

Pairs trading: the case of Norwegian seafood companies

Andreas Mikkelsen^a

^aNord University Business School, 8049 Bodø, Norway

ARTICLE HISTORY

Compiled April 10, 2017

ABSTRACT

In this paper, I investigate the performance of a pairs trading strategy on 18 seafood company stocks traded in the Norwegian consumer goods sector on the Oslo Stock Exchange. I apply both high-frequency and daily data from January 2005 to December 2014. I use two approaches – a distance approach and a cointegration approach – and compare the results. For both the distance and the cointegration approaches, nonconvergence of the pairs is high, which may indicate that more fundamental information about the companies traded should be accounted for. None of the strategies evaluated had significant profits after accounting for transaction costs. It therefore remains unclear which approach is best suited for pairs selection. Using high-frequency data yielded empirical distributions that were symmetrical and had a lower degree of leptokurtosis compared to the daily data.

KEYWORDS

Pairs trading, Statistical arbitrage, High-frequency trading, Norwegian seafood companies, Cointegration

JEL CLASSIFICATION

G11, G15

I. Introduction

In this paper, I investigate a trading strategy called pairs trading. In short, this strategy relies on trading two stocks that have historically moved in very similar patterns. However, once the pattern is broken, it is believed to return after a short period. The strategy buys one stock long and the other stock short. A thorough explanation follows.

Pairs trading is a market-neutral strategy that exploits equilibria between securities. When a pair of securities diverges from its equilibrium, a long–short position is entered and, on convergence to the equilibrium, positions are reversed that, by configuration, lead to a profit. To find pairs suitable for trading, a distance approach and a cointegration approach are employed. I use a 12 month formation, or a training period, and a subsequent 6 month trading period; thus, the study is comparable with other empirical work. The Norwegian stock market is small in both trading volume and market capitalization relative to the US stock market. To evaluate whether the use of high-frequency data adds benefits to the analysis of an illiquid market, I use

both high-frequency trading data and daily data. Therefore, the research contribution is threefold: I evaluate whether the pairs trading strategy performs well on stocks in the Norwegian seafood sector using a 12 month formation period and subsequent 6 month trading period. Comparing pairs trading selection methods. Finally, I compare high-frequency to daily data frequency.

II. Pairs trading

Pairs trading, also known as statistical arbitrage, is a well-known speculative trading strategy. Invented by investment banks on Wall Street in the 1980s, it proved to be successful during that time.

The main concept of the trading strategy is quite simple: identify two traded securities that comove in a given period. If they diverge from the equilibrium, buy the security that is relatively underpriced and sell the security that is overpriced. One speculates that the securities will reverse to their equilibrium, at which point the trades are reversed. The trader nets a profit from the transactions, by configuration. In practice, the spread of the securities is being bought and sold. Pairs trading is a market-neutral strategy, meaning that profits are realized regardless of the direction in which the market is moving.

The theoretical framework from which we understand this phenomenon is the ‘Law of One Price’ (LOP), which states that two identical goods should sell at the same price. If this statement is not true, arbitrage opportunities will exist that enable the good to be bought in the market in which it is the least expensive and to be immediately sold in the other market for a profit. In the pairs trading setting, we observe securities that have a high degree of comovement and, thus, should be close substitutes. Thus, the strategy is viewed as an enforcement of the LOP and drives prices back to their equilibrium value.

Numerous methods exist for identifying tradable securities. In the literature, at least the following three selection methods for pairs trading are discussed:

- The distance approach
- Cointegration
- Stochastic spread method

The distance approach proposed by Gatev, Goetzmann, and Rouwenhorst (2006) is the chosen approach for the majority of the empirical work on pairs trading; see for example: Perlin (2009); Do and Faff (2010); Broussard and Vaihekoski (2012). 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’ (SSD) – in other words, pairs with a small distance between the normalized price series. Cointegration is also a technique that is quite common to use. Cointegration is based on work done by Engle and Granger (1987), who find that two or more time series can have common factors that drive their evolution. The common factors ensure that a linear combination of the two with a stationary long-term equilibrium and a finite variance can be found. Stochastic spread models for pairs trading have also been suggested; see Elliott, Van Der Hoek, and Malcolm (2005) and Do, Faff, and Hamza (2006). Liew and Wu (2013) suggest that because asset returns are not normally distributed, the linearity assumption implied in the distance and cointegration approach does not hold. Therefore, copulas are better suited to model dependency structures between two stocks, and a copula approach for pairs trading is

sensible.

Common to all approaches is that data sample is ‘split’ in two parts. The first part consists of a formation, or training, period. In this period, I evaluate the pairs that show the best characteristics for trading. In the second subsequent part, which does not overlap with the first, the pairs are traded.

Once a selection approach is chosen and pairs are found that seem promising for trading, a trading algorithm must be operationalized. However, doing so is not straightforward: when is a trade opened in a pairs trading setting? In the literature, a trigger based on a SD metric is commonly used to decide when to enter a trade. There is not a ‘correct’ way of choosing the threshold, 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 possibly eradicates any profits earned from the strategy. Naturally, a very high trigger results in fewer trades; however, if too high, 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 long–short positions. Huck (2013) demonstrates how pairs trading returns are sensitive to the length of the formation period and the trading trigger.

The developed global financial markets should provide endless opportunities for trading because one can find pairs suitable for trading among derivatives, ETFs, commodities, equities, and so forth. However, studies of the phenomenon suggest that this may not be the case, at least not for the US market. In fact, the strategy was profitable from the 1960s, but with declining profits until 2000 according to Gatev et al. (2006). From 2000 onwards, Do and Faff (2010) show that the strategy is still profitable, but with further declining performance. However, when accounting for trading costs, Do and Faff (2012) find that, on average, pairs trading is unprofitable. Possibilities still exist to earn profits if pairs are carefully selected within narrowly defined industries and by utilizing the SSD criterion in addition to a ‘Number of Zero Crossing’ (NZC) criterion. The NZC measures how many times the sign of the spread changes. Therefore, a good candidate for trading has a high NZC, indicating good mean-reverting attributes, in addition to a low SSD. The decline in profitability is explained by two factors. First, increasing competition for these ‘free lunch’ pairs trading configurations has eliminated or limited them. The second explanation is that arbitrage risk has increased, such as the risk related to trade execution or nonconvergence of the pairs once opened. Therefore, finding long-term stable equilibriums between securities is difficult. Having observed the decline in profitability, Rad, Low, and Faff (2016) compares the distance, cointegration and copula approaches, to see if the more sophisticated cointegration and copula approaches performs better than the distance approach. They use daily data from the US, finding that the copula approach does not perform as well as the other two approaches.

An exception to the declining performance is during bear markets, for which Do and Faff (2010) finds that the strategy performs particularly well, which was true during the dot-com bubble of 2000 to 2002 and the Global Financial Crisis (GFC) of 2007 to 2009. They explain the performance with reduced market efficiency, which confirms the finding of Gatev et al. (2006), who show that their strategy performed exceptionally well during the 1969 to 1980 stock market decline. It is well known that there exists a leverage effect in the financial markets, where a negative correlation exists between past returns and future volatility. Therefore, a negative shock has a stronger price impact than an equivalent positive shock. This phenomenon results from leveraged investors possibly receiving margin calls or an increase in collateral requirements when markets plummet. Thus, in such a market, arbitrageurs may face

limitations on access to capital and be unable to enforce the LOP, making pairs trading profitable.

Given that a number of studies find that pairs trading is profitable, one may ask what drives the profitability. Engelberg, Gao, and Jagannathan (2009) studies profits and drivers of profits in pairs trading of US stocks. Drivers of profitability may include liquidity provisions and price discovery, where an immediate liquidity provision has a short-term effect on profits, but also a longer-term effect on illiquid stocks. Liquidity also has a stronger effect on small cap stocks than large cap stocks. Price discovery also affects profits. To test for the price discovery effect, they study how quickly stocks react to new information and find that industry-specific information diffuses into prices at different rates for different stocks, a differential that creates profitable trading opportunities. In contrast, firm-specific information may cause permanent shifts in equilibrium and negatively affects profits. In line with their hypothesis regarding price discovery, they find that both stocks in a pair being held by a large number of institutional holdings have a negative effect on profits. That both stocks are covered by the same analyst also has a negative impact on profits because of the lower difference in diffusion rates in both cases.

Studies from non-US markets, which are smaller concerning both liquidity and market capitalization, reveal promising results and mostly find positive returns; see Hong and Susmel (2003), Perlin (2009), Bolgün, Kurun, and Güven (2010), Broussard and Vaihekoski (2012), Mashele, Terblanche, and Venter (2013) and Li, Chui, and Li (2014). Bogomolov (2011) studies the Australian market using the distance, cointegration, and stochastic spread method. The returns on these strategies from the Australian market were not robust for trading costs.

The studies that implement high-frequency data are very limited. Nath (2003) studies pairs trading on US treasury securities using the distance approach with positive performance. Bowen, Hutchinson, and O’Sullivan (2010) uses the distance approach on FTSE100 equities between January 2007 and December 2009 using 60-minute data. The formation period is 264 hours and the following trading period is 132 hours. They find that the strategy’s returns are sensitive to both transaction costs and speed of execution. When waiting one period for the trade and adding transaction costs of 15 basis points, profits are eliminated. Dunis, Giorgioni, Laws, and Rudy (2010) studies the equities in the Eurostoxx50 index between 3 July 2009 and 17 November 2009. They use 5-, 10-, 20-, 30-, and 60-minute data, in addition to daily data. They also use a cointegration approach, in addition to employing the Kalman filter for time-varying coefficient estimation. They also seem to use a one-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 find positive results. Finally, Kim (2011) studies equities listed on the KOSPI100 index in the Korean market. A cointegration approach is chosen for this study as well, with the Kalman filter used to estimate the time-varying coefficients. The data are in 30-minute intervals, with a 2 week formation period and a 1 week trading period. The findings suggest that the strategy results in positive excess returns after transaction costs are taken into account and trading occurs one period after the trading signal. The strategy performs better during bear markets, in this case during the GFC. In addition two recent papers are relevant for high-frequency pairs trading, but with slightly different scope than the previous papers, namely Fallahpour, Hakimian, Taheri, and Ramezanifar (2016) and Liu, Chang, and Geman (2017). The first paper shows how a reinforced learning approach is superior in retrieving the best parameters for a cointegration approach strategy, on a S&P500 dataset from June 2015 to January 2016.

While, Liu et al. (2017) use a more dynamic way to identify mispricings between pairs, than the distance and cointegration method. They use this methodology on oil companies listed on NYSE, with data in 5-minute intervals, in 2008 and from June 2013 to April 2015, with positive results in both periods. Finally, a thorough pairs trading literature review is found in Krauss (2016).

III. Data and methodology

The dataset consists of high-frequency tick-by-tick transaction data of the consumer goods sector, sector id 30, acquired from the Oslo Stock Exchange. The reason this particular sector is selected is that a large number of the companies are in the fishing industry, in either commercial fishing or aquaculture (fish farmers). For the pairs trading strategy to work, the securities should be close economic substitutes and, as such, have common factors driving the price evolution. Fish farmers share many risk factors, such as income, for which the sale prices of salmon fluctuate on commodity exchanges such as FishPool; cost of fish fodder; diseases; and political risk, such as the 2014 Russian import embargo on Norwegian seafood. One can also argue that white fish and salmon are substitutes: if one becomes too expensive, consumers shift their demand to the other. In that sense, that an equilibrium exists – even between commercial fishing companies and aquaculture companies – is not unlikely. Three companies in the sector were processed food and dairy producers and, therefore, not obviously linked to the fishing companies. For this reason, they were omitted from the sample. Fig. 1 shows the cumulative logarithmic return series of the fish farmers Grieg Seafood, Lerøy Seafood, Marine Harvest, and Salmar. A review of the figure indicates that series evolution seems to be strongly correlated.

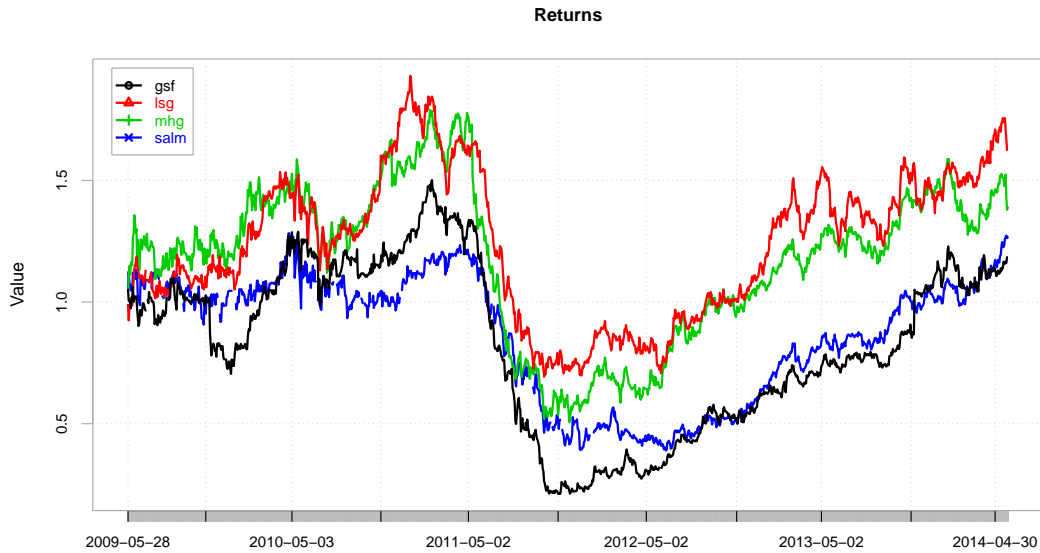


Fig. 1. Return series Grieg Seafood, Lerøy Seafood, Marine Harvest, and Salmar from June 2009 to June 2014

The range is a ten-year period from January 2005 to December 2014 and consists

of approximately 9 million trades in 18 shares. The first half of the period saw approximately 2.5 million trades and the second half of the period saw a more than doubling of volume, to 6.3 million trades – indicating a substantial increase in liquidity. The data are transformed from trade-by-trade data to five-minute intervals in OHLC (Open-High-Low-Close) format to ensure evenly spaced time series. If no trading occurred during the interval, the previously traded price is used to prevent missing values. To account for dividends and corporate events, an adjusted price for each share is calculated using adjustment factors from the TITLON (nd) database.

As previously mentioned, daily data are also used. The daily data consist of adjusted daily closing prices for the same date range as the high-frequency data downloaded from the TITLON (nd) database.

Pair formation

High-frequency data. Following the distance approach methodology by Gatev et al. (2006), the prices are normalized to start at NOK 1. The spread is then calculated by subtracting the normalized price of one stock from the other. The difference is squared to ensure that both negative and positive deviations are accounted for. The pairs with the smallest SSD, or distance, in the spread are considered suitable for trading because they have a high level of comovement. As mentioned, the Norwegian stock market is much smaller than the US stock market, both with respect to the trading volumes and market capitalization of the companies listed. Further, the sheer number of companies in each sector in the US market is important when seeking potential economic substitutes. In the case of the Norwegian market, the consumer goods sector, including mostly seafood companies, seemed to be the most obvious option to ensure that the stocks in a pair are economic substitutes to some degree.

For this reason, the sample is quite small; therefore, I cannot afford to exclude companies if they do not have trading activity for a day – as is the case is with studies from the United States – simply because no stocks would remain in the sample. Further, because of the small sample of stocks, I don’t exclude pairs with nonoverlapping time series. Nonoverlapping time series can be the result of new listings on the stock exchange or the delisting of a stock in the middle of a formation period. In this case, one of the stocks may have a time series from January to October, whereas the other stock may have a time series from March to December. Thus, when a nonoverlapping time series exists, the result is missing values where they do not overlap when comparing the two series. There is no selections on minimum market capitalization on the stocks.

The missing values are ignored in the SSD calculations, and the spread is only computed when observations for both stocks exist at a given time stamp. To ensure minimum liquidity, I require more than 10 800 observations for each stock in the formation period. This amount equals approximately 50% of the full set of possible observations in a 1 year formation period. Using a 1 year formation period allows us to compare our results to previously mentioned studies to determine whether our sample exhibits the same characteristics as shown in other studies.

$$SSD = \frac{1}{n} \sum_{t=1}^n (P_t^1 - P_t^2)^2 \quad (1)$$

where P_t^i is the normalized price for stock i at time t . In our sample, 1083 pairs

passed the selection criteria, with at least 50% of the time series overlapping each other. SSD values were calculated and averaged 60 pairs per trading period. For the trading strategy to be successful, it is wise to choose the best pairs – in this case, the pairs with the shortest distance between stocks as measured by the SSD. In this example, I select the top five pairs from every period regardless of the amount of pairs available. However, in some periods, such a large number of pairs is not available; in this case, I choose all available pairs. Thus, the trading sample consists of 46 pairs instead of the 90 pairs possible, if five pairs were available every trading period.

Daily Data. For the daily data, I follow the same procedure as for the high-frequency data. Naturally, the daily data have fewer observations than the high-frequency data for the same period. Therefore, I require that the time series for both stocks in a pair has at least 120 observations during the formation period, which translates to approximately 6 months with 20 trading days in each month, or 50% of a full sample.

Cointegration. A cointegration approach is also applied. I require that the missing values be less than approximately 50% if the formation period is similar to the distance approach. For the 12-month formation, this equals 10 800 observations. If both series have fewer missing values than the requirement, I create a matrix and include only those observations for which data for both time series exist at a given timestamp. This is a requirement for the cointegration test.

I test whether pairs are cointegrated using the Johansen (1988) procedure. The test used is a max eigenvalue test and is performed without a trend or constant specified in the formulation. The number of lags used in the VAR specification is based on the AIC criteria. If cointegrated, a linear combination of the two stocks that specifies the long-run equilibrium μ can be found, with the cointegration vector β and the stationary residual series ϵ as shown in Equation 2. Contrary to the distance method for which only the top five pairs are chosen in each period, all pairs that are found to be cointegrated at a 10% significance level or better are chosen for trading.

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

The cointegration vector normalized with regard to the first stock is used to determine the mean of the system during the formation period. A divergence of two SD from the mean is a trading signal. Because the cointegration vector indicates the long-run equilibrium of the system – one unit of P_t^1 and β units P_t^2 – I also use this relationship when entering trades as opposed to the distance approach, which uses even long/short legs. For example, if β is 2, I long 1 unit of stock A and short 2 units of stock B.

Trading period

As shown in Fig. 2, the formation period is twice the length of the trading period. The first 12 months from January 2005 to December 2005 are used for pairs formation. The first possible month for trading is January 2006.

Gatev et al. (2006) roll each trading period forward by one month, which results




Fig. 2. Formation and trading periods

in up to six unique portfolios in a given month using a 6 month trading period. Because of the limited liquidity of the Norwegian stock market relative to the US market, such a trading regime is likely to be a theoretical construction because trading frictions may make it impossible to enter or exit positions when needed. Given the possible restrictions and low number of shares in the sample, overlapping trading periods are not used in this study. Thus, the 1 year formation periods are from January to December or July to June in the calendar year, as recommended in Broussard and Vaihekoski (2012). The obvious drawback is that every month will contain fewer pairs in the portfolio relative to having six portfolios and rolling 1 month forward.

An indicator function is employed to create signals for trades. This indicator function uses the adjusted closing price, thus avoiding that spikes that originate from corporate events or dividend payments create a signal. Trades are performed on the next timestamp using the open price. The reason for choosing the open price as the trade price is that I want trades to occur as soon as possible after a trading signal has been triggered. This situation has an effect only when using high-frequency data because our daily data consist only of adjusted close prices. In this case, the trade is executed on the subsequent timestamp after the signal, in other words, the next day's adjusted close price.

For the distance approach, the methodology Gatev et al. (2006) proposed is applied. A trade is initiated if the spread of the pair has diverged by more than two SD from the mean spread, as recorded in the formation period. When a trade is initiated, both legs are equally weighted with NOK 1 long and NOK 1 short. The positions are reversed if the spread crosses the historical mean from either direction or at the end of the trading period. If a stock is delisted from the stock exchange or the sector, the position is reversed as well. No stop-loss mechanism exists. In addition, only one open long/short position in a pair is allowed at a time. Transaction costs are included in the study. The rate used is 0.039% of the total trade value per trade, which is the rate offered to private investors by Internet brokers operating in Norway. Thus, one might consider this rate the ceiling for transaction costs. At present, two times the transaction costs is accounted for when entering the position, whereas no cost is accounted for when exiting, which may introduce a small bias. Further, in the case of selling short, it is common practice for the broker to require a margin deposit or other collateral to reduce its credit risk. I have not accounted for this requirement for two reasons. First, institutional investors may put up other liquid assets as collateral for short sales. Second, I use a very high rate for transaction costs, making the return estimates very conservative.

Returns for the pairs are calculated on a daily basis. Calculating the return per pair $R_{pair,t}$ is done by taking the change in net equity (profit or loss), labelled ΔPL , on both the long and the short position from time $t - 1$ to t . The change in equity is then

divided by the initial exposure on each position, in addition to transaction costs, as shown in Equation 3. Thus, a return series for both the long and the short positions is available. A pair may trade numerous times per trading period. For instance, a position may be closed early in the trading period, earn a profit or a loss, and later open and close again. In this case, the cash sitting in the account between trades does not accrue interest.

The daily average return R_t is what Gatev et al. (2006) labels a fully invested return. I summarize the return per pair from all available pairs at time t and divide it by the number of pairs trading at time t , as shown in Equation 4. Finally, the daily returns are compounded to obtain monthly returns, labelled R_m , as shown in Equation 5.

$$R_{pair,t} = \frac{\Delta PL_{long}}{InitialExposureLong + \frac{Txn}{2}} + \frac{\Delta PL_{short}}{InitialExposureShort + \frac{Txn}{2}} \quad (3)$$

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

$$R_m = \prod_{t=1}^n (1 + R_t) - 1 \quad (5)$$

IV. Results

Distance approach: High-frequency data

The sample consists of 46 pairs ready for trading. However, as shown in panel A of Table 2, only 28 of the pairs were trading. This scenario is also evident when analysing the trading data from the first trading month of January 2006 until December 2009, which shows that trading occurred in only 3 of the 48 months. In the subsequent period, 13 months do not have trading activity, as shown in Fig. 3. The lack of trading may be attributed to the low trading volume before 2010 because the stocks in a pair need to have overlapping time series, when volume is low, the probability is lower that pairs with overlapping time series exist. Considering that only 18 stocks are in the sample and were trading at some point during our sample, the actual number of stocks available for trading at any given month is lower than this number because of IPOs and delistings. Choosing shorter formation and trading periods is likely to produce more trading pairs because the chance is higher that the time series overlap.

To test whether the trading strategy has produced any significant profits, I perform a Newey–West t-test using four lags and three degrees of freedom on the monthly return series. The results provide the estimate that the mean monthly excess return is -0.7% (t-statistic: -1.201), when including all months in the sample, also those without trading, which is not significantly different from zero. If I only look at the results in the months in which trading occurs, the estimated excess return is actually

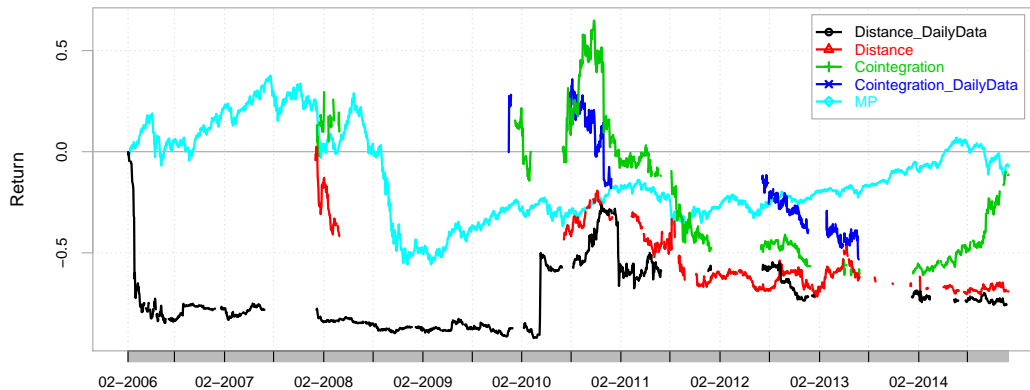


Fig. 3. Performance of the pairs trading strategy

worse, at -1.75% . Fig. 3 shows a chart of the cumulative return of the strategy. The market premium is used as a benchmark.

The performance of the strategy is as expected when reviewing the literature from the United States, where Gatev et al. (2006) and Do and Faff (2010) find a declining trend in pairs trading performance. Do and Faff (2012) finds that pairs trading was on average not profitable in recent years, which is also evident in Huck and Afawubo (2015) when including transaction costs for the distance approach.

In contrast, the results are quite different from Broussard and Vaihekoski (2012) from the Finnish market, which is very surprising. One would expect the Norwegian market to mimic the Finnish market closer than the US market. However, this result may be because the study on the Finnish market consisted of trading on the same underlying in multiple asset classes, which were closer economic substitutes. Moreover, the sample and sample period in this study might be nonrepresentative of the Norwegian market. Further, when reviewing Engelberg et al. (2009), one assumes that the Norwegian market would be a good candidate for pairs trading given the lower degree of liquidity and smaller market capitalization of the companies listed in this study, factors that they found increased the pairs trading profits. What one considers a small cap stock in the United States may very well be a large cap stock in Norway; therefore, the Norwegian market may be too small and too illiquid for the strategy to work as well.

Diversification plays an important role in pairs trading, as it does in other trading strategies. When the number of pairs in a portfolio increases, Gatev et al. (2006) finds that the SD of the portfolio decreases. In addition, they find that the minimum realized return increases, whereas the maximum return is stable. Considering that our strategy allows for a maximum of five pairs at a time, lack of diversification may be the reason that the SD of the monthly returns is 6.92% , which is three to four times higher than that found in other studies. This SD is also higher than that for the monthly returns of the market index (not reported), which is not very lucrative.

The average holding period is 56 days, which equals approximately three calendar months. This length is less than the length of the average holding period of 3.75 months as reported by Gatev et al. (2006) and longer than the 22–37 days as reported

Table 1. Monthly return distribution

	Distance 1 year	Distance Daily 1 year	Cointegration 1 year	Cointegration Daily 1 year
N	50	80	44	18
Average excess return	-0.00729	0.01573	0.00136	-0.00527
Standard error (Newey-West)	0.00607	0.03062	0.00753	0.00449
t-Statistic	-1.201	0.514	0.180	-1.174
Excess return distribution				
Median	0.0000	0.0000	0.0000	0.0000
Standard deviation	0.0692	0.3485	0.0679	0.0464
Skewness	-0.021	8.650	-0.250	1.235
Excess Kurtosis	6.10	86.000	6.234	17.367
Minimum	-0.276	-0.716	-0.284	-0.160
Maximum	0.31	3.540	0.269	0.297
Observations with excess return < 0	0.56	0.52	0.48	0.67

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

by Do and Faff (2010). It is reasonable to consider pairs trading in this setting as a medium-term strategy. The win/loss ratio is 0.31, not a very impressive result for a dollar-neutral strategy. An important observation is that the low number of winning positions is largely attributable to a high number of positions that barely broke even before transaction costs, but that incurred a loss after adding transaction costs.

Both Gatev et al. (2006) and Do and Faff (2010) report that the performance of the pairs trading strategy has declined from 2000 onwards. One of the significant risk factors contributing to the bad performance is the risk of nonconvergence, meaning that once a position is opened, it will never return to the equilibrium as measured in the formation period. This phenomenon is also confirmed by Engelberg et al. (2009). They find that a large portion of pairs trading profits are made in the first days after divergence, whereafter the profits declines. If pairs do not converge during the first few days after a divergence, a convergence is unlikely to happen.

Out of the 75 round trips in the sample, I find that 23 positions are closed because the trading period ends or one of the stocks in the pair is delisted. This situation translates to 30.67% of open positions not converging.

A review of only the percentage of nonconvergent positions may not reveal the full picture. In other words, if I have four pairs, the first pair has five round trip trades and the rest have one nonconvergent trade, then approximately 28% of the positions are nonconvergent. At the same time, 75% of the pairs are nonconvergent, which is naturally quite useful to know. A closer examination of the performance of individual pairs, in panel B of Table 2, shows that 57% of the pairs that trade never converge to their equilibrium.

Compared with Do and Faff (2010), they find that 40% of the pairs in their sample between 2003 and 2009 are nonconvergent. Thus, divergency risk is significantly important. If I have opened a position and the pairs further diverge from their assumed equilibrium, this situation leads to a loss on the position. For instance, the 'overpriced' stock that is shorted increases in value or the 'underpriced' stock, in which I have a long position, depreciates. Because 57% of the pairs never converge, this may contribute to the following. First, the holding period is increased because the pairs will be kept until the end of the trading period. Second, the high number of monthly observations with negative returns contributes to the conclusion that implementing stop-loss rules in the trading algorithm might be a good idea. For instance, if pairs continue to diverge through a 3 or 4 SD band, one may reverse the trades to prevent further losses. Another option is to have a maximum number of days that the position may be open, as suggested by Engelberg et al. (2009).

$$R_t^e = \alpha + \beta_1 MP_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 PR1YR_t + \epsilon_t \quad (6)$$

Table 2. Descriptive statistics. In Panel A, summary trading statistics. In Panel B, the pairs composition of the strategies with regards to convergence and roundtrips.

(a)

	Distance 1 year	Distance Daily 1 year	Cointegration 1 year	Cointegration Daily 1 year
Total pairs	46	84	74	36
Pairs trading	28	54	40	11
Trades	300	264	244	52
Roundtrips	75	66	61	13
Holding period	55.54	61.27	62.41	75.33
Biggest winner	0.62	2.26	0.77	0.35
Biggest loser	-0.18	-0.49	-0.26	-0.27
Win %	0.2360	0.2404	0.5201	0.1288
Loss %	0.7640	0.7596	0.4487	0.8712
W/L ratio	0.309	0.316	1.159	0.148

(b)

	Distance 1 year	Distance Daily 1 year	Cointegration 1 year	Cointegration Daily 1 year
% Position closed at roundtrip	69.33	36.36	50.82	15.38
% Position closed at end of period	30.67	63.64	49.18	84.62
% Non-convergence pairs	57.14	72.22	52.50	90.91
% Multiple roundtrip pairs	28.57	12.96	10.00	9.09
% Single roundtrip pairs	14.29	14.81	37.50	0.00

To further investigate the risk factors that affect pairs trading, I regress the daily return series against the Fama and French (1996) three-factor model, in addition to a Carhart (1997) momentum factor. This is shown in Equation 6. The factor indeces are retrieved from Ødegaard (nd) and the market return is calculated using the Oslo All Shares Index. For the risk-free rate, I use the daily yield on the three-month NIBOR (Norwegian Interbank Offered Rate) interest rate. I subtract the risk-free rate from the market return and retrieve the market premium. Pairs trading is a market-neutral strategy, and this is shown in panel A of Table 3. The four factors explain only a small portion of the returns ($R^2 = 0.005$) and none are significant.

Being able to separate the returns from the long and short leg of the transactions, it's natural to see how each leg separately has contributed towards the returns. If pairs trading is a simple reversal strategy, one can according to Gatev et al. (2006) expect that both legs have equal contribution to the profits or losses. The expectation is due to that a pair may either remain in equilibrium, otherwise it can diverge from its equilibrium in three ways:

- One of the stocks diverges upwards, whereas the other remains stable
- One of the stocks diverges downwards, whereas the other remains stable; and
- Both stocks diverge in opposite directions.

Thus, in these scenarios, it is equally likely that either stock may cause divergence. In contrast, if the profits are driven from other sources as suggested by Gatev et al. (2006), Engelberg et al. (2009) and Do and Faff (2010), I expect the legs to have uneven contributions. The estimates of the long and short returns are presented in panels A and B in Table 4. In the table, it is shown that the returns from the long position contribute slightly more to the total returns of the two. The estimated average monthly long return is -0.011% , whereas the short positions contribute on average -0.733% . However, the returns from the long and short positions are not significantly different.

$$R_t^e = \alpha_{nc}D_{nc} + \alpha_cD_c + \beta_1MP_t + \epsilon_t \quad (7)$$

Table 3. Multifactor market model from Equation 6. Standard errors using the Newey-West method with 4 lags on the monthly excess return series. In Panel A, the distance approach using high-frequency (left) and daily data (right). In Panel B, the cointegration approach using high-frequency data (left) and daily data (right).

		Distance 1 year				Distance Daily 1 year			
	Estimate	Std. Error	t-value	Pr(> t)	Estimate	Std. Error	t-value	Pr(> t)	
Intercept	-0.0007	0.0010	-0.671	0.50	0.0005	0.0018	0.281	0.7791	
Market Premium	0.1527	0.1855	0.823	0.41	-0.2077	0.1869	-1.112	0.2665	
SMB	0.1913	0.1719	1.112	0.27	-0.0529	0.1839	-0.288	0.7737	
HML	-0.1276	0.1788	-0.71	0.48	0.0448	0.1071	0.42	0.676	
PR1YR	-0.1941	0.1718	-1.13	0.26	-0.3158	0.1517	-2.08	0.0375 **	
R ²	0.0049				0.0061				

		Cointegration 1 year				Cointegration Daily 1 year			
	Estimate	Std. Error	t-value	Pr(> t)	Estimate	Std. Error	t-value	Pr(> t)	
Intercept	0.0003	0.0010	0.326	0.7442	-0.0019	0.0018	-1.033	0.30	
Market return	0.2028	0.2282	0.889	0.3745	0.7403	0.4606	1.607	0.11	
SMB	0.0733	0.2849	0.257	0.7971	-0.3694	0.4261	-0.867	0.39	
HML	-0.0150	0.1588	-0.09	0.9247	0.3712	0.4644	0.80	0.42	
PR1YR	0.3416	0.1526	2.24	0.0255 **	-0.2880	0.2988	-0.96	0.34	
R ²	0.0108				0.0511				

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1

$$R_t^e = \alpha + \beta_1 MP_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 PR1YR_t + D_t[\delta_1 + \delta_2 MP_t + \delta_3 SMB_t + \delta_4 HML_t + \delta_5 PR1YR_t] + \epsilon_t \quad (8)$$

As a further robustness test of the excess returns, I follow the methodologies of Mill (2006) and Ang (2015) and test two augmented factor models. The first model shown in Equation 7, includes a dummy variable, which is one when markets are in crisis (downward moving) and zero otherwise. Hence, I can test whether the alphas in excess of market premium of the strategies are similar during different market conditions. The second model shown in Equation 8 is a conditional augmented market model, to test whether the factor loadings in Equation 6 are different during crisis or not.

Since the pairs are composed to a large degree of fish farmers, I use the fish price index from Fishpool¹, in addition to the Oslo All Share Index (OSEAX). The indeces captures industry specific risk (income) and market risk respectively. Defining times of market crisis or recessions is not trivial and outside the scope of this paper. For the OSEAX, I have used a ‘laymans’ definition of bear markets, in which a 20 % drawdown from the peak is considered a bear market. For the more volatile fish price index, a 30 % drop is chosen. The dummy variable is one, if either of the indeces are in a ‘crisis’ and the periods are highlighted in Fig. 4. Thus, with the augmented model in Equation 7, we are able to test whether the factor loadings of the strategies are different in bull and bear markets.

From Table 5, we can see that the non-crisis alpha is insignificant, while the crisis-alpha is significant and the coefficient is negative. Thus the strategy performs differently under different market conditions. It is also a further indication that the long returns are indeed contributing more than the short returns. The reason for this is that during normal market conditions the alpha is higher than during times of crisis where the short position would earn more as markets move downwards.

¹<http://fishpool.eu/price-information/spot-prices/fish-pool-index/>




Fig. 4. Panel A: Oslo All Share Index from January 2005 to December 2014. Panel B: Fish Pool index denominated in NOK, from January 2006 to December 2014. Highlighted is the periods with bear markets. This is defined as a 20 % and 30 % drop from peak for the OSEAX and Fish Pool index respectively. When either of the indices are highlighted the dummy variable in Equations 5 and 6 is set to 1.

The results from the conditional augmented market model is shown in Table 6. The momentum factor has the same sign as in the unconditional model, but with a greater magnitude and is now significant. This is also the case during times of crisis, now that the markets are moving in the opposite direction, it is not surprising that the factor loading also has the opposite sign. The magnitude of the factor loading is much greater during times of distress, which suggests that the returns are more sensitive to the momentum factor during crisis than in normal market conditions. Finally, the SMB factor is significant in explaining excess returns during crisis and has the same sign as in the unconditional model. Overall, using an augmented model still only explains a small portion of the pairs trading returns ($R^2 = 0.022$, not reported), similar to the unconditional model. The strategy remains market neutral in both bull and bear markets.

Distance approach: Daily data

Regarding the daily data, 84 pairs are eligible for trading, still with a maximum of five pairs per trading period. The increase in the number of pairs available is attributable to the concept that more pairs are chosen from the start of the sample compared with the case using high-frequency data. The average SSD for the pairs using daily data is 0.243, which is higher than the case using high-frequency data, which is 0.101. In this sense, the use of high-frequency data may be beneficial for identifying pairs. Out of the pairs selected in the formation period, 54 are trading, which is a nearly identical ratio as that for the high-frequency selection. Although an increase in pairs exists, the number of trades is 264, which is less than the high-frequency case. The holding period is an average of 61 days, which is approximately half the trading period. Additionally, in this case, the win/loss ratio is not particularly good, at only 0.316. To summarize the descriptive statistics in Table 2, they are for the most part very similar to the case using high-frequency data, with some differences regarding the number of pairs available and the average SSD for the pairs. One aspect that stands out is that there seems to be an increase in extreme profits and losses, where both the maximum position profit and loss has increased approximately three times when using daily data compared with high-frequency data.

Table 4. Individual leg return contribution

(a) Long returns				
	Distance 1 year	Distance Daily 1 year	Cointegration 1 year	Cointegration Daily 1 year
Estimate	-0.0001	0.0221	0.0060	0.0009
Standard error (Newey-West)	0.0067	0.0183	0.0087	0.0021
t-Statistic	-0.0166	1.2058	0.6893	0.4111
Pr(> t)	0.9878	0.3143	0.5401	0.7086

(b) Short returns				
	Distance 1 year	Distance Daily 1 year	Cointegration 1 year	Cointegration Daily 1 year
Estimate	-0.0073	-0.0162	-0.0040	-0.0058
Standard error (Newey-West)	0.0033	0.0129	0.0057	0.0050
t-Statistic	-2.2046	-1.2579	-0.7139	-1.1434
Pr(> t)	0.1147	0.2974	0.5268	0.3359

Signif. codes: '***' 0.01 '**' 0.05 '*' 0.1

This aspect is also evident when looking at the monthly returns in Table 1, where I find that for the more than 75 months in which trading occurs, the estimated average monthly return is 1.57%. To obtain the average monthly return estimate, a Newey–West t-test is performed using four lags and three degrees of freedom on the monthly return series – similar to the high-frequency case. Obviously, this result is not significant given that the SE is twice the size of the estimate. As in the case of the best and worst positions, the best and worst monthly returns are much higher in absolute value, where the best month yielded a 354% return and the worst month yielded -71.6% . This result is naturally reflected in the moments of the empirical distribution, for which the SD is very high at 34.8%. The skewness and excess kurtosis are 8.65 and 86.05, respectively, quite extreme and primarily the result of one outlier.

In the previous sections, I have identified the risk of nonconvergence as a risk factor for pairs trading. Table 2 provides an overview of whether or not pairs were convergent. When using daily data, out of the total number of positions opened, only 36.36% are closed because of a trading signal when pairs converge. Thus, the majority of the positions are kept until the trading period ends, and most likely incurs a loss. Using daily data represents an increase of almost two times in nonconvergent positions compared with when using high-frequency data. A position in a pair that is opened and kept until the end of the trading period is labelled a nonconvergent pair. A single round trip pair has exactly one convergence during the trading period. In addition, it may or may not open a second position that is closed at the end of the trading period. Multiple round trip pairs have more than one convergence and may or may not open a position that is kept to the end of the trading period. An increase in nonconvergence pairs occurs, and 72.22% of the pairs in the sample never converge once a position is opened, a 15 percentage point increase compared with the high-frequency data.

It is interesting to observe each leg individually to investigate whether the long and short positions have equal return contributions. In panel A and B of Table 4, I find that the long position generated an average profit of 2.21% and the short position generated an average loss of -1.62% . When testing whether the means are equal, I obtained a p-value of 0.101, indicating that I barely cannot reject the null hypothesis that the means are equal on a 10% significance level. However, the long position seems to generate higher returns than the short position. Therefore, when testing whether the long position is greater than the short position, I obtain a p-value of 0.0504 and are able to reject the null hypothesis that the true difference in means is greater than 0 on at least the 10% significance level. Therefore, the conclusion is that the long position makes a more significant contribution than the short position. That the stocks that I go long in are the relative losers in the pairs indicates that the losers have had a mean

Table 5. Non-crisis and crisis alpha market model from Equation 7. Standard errors using the Newey-West method with 4 lags.

	α_c	α_{nc}	β_1
Distance 1 year	-0.015*	-0.001	0.152
Distance Daily 1 year	-0.002	-0.004	-0.016
Cointegration 1 year	0.005	-0.001	0.192
Cointegration Daily 1 year	-0.008	0.001	1.366***

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1

Table 6. Conditional augmented market model from Equation 8. Standard errors using the Newey-West method with 4 lags.

	α	β_1	β_2	β_3	β_4	δ_1	δ_2	δ_3	δ_4	δ_5
Distance 1 year	0	0.079	0.049	0.094	-0.448*	-0.005	0.592	1.202*	-1.054	2.345**
Distance Daily 1 year	-0.004	0.069	0.066	0.06	0.077	0.003	0.246	0.588	-0.359	-0.538
Cointegration 1 year	-0.001	0.211	-0.107	0.201	0.167	-0.005	-2.567**	-4.749**	-1.325	-2.35
Cointegration Daily 1 year	0.002	1.116	-0.46	1.071	-0.421	0.012	-5.026**	-3.281***	-1.008	-0.819

Signif. codes: ‘***’ 0.01 ‘**’ 0.05 ‘*’ 0.1

reversal of a greater magnitude than what the ‘winner’ stocks have had.

As described in a previous section, I follow the same procedure as with the high-frequency data and perform a regression of the daily return series against the Fama and French (1996) three-factor model, and with a Carhart (1997) momentum factor. Once more, the factors explain only a small part of the variation in the return series ($R^2 = 0.006$) and that pairs trading is market neutral. The momentum factor loading is significant at 5% and has a negative sign, as expected. The estimates of the crisis and non-crisis alpha market model in Equation 7, are not significant, as shown in Table 5. Thus, there is no evidence that the daily strategy performs differently under different market regimes. Further, the estimates of the conditional augmented market model in Equation 8 is shown in Table 6. None of the coefficients are significant, and thus fails to explain factor exposure in varying market regimes, in addition to explain the strategies excess returns.

Cointegration approach: High-frequency data

In the case of a cointegration approach, there is an increase in the pairs available for trading, the 74 pairs represent a 61% increase from that of the distance approach, where the maximum number of pairs is restricted to five per period. However, when observing the number of pairs that actually trade, the ratio is quite similar regardless of the strategy chosen. In this case, 54% of the pairs selected in the formation period trade in the subsequent trading period. Compared with the distance approach, despite the higher number of pairs available, the number of trades and round trips for the cointegration approach is lower. The difference in the win–loss ratio is 4 percentage points in favour of the cointegration approach. A slight increase of seven days occurred in the holding period, with an average holding period of 62 days.

The decrease in trades is also evident in the number of months in which trading occurs, 44 months, representing a decrease of 6 months compared with the distance approach. The average return is as previously estimated using the Newey–West t-test with four lags and three degrees of freedom on the monthly return series. The average return is 0.136%, which is not significant, as is shown by the low t-statistic (0.18) in Table 1. The estimates for the return on the long and short leg are 0.6% and -0.4% , respectively. Although the long leg seems to have a higher contribution, it is not significantly greater than the contribution from the short leg (p-value = 0.201). The empirical distribution seems to be quite similar to that of the distance

approach. The SD is 6.79%, the skewness coefficient is -0.25 , and the excess kurtosis is 6.23. However, a difference exists regarding the number of months with negative returns. Only 47.73% have negative returns, which is a decrease of 8.27 percentage points compared with the distance method. The lower amount of negative returns may be the result of a decrease in nonconvergent pairs, as shown in Table 2. Using the cointegration approach, only 52.5% of the pairs that trade are nonconvergent, a decrease compared with the distance approach. On the other hand, the number of pairs that have multiple round trips decreased and the amount of positions closed at the end of the trading period is higher. Therefore, the case might be that the cointegration approach finds pairs with a more stable equilibrium compared with the distance approach. In contrast, there might be a higher degree of mean reversion in the pairs found in the distance approach, resulting in a higher rate of multiple round trip pairs and positions closed at round trip.

Once again, I am interested in testing whether pairs trading is market neutral, and whether there are significant factor loadings, as shown in Table 3. Overall, the factors explain a small portion of the variation in the daily return series of the pairs trading strategy ($R^2 = 0.011$). Only the momentum factor has a significant factor loading (t-statistic = 2.418), at the 5% level. The estimated coefficient is positive, which is the opposite sign from what I expected. Because I am selling the ‘winning’ stock short, a return continuation resulting in a further increase in the ‘winning’ stock price would incur a loss on such a position.

There is no evidence that there is difference in the alphas in times of crisis and non-crisis, in both cases the coefficients are insignificant as shown in Table 5. In the case of the conditional market model, both the MP and SMB factors during crisis, δ_2 and δ_3 respectively, have significant factor loadings, quite high in magnitude. Thus, one way of interpreting these results is that the strategy is no longer market-neutral in times of crisis. However, the negative sensitivity would imply that the returns will be higher, everything else equal, if the market crashes.

Cointegration approach: Daily data

A cointegration approach on these stocks using daily data was not very fruitful, although some insights may be had. As shown in Table 2, only 36 pairs were found to be cointegrated during the entire sample period. Out of these pairs, only 11 traded, approximately 31%. The number of trades was 52, or 13 round trips. Most of the trades resulted in a loss, which shows a very low win–loss ratio at 0.15. At the same time, the most profitable position returned 34.64%, only approximately half the performance of the other strategies’ most profitable positions. The low number of winning positions is also reflected in the average return, which is estimated at -0.53% , not significantly different from 0 (t-statistic: -1.17). The return distribution mimics that of the distance approach using daily data, albeit less extreme. The SD is 4.64%. The skewness is 1.24, thus a skewed distribution is observed only when using daily data. Excess kurtosis is also high at 17.37. As shown in Table 2, 84.62% of the positions are closed at the end of a period, and 90.91% of the pairs that are trading never converge back to the equilibrium. Once more I regress the daily return series against the previously mentioned factor indices shown in Table 3. Overall, the factors explain 5.11% of the variation in the return series, which is the highest R^2 for all cases. Also in this case, the strategy is market-neutral, although just barely, with a higher factor loading than the other strategies.

This is however not true anymore, when testing whether the alphas are different during crisis and non-crisis periods. As can be seen in Table 5, none of the alphas are significant, but the MP is, with a slightly higher factor loading than in the multifactor model. The conditional augmented market model in Table 6, shows similar results as for the high-frequency cointegration approach. Also for the daily data, the MP and SMB factors are significant in times of crisis. Thus, it seems that both cointegration approaches are sensitive to the market factor in times of crisis. Maybe this has something to do with the position sizing. The cointegration vector is used for position sizing; therefore, one explanation may be that the equilibrium between the stocks is different during times of crisis, making one leg dominant and thus exposed to the market factor.

V. Conclusion

I have tested a pairs trading strategy on a sample of Norwegian seafood company stocks during a 10-year period (2005–2014). The purpose of the study was threefold. First, this study evaluated whether pairs trading is profitable in the Norwegian equity market using a 1 year formation period and a 6 month trading period. Second, selection methods were compared, specifically the distance approach and the cointegration approach. Lastly, the use of high-frequency data was compared with the use of daily frequency data.

Regarding the performance of the strategy for a 1 year formation period, none of the strategies had profits significantly different from zero. Thus, the short answer is that pairs trading is not profitable in the Norwegian market. In general, this result is to be expected given that other studies find a decline in pairs trading profits and that such trading is not profitable on average, i.e. Do and Faff (2012). In contrast, it is known that pairs trading thrives in bear markets. Our sample includes the GFC, during which the Norwegian market plummeted and lost more than half its value, as is seen in Fig. 3. For some reason, hardly any trading occurred during this period in this study, apart from the distance method using daily data. The result also differs from studies on non-US markets that find economically and statistically significant profits, such as Perlin (2009), Broussard and Vaihekoski (2012), Mashele et al. (2013) and Li et al. (2014), which investigate markets in South America, Europe, Africa, and Asia, respectively.

Given that neither the distance nor the cointegration approach produced significant profits when implementing the strategy, it is not obvious whether one method is preferred over the other. As shown in Fig. 3, the cointegration approach is preferred when looking at cumulative profits, although it is outperformed by the market portfolio. One aspect that may favour the cointegration approach is that, relative to the distance approach, it has a lower percentage of months with negative returns. The lower amount of negative returns is likely linked to the lower percentage of pairs trading that is nonconvergent. However, these indicators are only weak, and one may not conclude one way or another on the basis of the empirics of the study.

Similarly, with respect to a comparison of high-frequency data with daily data, none of the strategies had significant profits. However, one may argue that when using the same formation and trading period for high-frequency data and daily data, the comparison is not quite apples to apples. For example, the previously mentioned studies that employ high-frequency data use a much shorter formation and trading period. One point that stands out when using daily data is the empirical distributions. For both the distance and cointegration approaches, I find in particular a high excess

kurtosis, thus highly leptokurtic distributions. In addition, for both cases with daily data the distributions are positively skewed. The high-frequency data distributions on the other hand, have a smaller excess kurtosis. In addition, both distributions are relatively symmetrical with skewness coefficients close to zero. Therefore, in this instance, the empirical distribution is closer to the normal distribution using high-frequency data than using daily data.

References

- Ang, W. R., 2015. Sustainable investment in korea does not catch a cold when the united states sneezes. *Journal of Sustainable Finance & Investment* 5 (1-2), 16–26.
- 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.
- Bowen, D., Hutchinson, M. C., 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.
- 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., Hamza, K., 2006. A new approach to modeling and estimation for pairs trading. In: *Proceedings of 2006 Financial Management Association European Conference*.
- Do, B., Faff, R. W., 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*.
- Elliott, R. J., Van Der Hoek, J., Malcolm, W. P., 2005. Pairs trading. *Quantitative Finance* 5 (3), 271–276.
- 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*.
- Engle, R. F., Granger, C. W., 1987. Co-integration and error correction: representation, estimation, and testing. *Econometrica: journal of the Econometric Society*, 251–276.
- 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., 1996. Multifactor explanations of asset pricing anomalies. *The journal of finance* 51 (1), 55–84.
- 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.
- Hong, G., Susmel, R., 2003. Pairs-trading in the asian adr market. *University of*

- Houston.–2004.
- 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.
- 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 at SSRN 1913707.
- Krauss, C., 2016. Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of Economic Surveys*.
- 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 & Hedge Funds* 19 (1), 12–30.
- 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.
- Mill, G. A., 2006. The financial performance of a socially responsible investment over time and a possible link with corporate social responsibility. *Journal of Business Ethics* 63 (2), 131–148.
- 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.
 URL 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.
- TITLON, nd.
 URL <https://titlon.uit.no/>