dependence modelling with bivariate copulas: Statistical arbitrage pairs trading on the s&p 100. Applied Economics, 49(52), 5352-5369.

Li, M.L., Chui, C.M., Li, C.Q. (2014), Is pairs trading profitable on china ah-share markets? Applied Economics Letters, 21(16), 1116-1121.

Liew, R.Q., Wu, Y. (2013), Pairs trading: A copula approach. Journal of Derivatives and Hedge Funds, 19(1), 12-30.

Lintilhac, P.S., Tourin, A. (2017), Model-based pairs trading in the bitcoin markets. Quantitative Finance, 17(5), 703-716.

Liu, B., Chang, L.B., Geman, H. (2017), Intraday pairs trading strategies on high frequency data: The case of oil companies. Quantitative Finance, 17(1), 87-100.

Mashele, H., Terblanche, S., Venter, J. (2013), Pairs trading on the johannesburg stock exchange. Investment Analysts Journal, 78, 13-26.

Miao, G.J. (2014), High frequency and dynamic pairs trading based on statistical arbitrage using a two-stage correlation and cointegration approach. International Journal of Economics and Finance, 6(3), 96.

Mikkelsen, A. (2018), Pairs trading: The case of norwegian seafood companies. Applied Economics 50(3), 303-318.

Næs, R., Skjeltorp, J., Ødegaard, B.A. (2009), What Factors Affect the Oslo Stock Exchange. Norges Bank (Central Bank of Norway), Working Paper.

Nath, P. (2003), High Frequency Pairs Trading with us Treasury Securities: Risks and Rewards for Hedge Funds. SSRN eLibrary.

Ødegaard, B. (nd.), Asset Pricing Data at OSE. Available from: http://www1.uis.no/ansatt/odegaard/financial_data/ose_asset_pricing_data/index.html.

Perlin, M. (2009), Evaluation of pairs-trading strategy at the Brazilian financial market. Journal of Derivatives and Hedge Funds, 15(2), 122-136.

Rad, H., Low, R.K.Y., Faff, R. (2016), The profitability of pairs trading strategies: Distance, cointegration and copula methods. Quantitative Finance, 16(10), 1541-1558.

Stu¨binger, J., Bredthauer, J. (2017), Statistical arbitrage pairs trading with high-frequency data. International Journal of Economics and Financial Issues, 7 (4), 650-662.

Stu¨binger, J., Endres, S. (2018), Pairs trading with a mean-reverting jumpdiffusion model on high-frequency data. Quantitative Finance, 1-17. Available from: https://www.doi.org/10.1080/14697688.2017.1417624

Stu¨binger, J., Mangold, B., Krauss, C. (2016), Forthcoming. Statistical arbitrage with vine copulas. Quantitative Finance.

TITLON, (nd.), Available from: https://www.titlon.uit.no/.

Verousis, T., ap Gwilym, O. (2010), An improved algorithm for cleaning ultra-high-frequency data. Journal of Derivatives and Hedge Funds, 15(4), 323-340.

Vidyamurthy, G. (2004), Pairs Trading: Quantitative Methods and Analysis. Wiley: Wiley Finance.

Yang, J.W., Tsai, S.Y., Shyu, S.D., Chang, C.C. (2016), Pairs trading: The performance of a stochastic spread model with regime switching-evidence from the s&p 500. International Review of Economics and Finance, 43, 139-150.