

NORWEGIAN SCHOOL OF ECONOMICS

Statistical Arbitrage Pairs: Can Cointegration Capture Market Neutral Profits?

Christoffer Haakon Hoel

Bergen

Spring, 2013

Supervisor: Associate Professor Jørgen Haug

Master Thesis, MSc Economics and Business Administration, ECO

This thesis was written as a part of the Master of Science in Economics and Business Administration at NHH. Please note that neither the institution nor the examiners are responsible – through the approval of this thesis – for the theories and methods used, or results and conclusions drawn in this work.

This page intentionally left blank

Abstract

We back-test a statistical arbitrage strategy, pairs trading, over the ten year period 01.01.2003 – 31.12.2012 at the Oslo Stock Exchange. We construct an unbiased dataset, where stocks are matched into pairs using a cointegration approach and traded according to a set of pre specified rules. The strategy yields consistent negative returns independent of parameterisation of entry- and exit thresholds. Our findings are in line with previous literature, where we support the view that absence of profits is not necessarily due to increased activity among hedge funds, but rather changes in fundamental factors governing the relationships between stocks.

Preface

This thesis is divided into two parts. Part one outlines the background and theoretical framework of pairs trading. In addition we conduct a Monte Carlo simulation where we show that profits to a pairs trading strategy is negatively related to the correlation between the assets in a pair. Part two is an empirical back-test applying the theory discussed under part one.

All data in this thesis may be made available upon request: chhoel@gmail.com

Acknowledgements

I am grateful to Nils Diderik Algaard from NHH Børsprosjektet for help with constructing a historical constituent list over stocks quoted on the Oslo Stock Exchange. I will also thank my supervisor Jørgen Haug for valuable input early in the process.

Contents

1. Introduction	1
1.1 Deterministic and Statistical Arbitrage.....	1
1.2 The History of Pairs Trading.....	3
1.3 Literature Review.....	3
2. The Fundamentals of Pairs Trading.....	5
2.1 The Basic Idea	5
2.2 Various Approaches to Pairs Trading.....	6
2.2.1 <i>The Distance Approach</i>	6
2.2.2 <i>The Stochastic Approach</i>	7
2.2.3 <i>The Cointegration Approach</i>	8
3. A Cointegration Approach.....	9
3.1 Stationarity	9
3.2 Cointegration.....	10
3.3 Pairs Trading and Cointegration	11
3.3.1 <i>Estimation Procedure</i>	11
3.3.2 <i>Price Series</i>	12
3.3.3 <i>Trading Thresholds</i>	14
3.3.4 <i>Interpretation of the Hedge Ratio</i>	15
4. Simulation – Correlation and Cointegration in Pairs.....	16
4.1 Model and Parameter Overview.....	17
4.2 Results.....	17
5. An Applied Pairs Trading Strategy	19
5.1 Introduction and Specifications.....	19
5.1.1 <i>Formation Period</i>	20
5.1.2 <i>Trading Period</i>	21
5.2 Results.....	23
5.2.1 <i>Unrestricted Portfolio</i>	23
5.2.2 <i>Restricted Portfolio</i>	26
5.2.3 <i>Risk Decomposition</i>	28
5.3 Discussion of Findings and Concluding Remarks.....	32
References	33
Appendices.....	35

Part One

1. Introduction

1.1 Deterministic and Statistical Arbitrage

The notion of arbitrage can perhaps be considered the Holy Grail of investing, as it is the possibility of a risk-free profit at zero cost due to mispricing of assets; construct a self-financed portfolio that has a positive probability of a positive payoff, and a zero probability of a negative payoff, for all future states in time. Such an arbitrage is often termed a *deterministic* or *pure* arbitrage, and is inconsistent with equilibrium pricing yet important for asset pricing theories such as the Arbitrage Pricing Theory (Huberman 1982). In contrast, a *statistical* arbitrage represents an opportunity in which there is a statistical relative mispricing between assets based on their expected values. A position can then be taken in order to capitalise on this relationship. However, unlike a deterministic arbitrage, such a position is not riskless. The expected payoff is positive, but so is also the probability of a negative payoff. Only when time approaches infinity and the trading strategy is continuously repeated will the probability of a negative payoff approach zero – much like a martingale betting system¹. Hogan *et al.* (2004) defines a statistical arbitrage as having the following properties:

$$v_0 = 0 \tag{1.1}$$

$$\lim_{t \rightarrow \infty} E^P[v_t] > 0 \tag{1.2}$$

$$\lim_{t \rightarrow \infty} p(v_t < 0) = 0 \tag{1.3}$$

$$\lim_{t \rightarrow \infty} \frac{\text{Var}^P(v_t)}{t} = 0 \text{ if } p(v_t < 0) > 0, \forall t < \infty \tag{1.4}$$

(1.1) it is a zero cost self-financing portfolio, (1.2) it has positive expected discounted profits and (1.3) a probability of loss converging to zero in the limit, and (1.4) a time-averaged variance converging to zero if the probability of a loss does not become zero in finite time. The fourth condition only applies if there is a positive probability of a negative outcome. Consider the case of $p(v_t < 0) = 0$ for all $t \geq T$ for some $T < \infty$. That is, the probability of a loss is zero for $t \geq T$, so that a deterministic arbitrage opportunity is available. The economic interpretation of this condition is that a statistical arbitrage opportunity will eventually return a risk free profit in the limit. In that sense, its properties will become similar to a deterministic arbitrage as time increases.

¹ E.g. in a game of Roulette a gambler would double his stake after every loss so that the first win would cover all previous losses and leave him with a profit equal to the initial stake.

Let us give an example, following that of Hogan *et al.* (2004).

Example – Assume that a trading strategy generates a profit over the time interval $[t-1, t]$ that can be written as

$$v_t - v_{t-1} = \mu + \sigma z_t$$

where $\mu > 0$, $\sigma > 0$ and z_t is i.i.d. $\sim \mathcal{N}(0, 1/t)$. For every time interval this strategy will have positive expected discounted profits with random noise; in other words, the profit will oscillate around the mean value. For simplicity assume a zero discount rate. Supposing that $v_0 = 0$ the cumulative profits at time T is

$$v_T = \sum_{t=1}^T (v_t - v_{t-1}) = \mu T + \sigma \sum_{t=1}^T z_t$$

We notice that $E[v_T] = \mu T$ and $\text{Var}(v_T) = \sigma^2 \sum_{n=1}^T 1/n$ converge to infinity as $T \rightarrow \infty$. Even so, the time averaged variance $\sigma^2 \frac{\sum_{n=1}^T 1/n}{T}$ will converge to zero as $T \rightarrow \infty$, precisely because the variance is a concave function of time. Hence, the example is a statistical arbitrage.

Although we will thoroughly cover the strategy of pair trading later, an important question is whether or not it can be considered a statistical arbitrage according to the above definition. Firstly, v_t can be thought of as a long-short portfolio consisting of two stocks whose weights can be determined so that it is a self-financed position². Secondly, it is clear that for a rational investor the expected value of an investment will be positive, if not she would not invest (assuming she is not a risk-seeker). For pairs trading, it is clear that when a position is entered into the expected profit is positive due to the mean reverting nature of which the strategy is defined. By these arguments we therefore state that requirements (1.1) and (1.2) are satisfied. Furthermore, requirements (1.3) and (1.4) are by Chiu and Wong (2012) proven to be fulfilled for a pairs trading strategy in an economy where assets are cointegrated. Because the basis of pairs trading relies on cointegration and error correctional behaviour of assets, we conclude that such a trading strategy may indeed be considered a statistical arbitrage.

In practise, statistical arbitrage is often used synonymously with the term *quantitative trading* to describe any quantitative trading strategy that searches to mitigate risk almost entirely. A less stringent layman’s definition would be to say that statistical arbitrage is the process in which one uses heavily quantified techniques seeking to profit from the relative price discrepancies between assets, where risk is believed to be so small that it is negligible. However, because the statistical relationship between two or more assets may not necessarily continue to hold into the future due to possible changes in underlying fundamental variables,

² In reality, it is seldom the case that the proceeds from the short sell can be used to cover the long position due to margin requirements.

statistical arbitrage is certainly not without risk. The 1998 bailout of Long Term Capital Management is evidence of just that³.

1.2 The History of Pairs Trading

In the early 1980's the Wall Street quant Gerry Bamberger of Morgan Stanley had the idea that it could be profitable to hedge positions within an industry group according to a set of pre specified rules (Wilmott 2005). This idea was further developed by his colleague Nunzio Tartaglia who led a team of mathematicians, physicists and computer scientists, who developed algorithms for which trades could be automatically executed (Vidyamurthy 2007). This has later become known as the black box of Morgan Stanley, and proved highly profitable in the years that came. One of the strategies the team developed was rather simple intuitively, yet intricate: find two securities whose prices seem to move together due to an underlying relationship, and when an anomaly in the relationship is noticed the pair is traded believing that the relationship will restore itself. This strategy has since been named pairs trading.

As a result of the interest in the quantitative work at Morgan Stanley and the group led by Tartaglia eventually dissolving, new hedge funds emerged. Together with increased academic interest, quantitative trading and statistical arbitrage became well known in the financial industry, and pairs trading is used extensively among institutional investors today (Pole 2007).

1.3 Literature Review

While statistical arbitrage and pairs trading have been around for over 30 years, few papers on the subject have been published in top tier academic journals. Here we give an overview of the most prominent literature.

Gatev *et al.* (2006) is perhaps the most cited paper on pairs trading. They back test a simple trading algorithm with daily data in the period 1962-2002 using S&P 500 constituents and find average annualised returns of up to 11% for portfolios of pairs. Although the proposed strategy is profitable the authors note that returns have declined in recent years, possible due to increased competition among hedge funds and/or a reduction in the importance of an underlying common factor that drives the returns in a pairs trading strategy. Furthermore, a thorough analysis of the risk characteristics shows that returns have a high risk-adjusted alpha and an insignificant exposure to sources of systematic risk.

Andrade *et al.* (2005) replicate the study of Gatev *et al.* (2006) (using their working paper from 2003) on the Taiwanese stock market from 1994 to 2002, which produces similar results with average annualised returns of 10%. Perlin (2009) tests a trading strategy much

³ LTCM was a Wall Street hedge fund using quantitative techniques to uncover statistical arbitrage opportunities in the bond and equity markets.

alike Gatev *et al.* (2006) on the Brazilian stock market using daily, weekly and monthly data, where daily data yields significantly higher returns than that of lower frequency strategies. Also, his results indicate that returns are sensitive to the parameterisation of entry and exit thresholds.

Do and Faff (2010) reproduce the paper of Gatev *et al.* (2006) with near identical results as the original paper. Expanding the study to the first half of 2008, they find that returns to the strategy continue to decline at an accelerating rate. Contrary to the general belief that increased hedge fund activity reduces profit potential, they claim that it can be attributed to changes in the nature of the “Law of One Price” as an increasing proportion of pairs do not converge upon divergence; signalling a change in underlying common factors in which the trading algorithms are formed. This has the implication that pairs of stocks historically found to be close substitutes may no longer be so in forthcoming time periods. In a recent paper Do and Faff (2012) conclude that inclusion of trading costs severely impact profits, and together with narrowed trading opportunities have rendered pairs trading largely unprofitable after 2002.

Bowen *et al.* (2010) back-test a pairs trading algorithm using intraday data over a twelve month period in 2007, and conclude that returns are highly sensitive to the speed of execution. Moreover, accounting for transaction costs and enforcing a ‘wait one period’ restriction, excess returns are completely eliminated.

Engelberg *et al.* (2009) seek to explain the nature behind pairs trading profits, and find that possibilities for profit are greatest soon after equilibrium divergence, and that the divergence is strongly related to how information disperses through the stocks that form the pair. Idiosyncratic liquidity shocks result in higher profitability than idiosyncratic news and when there is information common to both legs of the pair, profit possibilities may arise when the information is more quickly incorporated into one stock than the other.

Besides the works mentioned above, there are few papers addressing the actual performance of pairs trading. Most of the available literature is purely theoretical and deal with the underlying technicalities and modelling, not how the actual models would have performed in the long run. In this second category Vidyamurthy (2007), Lin *et al.* (2006) and Elliot *et al.* (2005) are noteworthy. While the first two give a thorough and detailed presentation of pairs trading from a cointegration viewpoint, the latter details how stochastic spread models can be useful in modelling the dynamics between assets in a pair.

We will briefly comment further on the mentioned papers in the next section.

2. The Fundamentals of Pairs Trading

2.1 The Basic Idea

The essence of pairs trading is quite simple, and builds on the premise of relative pricing. If there exists equilibrium between two assets and an anomaly is observed in the relationship, one can seek to profit from the comparative mispricing by selling the relative overvalued asset and simultaneously buying the undervalued asset. When equilibrium is again restored both positions are unwound and the investor makes a profit. This profit can naturally stem from either the long or short leg of the trade, or both.

Consider the two series of simulated stock prices depicted in figure 1.1 below. Even though they seem to follow a random walk process (with drift), they clearly share a common underlying factor thereby never drifting too far apart from each other.




Figure 1.1. Simulated stock price series.

The distance between the two stocks is referred to as the *spread*, and can be thought of as being a *synthetic asset*. The magnitude of the spread indicates the degree of relative mispricing between the stocks, thus generating buy and sell signals. As illustrated by a dummy variable in figure 1.1, positions are opened when the spread crosses a given threshold and is closed upon mean reversion. Figure 1.2 shows the modelled spread series associated with the simulated stock prices above, along with examples of entry thresholds specified by the stippled lines. How the spread can be modelled will be discussed in the next subsection.

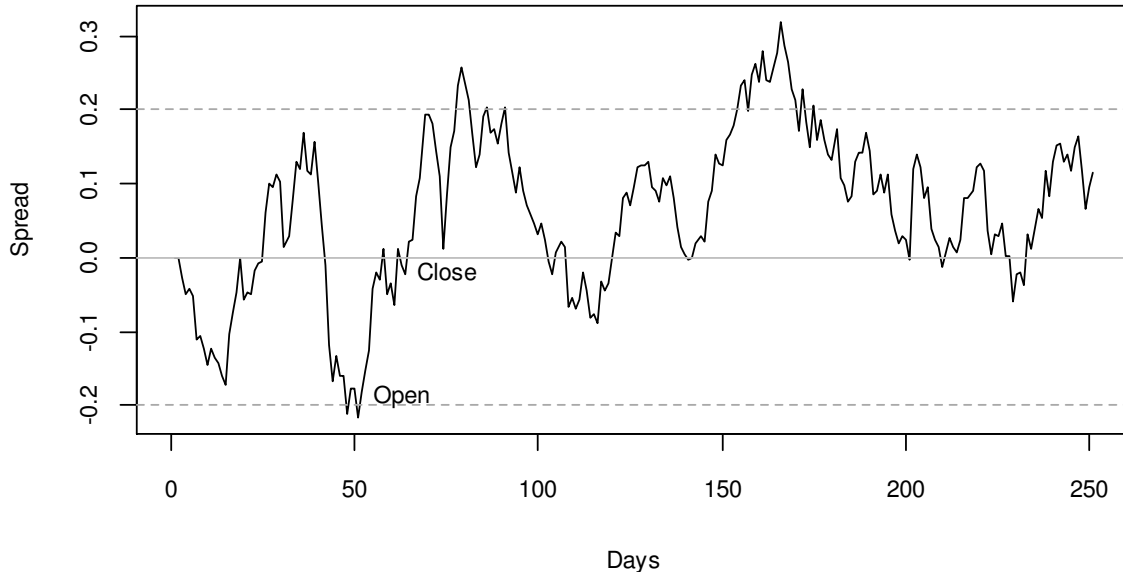


Figure 1.2. Spread series from simulated stock prices.

2.2 Various Approaches to Pairs Trading

Broadly defined, there are three⁴ different approaches to pairs trading: the *distance* approach, the *stochastic* approach and the *cointegration* approach. These methods all vary with regard to how the spread of the stock pairs is defined. Below we give a short introduction.

2.2.1 The Distance Approach

The distance approach is used among others by Gatev *et al.* (2006), Andrade *et al.* (2005), Engelberg *et al.* (2009), Perlin (2009), Do and Faff (2010, 2012) and Bowen *et al.* (2010). By this approach the *distance* between two stocks, which is the squared difference between the two normalised price series, measures the co-movement in the pair. The normalised price series for a stock is given by its cumulative total returns index, as shown in equation (1.5):

$$\tilde{P}_t = \prod_{\tau=1}^t (1 + r_\tau) \quad (1.5)$$

The normalized series begin the observation period with a value equal to one, and increases or decreases each day given its return. Stocks are matched into pairs by computing the distance (D) according to equation (1.6):

$$D^{AB} = \sum_{t=1}^T (\tilde{P}_t^A - \tilde{P}_t^B)^2 \quad (1.6)$$

⁴ A fourth approach, the *Combined Forecast* approach is suggested by Huck (2009, 2010) as the sole promoter.

When the distance measure has been computed for all stock pairs in question, one typically rank pairs based on minimum distance, where usually a certain number of pairs with the lowest value will be used for trading. The spread is simply defined as one stock price subtracted by the other, where trades are opened according to the rule in (1.7):

$$\begin{aligned} \tilde{P}_t^A - \tilde{P}_t^B \geq q &\quad \rightarrow \quad \text{short position} \\ \tilde{P}_t^A - \tilde{P}_t^B \leq -q &\quad \rightarrow \quad \text{long position} \end{aligned} \tag{1.7}$$

where q represents a threshold value.

Notably, the distance approach is a model free approach and exploits a statistical relationship among two stocks at the return level. As Do *et al.* (2006, 4) notes, it therefore has the advantage that it is not prone to model misspecification or misestimation. However, it makes the assumption that the returns of the two stocks are in parity, or equivalently that the level distance is static through time, something that may hold true for only brief periods of time and “for a certain group of pairs whose risk-return profiles are close to identical”. Additionally, because it is parameter free, it also lacks forecasting capabilities.

2.2.2 The Stochastic Approach

Papers included in this category are Elliot *et al.* (2005), Do *et al.* (2006) and Mudchanatongsuk *et al.* (2008). The common approach is outlined by Elliot *et al.* (2005) where the price difference between two assets is modelled in continuous time, and assumed to be driven by a state process and some additional measurement error:

$$y_k = x_k + D\omega_k \tag{1.8}$$

where x_k represents the value of some variable at time $t_k = k\tau$ for $k = 0, 1, 2, \dots$, ω_k is i.i.d. Gaussian $\sim \mathcal{N}(0,1)$ and $D > 0$. x_k is assumed to follow the process given by (1.9):

$$x_{k+1} - x_k = (a - bx_k)\tau + \sigma\sqrt{\tau}\varepsilon_{k+1} \tag{1.9}$$

where $\sigma \geq 0$, $b > 0$, $a \in \mathcal{R}$ and ε_k is i.i.d. Gaussian $\sim \mathcal{N}(0,1)$ and independent of ω_k from (1.8). The process described by (1.9) will mean revert around $\mu = a/b$ with “power” b . We denote $Y_k = \sigma\{y_0, y_1, \dots, y_k\}$ representing the information from observing y_0, y_1, \dots, y_k . The conditional expectation

$$\hat{x}_k = E[x_k | Y_k] \tag{1.10}$$

will be the estimate of the hidden state process of (1.9) through the observed process of (1.8). Note that (1.9) can be rewritten as:

$$x_{k+1} = A + Bx_k + C\varepsilon_{k+1} \tag{1.11}$$

where $A = a\tau \geq 0$, $0 < B = 1 - b\tau < 1$ and $C = \sigma\sqrt{\tau}$. One can regard $x_k \cong X(k\tau)$ where $X(t)|t \geq 0$ satisfies the stochastic differential equation

$$dX(t) = (a - bX(t))dt + \sigma dW(t) \tag{1.12}$$

An Ornstein–Uhlenbeck process is then used as an approximation to (1.12) in order to estimate A , B and C so that an estimate of (1.10) can be obtained. The trading dynamic is similar to that of (1.7). A trade is opened when $y_k \geq q\hat{x}_{k|k-1} = qE[x_k|Y_{k-1}]$, as the spread is considered too large: the trader takes a short position, and profits when a correction occurs. Similarly, she takes a long position if $y_k \leq -q\hat{x}_{k|k-1}$. Again, q is the threshold value for when trades are opened.

The advantages of using the stochastic approach is firstly that it captures mean reversion, the main building block of pairs trading, and secondly that it is convenient for forecasting. Specifically, expected holding period and expected return can be calculated explicitly using First Passage Time results for an Ornstein–Uhlenbeck process. Conversely, Do *et al.* (2006) argues that the model suggested by Elliot *et al.* (2005) has a fundamental issue, much like the *distance* approach, in that it restricts the long-run relationship between the securities to one of return parity. This problem may be overcome by using a transformed price series.

2.2.3 The Cointegration Approach

The cointegration approach is suggested by Lin *et al.* (2006), Vidyamurthy (2007) and Galenko *et al.* (2012). This approach uses a regression⁵ based framework to estimate the spread between two stocks as shown by equation (1.13):

$$P_t^A - \hat{\beta}P_t^B = \hat{\mu} + \hat{\varepsilon}_t \tag{1.13}$$

where $\hat{\beta}$ is the estimated coefficient from a regression of stock B on stock A, $\hat{\mu}$ is the estimated intercept⁶ and $\hat{\varepsilon}_t$ is the estimated error term, i.e. the residuals from the regression. If the spread is found to be stationary it will fluctuate around the estimated long-run equilibrium $\hat{\mu}$. Trading thresholds can then be constructed such that trades are triggered in the same way as (1.7): if $P_t^A - \hat{\beta}P_t^B - \hat{\mu} \geq q$ a short position is taken. Likewise, if $P_t^A - \hat{\beta}P_t^B - \hat{\mu} \leq -q$ a long position is taken.

The cointegration approach has its strengths in that it is a relatively simple framework where parameters are easily estimated using regression analysis, and that it explicitly models the mean reverting properties of the spread. On the other hand, Do *et al.* (2006) states that it

⁵ Note that the *Johansen* (cointegration) *test* uses VAR (vector autoregressive) models instead of regression.

⁶ A regression where the intercept is forced to equal zero is also possible.

is difficult to associate cointegration with asset pricing theories, although Vidyamurthy (2007) makes an attempt to link it to Arbitrage Pricing Theory.

*

Among academics the distance approach is the most widespread methodology, and Gatev *et al.* (2006, 803) claims that it “*best approximates the description of how traders themselves choose pairs*”. Even so, this thesis adopts a cointegration approach to pairs trading for three main reasons. Firstly, we cannot find any literature that back-test a long-term strategy based on cointegration, and it would therefore be interesting to see how its performance compares with the distance approach. Secondly, the stochastic approach seems to be little (if any at all) used in practice, and we cannot find a single paper that tests it on any actual data except simulated data. Thirdly, we will argue that cointegration is in fact, to some extent, the underlying basis for both the distance and the stochastic approach. Naturally, the pairs formed on the basis of the minimum distance criterion will most likely be cointegrated, namely because the spread oscillates about the equilibrium value. As we will later show, there is a clear resemblance between the state process in (1.9) and *error correction* models which can be deduced from cointegration. The next section details the cointegration based approach to pairs trading.

3. A Cointegration Approach

We begin this section by introducing the concepts of stationary time series and cointegration, before outlining in detail how these concepts can be used for trading pairs.

3.1 Stationarity

A stationary time series is characterised by the following properties for all t and $t - s$:

$$E(y_t) = E(y_{t-s}) = \mu \tag{1.14}$$

$$E[(y_t - \mu)^2] = E[(y_{t-s} - \mu)^2] = \sigma_y^2 \tag{1.15}$$

$$E[(y_t - \mu)(y_{t-s} - \mu)] = E[(y_{t-j} - \mu)(y_{t-j-s} - \mu)] = \gamma_s \tag{1.16}$$

where μ , σ_y^2 and γ_s are all constants. (1.14) through (1.16) therefore states that a stationary⁷ series has a constant mean, variance and autocovariance (Enders 2010). Obviously, (1.14) is the most important property in terms of pairs trading, or any other spread trading regime for that matter. If the spread between two assets are found to have a constant mean any deviations from this value can be traded against, as we illustrated in figure 1.2. (1.15) and (1.16) is perhaps of lesser importance for pairs trading, although a changing variance may affect profit potential through the magnitude of the oscillations about the mean. In

⁷ Strictly speaking, this is the definition of a *covariance-stationary* time series. However, the terminology of stationarity and covariance-stationarity is often used interchangeably.

econometrics the notion of stationarity is important, because if we want to understand the relationship between variables using regression we need to assume stability over time: by allowing the relationship between variables to change randomly in each time period, we cannot hope to learn much about how a change in one variable affects the other(s) (Wooldridge 2009).

Most non-stationary time series can be transformed into a stationary series. A common procedure is to difference the series, so that the values represent changes and not levels. A time series that becomes stationary after d times of differencing is referred to as an $I(d)$ series – integrated of order d . For instance, stock prices are often assumed to be $I(1)$ series (see e.g. Lanne (2002) and Lo (1991)).

There exist multiple statistical tests for determining whether a time series can be considered stationary, and this thesis adopts the framework of Said and Dickey (1984), namely the augmented Dickey-Fuller test (ADF-test). The ADF-test uses regression analysis in order to test for a *unit root*, i.e. non-stationarity, in an assumed underlying data generating process:

$$y_t = a_0 + a_1 y_{t-1} + a_2 y_{t-2} + \dots + a_{p-1} y_{t-p-1} + a_p y_{t-p} + a_{p+1} t + \varepsilon_t \quad (1.17)$$

(1.17) is a p th order autoregressive process: AR(p). Equation (1.17) can also be written as:

$$\Delta y_t = a_0 + \gamma y_{t-1} + a_2 t + \sum_{i=2}^p \beta_i \Delta y_{t-i+1} + \varepsilon_t \quad (1.18)$$

where

$$\gamma = -(1 - \sum_{i=1}^p a_i) \quad \text{and} \quad \beta_i = -\sum_{j=i}^p a_j$$

which is the equation used in the ADF-test⁸. Note that a_0 and/or a_2 can be set equal to zero depending on the assumptions behind the data generating process. The coefficient γ is tested with regard to the two hypotheses

$$\begin{aligned} H_0: \gamma = 0 & \rightarrow \text{non-stationary} \\ H_1: \gamma < 0 & \rightarrow \text{stationary} \end{aligned} \quad (1.19)$$

3.2 Cointegration

Introduced by Granger (1981) and further developed by Engle and Granger (1987), cointegration is the property in which two or more time series share a common stochastic trend. Consider two $I(d)$ series, x_t and y_t . It is generally true that a linear combination $z_t = x_t - ay_t$ will also be $I(d)$. Still, there is a possibility that z_t is $I(d-b)$, $b > 0$, though this is seldom the case. Now suppose $d = b = 1$, such that the two series are cointegrated: the long-run component of x_t and y_t cancels out so that z_t is stationary. The use of a constant

⁸ See Enders (2010, 215) for the transition from (1.17) to (1.18).

indicates that the relationship needs to be scaled so to attain $I(0)$ difference. Recall from equation (1.13) that if z_t is stationary it will consist of two parts

$$z_t = \mu + \varepsilon_t$$

so that it will oscillate around its equilibrium value μ . Engle and Granger (1987) suggested a two-step procedure to test for cointegration⁹. Consider two variables, x_t and y_t .

- 1) If both variables are integrated of the same order, say $I(1)$, the (possible) cointegration relationship can be estimated by a regression of the form

$$x_t = \mu + \beta y_t + \varepsilon_t$$

The residual series $\hat{\varepsilon}_t = x_t - \hat{\beta}y_t - \hat{\mu}$, previously denoted as the spread, is the estimated values of the deviations from the long-run relationship.

- 2) Test the $\hat{\varepsilon}_t$ -sequence for stationarity using the ADF-test. If the deviations are found to be stationary, x_t and y_t are cointegrated.

As noted by MacKinnon (1991) it is not possible to use the ordinary Dickey-Fuller test statistics. $\hat{\varepsilon}_t$ is generated from a regression equation and we do not know the true residual series ε_t , only its estimate. A problem arises because $\hat{\beta}$ and $\hat{\mu}$ are fitted so that they minimise the residual variance, thus making the procedure biased towards finding the most stationary relationship in the ADF equation. The test statistic used to test the magnitude of γ in (1.18) needs to reflect this – fortunately MacKinnon (1991) provides the necessary values.

3.3 Pairs Trading and Cointegration

Now that the concepts of stationarity and cointegration have been introduced, let us further detail how cointegration can be used for pairs trading.

3.3.1 Estimation Procedure

As we have seen, the notion of cointegration rests on a long-run relationship between the stochastic trends of two time series. An important issue is therefore how the possible relationship should best be estimated. Engle and Granger (1987) suggest using Ordinary Least Squares (OLS) regression, which seems to be the workhorse of choice among all the literature on cointegration based pairs trading. However, there are a few problems regarding OLS, cointegration and pairs trading. Notice the two regression equations below, where the relationship between x_t and y_t have been modelled in two separate ways:

⁹ This is known as the *Engle-Granger Two-Step Procedure*.

$$\begin{aligned}
 x_t = \hat{\mu}_0 + \hat{\beta}_0 y_t + \hat{\varepsilon}_{0t} &\quad \rightarrow \quad \hat{\varepsilon}_{0t} = x_t - \hat{\beta}_0 y_t - \hat{\mu}_0 \\
 y_t = \hat{\mu}_1 + \hat{\beta}_1 x_t + \varepsilon_{1t} &\quad \rightarrow \quad \hat{\varepsilon}_{1t} = y_t - \hat{\beta}_1 x_t - \hat{\mu}_1
 \end{aligned}
 \tag{1.20}$$

The OLS algorithm minimises the squared residuals of the dependent variable in the regression equation. This has the implication that the coefficients of the two regressions will not be the inverse of the other, i.e. $\hat{\beta}_0 \neq 1/\hat{\beta}_1$. This in turn has two effects. Firstly, cointegration analysis using OLS will be sensitive to the ordering of variables. It is a possibility that one of the relationships in (1.20) will be cointegrated, while the other will not. This is troublesome because we would expect that if the variables are truly cointegrated the two equations will yield the same conclusion. Secondly, the unsymmetrical coefficients imply that a hedge of long x_t / short y_t is not the opposite of long y_t / short x_t , i.e. the hedge ratios are inconsistent. Along with Teetor (2011) and Gregory *et al.* (2011) we propose that a better approach will be to use *orthogonal regression* – also referred to as *Total Least Squares* (TLS), *deming* and *errors-in-variables* (EIV) regression – in which the residuals of both dependent and independent variables are taken into account. That way, we incorporate the volatility of both legs of the spread when estimating the relationship so that hedge ratios are consistent, and thus the cointegration estimates will be unaffected by the ordering of variables. Appendix 1.1 illustrates the difference between OLS and orthogonal regression.

3.3.2 Price Series

Cointegration tests can be applied to both untransformed and transformed price series. A straightforward approach is to simply use the raw price series for a set of assets to test for cointegration between pairs. Then again, Do *et al.* (2006) notes that the long-term level difference of two stocks should not be constant except when they trade at similar price points; rather, it should increase as they go up and decrease as they go down¹⁰. A simple transformation of the price series by taking the natural logarithm overcomes this problem. To see this, define the spread between the level prices of two stocks as

$$s_t^{level} = P_t^A - \hat{\beta}P_t^B = \hat{\mu} + \hat{\varepsilon}_t$$

The prices at time $t + 1$ can be expressed as

$$P_{t+1}^A = P_t^A(1 + r_{t+1}^A) \quad P_{t+1}^B = P_t^B(1 + r_{t+1}^B)$$

where r^i , $i \in \{A, B\}$ is the discrete return. The spread at time $t + 1$ then becomes

$$s_{t+1}^{level} = P_t^A(1 + r_{t+1}^A) - \hat{\beta}P_t^B(1 + r_{t+1}^B)$$

¹⁰ Ref. the previous discussions relating to return parity with regard to the distance approach and the stochastic approach.

$$\begin{aligned}
 &= P_t^A - \hat{\beta}P_t^B + P_t^A r_{t+1}^A - \hat{\beta}P_t^B r_{t+1}^B \\
 &= s_t^{level} + P_t^A r_{t+1}^A - \hat{\beta}P_t^B r_{t+1}^B
 \end{aligned} \tag{1.21}$$

so that $s_{t+1}^{level} = s_t^{level}$ iff. $P_t^A r_{t+1}^A = \hat{\beta}P_t^B r_{t+1}^B$. Imagine that $r_{t+1}^A = r_{t+1}^B = r_{t+1}$. We now write

$$\begin{aligned}
 s_{t+1}^{level} &= s_t^{level} + (P_t^A - \hat{\beta}P_t^B)r_{t+1} \\
 &= s_t^{level}(1 + r_{t+1})
 \end{aligned}$$

indicating that the spread value will not be constant, but widens/narrows as prices increase/decrease. Rewriting (1.21) by forcing equality and substituting for P_t^A we see that

$$\begin{aligned}
 0 &= P_t^A r_{t+1}^A - \hat{\beta}P_t^B r_{t+1}^B \\
 &= (\hat{\mu} + \hat{\beta}P_t^B + \hat{\varepsilon}_t)r_{t+1}^A - \hat{\beta}P_t^B r_{t+1}^B \\
 &\iff \\
 r_{t+1}^B &= \frac{\hat{\mu} + \hat{\beta}P_t^B + \hat{\varepsilon}_t}{\hat{\beta}P_t^B} r_{t+1}^A \\
 &= \left(1 + \frac{\hat{\mu} + \hat{\varepsilon}_t}{\hat{\beta}P_t^B}\right) r_{t+1}^A
 \end{aligned} \tag{1.22}$$

there is a specific relationship between the individual returns that is required if the long-term level distance between the two assets are to be constant.

Now suppose that the spread is defined as the logarithm of prices and that at time $t + 1$ prices can be expressed as

$$P_{t+1}^A = P_t^A e^{r_{t+1}^A} \quad P_{t+1}^B = P_t^B e^{r_{t+1}^B}$$

where r^i , $i \in \{A, B\}$ is the continuous return.

The spread at time $t + 1$ now becomes

$$\begin{aligned}
 s_{t+1}^{log} &= \ln(P_{t+1}^A) - \hat{\beta} \ln(P_{t+1}^B) \\
 &= \ln(P_t^A) - \hat{\beta} \ln(P_t^B) + r_{t+1}^A - \hat{\beta} r_{t+1}^B \\
 &= s_t^{log} + r_{t+1}^A - \hat{\beta} r_{t+1}^B
 \end{aligned} \tag{1.23}$$

so that $s_{t+1}^{log} = s_t^{log}$ iff. $r_{t+1}^A = \hat{\beta} r_{t+1}^B$, and the spread will be independent of the price levels.

Log-transformation of prices is the approach used in part two of this thesis.

3.3.3 Trading Thresholds

Logically, the construction of trading thresholds is crucial to the performance of a pairs trading strategy, as it dictates when positions are both entered into and unwound. Entry-thresholds decide when trades are triggered, and exit-thresholds decide when trades are unwound. For entry-thresholds there is generally a trade-off between profits per trade and the number of trades. *Ceteris paribus*, a high threshold will certainly yield higher profits per trade than a lower threshold because the purchase, or sell, of the synthetic asset occurs farther away from equilibrium than if the threshold had been set lower. Conversely, a low threshold will yield a higher number of trades, simply because there is an increased probability that the spread will hit the trigger value. Likewise, the farther away the exit-threshold is from the trigger value the higher the profit potential, but the number of trades will be lower as the probability of exiting a position decreases.

The threshold can be constructed in a variety of ways where the most common method seems to be a static measure based on the historical standard deviation of the spread:

$$q = \delta \sigma_{spread} \quad (1.24)$$

Gatev *et al.* (2006), Andrade *et al.* (2005) and Do and Faff (2010) set $\delta = 2$, whereas Perlin (2009) and Bowen *et al.* (2010) experiment with a range of values. It is also possible to let q be a variable by defining σ as a rolling parameter with window size n ; this may allow us to better capture the profit potential of periods with higher volatility in the spread. In part two of this thesis we will experiment with multiple estimates for both δ and σ . Appendix 1.2 illustrates how various values for q impact number of trades and holding time.

Lin *et al.* (2006) suggest a cointegration coefficient weighting rule (CCW), and show how the threshold values for entry and exit determines profit per trade. Let us assume that two stocks, A and B, are cointegrated with the following relationship:

$$P_t^A - \hat{\beta} P_t^B = \hat{\mu} + \hat{\varepsilon}_t$$

Where P_t^A is the raw price of stock A at time t and $\hat{\varepsilon}_t$ is an $I(0)$ series. Now suppose that $\hat{\varepsilon}_t \geq q$, i.e. A is overvalued while B is undervalued, so a trade is opened. We sell one unit of A at price P_t^A and buy $\hat{\beta}$ units of B at price P_t^B . The position is unwound upon mean reversion by buying back one unit of A at price P_{t+n}^A and selling $\hat{\beta}$ units of B at price P_{t+n}^B . The minimum profit at time $t + n$ can then be expressed as:

$$P_{t+n} = \hat{\beta}(P_{t+n}^B - P_t^B) + (P_t^A - P_{t+n}^A)$$

By substituting for P^A we can write:

$$\begin{aligned}
 P_{t+n} &= \hat{\beta}(P_{t+n}^B - P_t^B) + [(\mu + \hat{\beta}P_t^B + \hat{\varepsilon}_t) - (\mu + \hat{\beta}P_{t+n}^B + \hat{\varepsilon}_{t+n})] \\
 P_{t+n} &= (\hat{\varepsilon}_t - \hat{\varepsilon}_{t+n}) \geq q
 \end{aligned}
 \tag{1.25}$$

So by trading the number of shares equal to the cointegration coefficient, the profit per trade will be at least q . The derivation of minimum profit for a lower trade when $\hat{\varepsilon}_t \leq -q$ is analogous to the above. Lin *et al.* (2006) considers cointegration using the raw price data. By using log-transformed data the expression in (1.25) is interpreted differently: instead of minimum profit per trade it now yields a “return-like” expression:

$$\begin{aligned}
 \overline{R}_{t+n} &= \hat{\beta}[\ln(P_{t+n}^B) - \ln(P_t^B)] + [\ln(P_t^A) - \ln(P_{t+n}^A)] \\
 \overline{R}_{t+n} &= (\hat{\varepsilon}_t - \hat{\varepsilon}_{t+n}) \geq q
 \end{aligned}
 \tag{1.26}$$

Vidyamurthy (2007, 81) claims that (1.26) is the return to a long-short portfolio consisting of short one share of stock A and long $\hat{\beta}$ shares of stock B – this is clearly wrong because the individual returns are not proportionally weighted. However, the expression is useful when filtering possible pairs with respect to bid-ask spreads, as we will see in part two.

3.3.4 Interpretation of the Hedge Ratio

The proportion of shares bought to shares sold may vary depending on investor preference. Gatev *et al.* (2006) and papers following their approach construct *capital neutral* portfolios, by using the proceeds from the short sell to invest in the long leg of the spread. At the time of investment the trader is therefore unexposed with a zero value portfolio, though such a position is seldom achievable due to margin requirements. Lin *et al.* (2006) construct *market neutral* portfolios by using the cointegration coefficient as the hedge ratio, so that exposure to systematic risk is mitigated. The interpretation of the coefficient in terms of a hedge ratio will vary depending on the price series used in the cointegration relationship. Certainly, if one uses raw prices the coefficient $\hat{\beta}$ can simply be defined as the number of shares to go long or short. On the other hand, if the prices are log-transformed prior to estimation the coefficient cannot be interpreted as the number of shares in the hedge, but rather it should be viewed as the relative weight in capital. In a log-log regression the coefficients are interpreted as: *a one per cent increase in independent variable x transmits to a $\hat{\beta}$ per cent increase in dependent variable y*. Or put differently, the coefficient is the estimated elasticity of y with respect to x (Wooldridge 2009). Let us give a simple example.

Example – Assume we have estimated the following relationship:

$$\ln(P_t^A) = \hat{\mu} + 1.50\ln(P_t^B) + \hat{\varepsilon}_t$$

A price increase of 10% in stock B will result in a price increase of approximately 15% in stock A. Let $P_t^A = 12$ and $P_t^B = 6$. If the hedge ratio had been interpreted as the number of

shares, the weights would be $\omega_A = P_t^A / (P_t^A + \hat{\beta}P_t^B) = 0.57$ and $\omega_B = \hat{\beta}P_t^B / (P_t^A + \hat{\beta}P_t^B) = 0.43$ and our long-short portfolio would yield an expected return of 4.3%. If we instead define the coefficient as the relative weight in capital, i.e.

$$\omega_A = P_t^A / (P_t^A + \hat{\beta}P_t^A) = 0.40 \text{ and } \omega_B = \hat{\beta}P_t^A / (P_t^A + \hat{\beta}P_t^A) = 0.60 \quad (1.27)$$

the long-short portfolio would make an expected return of 0.0%. This clearly illustrates that when working with log-transformed prices series the cointegration coefficients must be seen as a relative capital weight.

We would like to point out that if two stocks are negatively correlated, they will have a negative hedge ratio should they be cointegrated. If that is the case, the relationship becomes $P_t^A + \hat{\beta}P_t^B = \hat{\mu} + \hat{\varepsilon}_t$, so that the same position is taken in both stocks, i.e. both stocks are either bought or sold together. Still, this is seldom the case.

4. Simulation – Correlation and Cointegration in Pairs

In this section we illustrate a pairs trading example using simulated data under two conditions: 1) prices are cointegrated but returns are uncorrelated, and 2) prices are cointegrated *and* returns are correlated. Pairs of cointegrated stocks are simulated using an error correction model of the form:

$$\begin{aligned} \Delta x_t &= \alpha_x (x_{t-1} - \hat{\beta}y_{t-1}) + \varepsilon_{x_t} \\ \Delta y_t &= \alpha_y (x_{t-1} - \hat{\beta}y_{t-1}) + \varepsilon_{y_t} \end{aligned} \quad (1.28)$$

Where the correction factors $\alpha_x > 0$ and $\alpha_y < 0$, $\hat{\beta}$ is a cointegration coefficient, and ε_{x_t} and ε_{y_t} are white-noise error terms $\sim \mathcal{N}(\mu, \sigma^2)$. Granger¹¹ proves that for any pair of $I(1)$ variables, cointegration is equivalent to an error correction model such as (1.28). As a side note, notice the similarity between the state process in (1.9) and the error correction model: both contain correction factors working to adjust the spread should it not be in equilibrium.

The correlation between returns is modelled in the following way:

$$\begin{aligned} \varepsilon_{x_t} &= \mu + \sigma z_{1_t} \\ \varepsilon_{y_t} &= \mu + \sigma [z_{1_t}\rho + z_{2_t}(1 - \rho^2)^{0.5}] \end{aligned}$$

where z_{1_t} and z_{2_t} are $\sim \mathcal{N}(0, 1)$ and ρ is the correlation coefficient.

¹¹ In Engle and Granger (1987) – *The Granger Representation Theorem*.

4.1 Model and Parameter Overview

We conduct 500 simulations for each value of ρ , with 250 observations for each series. We then use the last 125 observations, approximately six months of daily prices, for trading in order to avoid any possible bias resulting from the model being in equilibrium at $t = 0$. Simulation is done using natural logarithms to circumvent the possibility of negative stock prices; the cointegration coefficient must be interpreted thereafter. Trading thresholds are set to two times the historical standard deviation of the spread from the first 125 observations, with positions unwound at mean reversion. Positions are opened when the spread crosses *down/up* through the threshold towards equilibrium, and all open trades are forced closed at the last trading day, possibly with a loss. In addition, all profits are reinvested during the trading period. The parameter values used during the simulation experiments are:

$$\begin{array}{cccccc} \alpha_x = 0.05 & x_0 = 15 & \mu = 0 & \hat{\beta} = 1.18 & & \rho = \begin{cases} 0.0 \\ 0.4 \\ 0.6 \\ 0.8 \end{cases} \\ \alpha_y = -0.05 & y_0 = 10 & \sigma = 0.04 & & & \end{array}$$

4.2 Results

Table 1.1 below presents the results from the simulations. There seems to be a clear link between the return to a pair and the correlation between their individual returns: the higher the correlation the lower the returns. The number of roundtrips, that is the number of times a position is entered and subsequently exited, appears to be fairly independent of correlation. We also note that the standard deviation of the spread decreases as the correlation increases.

<i>Statistic</i>	$\rho_{0.0}$	$\rho_{0.4}$	$\rho_{0.6}$	$\rho_{0.8}$
Average return	22.49 %	16.41 %	13.41 %	9.60 %
<i>Maximum</i>	73.76 %	48.46 %	44.08 %	31.72 %
<i>Minimum</i>	-5.72 %	-0.37 %	-4.08 %	-1.75 %
Average # roundtrips	2.29	2.16	2.17	2.16
Average # holding days	25.67	24.76	25.75	26.10
Avg. holding time per trade	10.88	11.13	11.64	11.92
% of pairs not open	11.20 %	12.80 %	9.60 %	10.80 %
Average return per trade ¹²	9.25 %	7.29 %	5.98 %	4.33 %
Avg. correlation	-0.0045	0.3978	0.5942	0.7987
Avg. historical std.dev. ¹³	0.1225	0.0960	0.0794	0.0564
<i>Avg. trading threshold</i>	0.2450	0.1919	0.1589	0.1129

Table 1.1. Simulation results.

¹² Computed as compounded return per trade: $(1 + r)^{(1/\#roundtrips)} - 1$.

¹³ The average of historical standard deviations of the spread, i.e. its first 125 observations.

Tests on the significance of the difference in means are presented in table 1.2. The relationship between return and correlation is statistically significant given any significance level – this is to be expected precisely as the standard deviation of the spread decreases with increased correlation, in that way reducing the magnitude of the mispricing and thus the profit potential. Furthermore, we see a tendency towards fewer roundtrips when the individual asset returns are correlated compared to uncorrelated. However, the degree of co-movement seems to be of little importance. Lastly, the average holding time per trade seems unaffected by correlation, although there is a significant difference between zero correlation and a correlation of 0.8. Even though cointegration does not necessarily imply correlation, in practise the vast majority of cointegrated pairs will also have high correlated returns. The results from this section indicate that traders searching for pairs using a correlation measure¹⁴ would instead be better off by focusing on cointegration and a *low* correlation.

<i>Variable</i>		$\rho_{0,0/0,4}$	$\rho_{0,0/0,6}$	$\rho_{0,0/0,8}$	$\rho_{0,4/0,6}$	$\rho_{0,4/0,8}$	$\rho_{0,6/0,8}$
Avg. return	t-stat	7,18	11,56	17,57	4,73	11,97	8,08
	<i>p-value</i>	0,00	0,00	0,00	0,00	0,00	0,00
# roundtrips	t-stat	1,44	1,39	1,42	-0,07	-0,04	0,02
	<i>p-value</i>	0,07	0,08	0,08	0,47	0,48	0,49
Avg. hold time per trade	t-stat	-0,46	-1,42	-1,81	-0,91	-1,33	-0,49
	<i>p-value</i>	0,32	0,08	0,04	0,18	0,09	0,31
Std. of spread	t-stat	20,79	35,73	59,28	6,76	45,30	30,06
	<i>p-value</i>	0,00	0,00	0,00	0,00	0,00	0,00

Table 1.2. One-tailed two-sample T-tests with assumed unequal variance

¹⁴ A quick Google search for “pairs trading correlation” shows that correlation is often used as a measure to identify possible pairs among practitioners.

Part two

Part two of this thesis back-tests a pairs trading strategy using a cointegration approach. We first sketch out the details, before presenting the results. Lastly, we discuss findings and compare our results with the aforementioned literature.

5. An Applied Pairs Trading Strategy

5.1 Introduction and Specifications

We test a pairs trading strategy on the Oslo Stock Exchange (OSE) over the ten year period 01.01.2003 – 31.12.2012, defining the space of available assets as all listed equities¹⁵. All stock price data is gathered from NHH Børsprosjektet through the Amadeus database, and is adjusted for both dividends and splits. The empirical studies mentioned in the literature review all use a formation period of one year followed by a trading period of six months¹⁶. Even so, seeing how we wish to ensure a long and stable cointegration relationship between pairs, we use a formation period twice as long; pairs are matched over a formation period of two years before being traded the next consecutive six months. Figure 2.1 illustrates the overlapping periods. The use of separate formation and trading periods ensures proper in- and out-of-sample data for the back-test, so that our results are not biased in terms of data snooping or survivorship. Stocks that are delisted in a trading period will still be included in both the formation and trading period; it is crucial that we behave as if we do not have any *a posteriori* information.

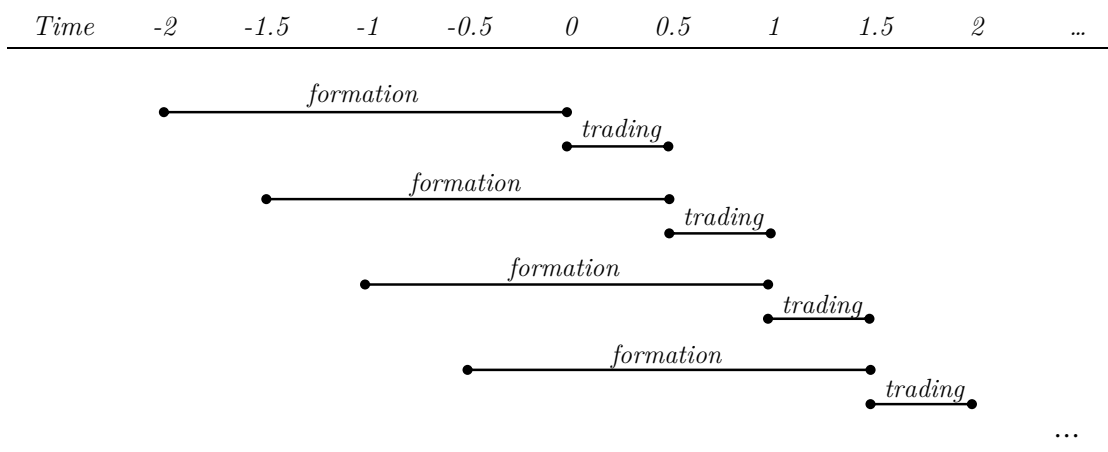


Figure 2.1. Illustration of overlapping formation- and trading periods.

¹⁵ Excluding equity certificates.

¹⁶ Except Bowen *et al.* (2010) who use intraday data.

We employ the cointegration framework presented in part one of the thesis. The relationship between two stocks is estimated by regression equation (2.1), before testing the residuals using the ADF-test with the test equation given by (2.2).

$$\ln(P_t^{A,m}) = \hat{\mu} + \hat{\beta}\ln(P_t^{B,m}) + \hat{\varepsilon}_t \quad (2.1)$$

$$\Delta\hat{\varepsilon}_t = a_1\hat{\varepsilon}_{t-1} + a_2\Delta\hat{\varepsilon}_{t-1} + e_t \quad (2.2)$$

Cointegration is modelled using the logarithm of the closing midprice, i.e. the average of closing bid and ask prices, whereas during the trading period we will use the actual bid-ask prices, so as to account for transaction costs (for simplicity we do not account for commission fees, which naturally would reduce profits). The lag length in the ADF-equation has been set to one. Although higher lag lengths could have been used to identify stationary series that would otherwise have been rejected at lower lag lengths, we only consider cointegration relationships that are “strong enough” to yield stationary residuals with lag length one. Also notice that in the test equation of (2.2) we have omitted both an intercept and a time trend. Because the residual series is from a regression and conversely should be stationary, there is no economic meaning in including these terms.

5.1.1 Formation Period

The formation period consists of a rigorous regime where we test for cointegration among all available assets. By considering all equities listed on OSE, this translates to 15.400 possible pairs for each period given an average number of listed stocks of 176 throughout the sample. We do some additional filtering before testing for cointegration by only considering liquid stocks. Specifically, we filter out stocks that have had one or more days where there has *not* been available *both* a bid and ask order at the end of the trading day. We therefore assume that we can trade at the closing bid-ask prices each day. After removing illiquid stocks we continue by testing for cointegration by following the steps below:

- 1.1 Test for cointegration using two years of historical data. Test both equations of (1.20). Discard pairs not cointegrated using both relations at the 5%-level.¹⁷
- 1.2 Test for cointegration using one year of historical data. Test both equations of (1.20). Discard pairs not cointegrated using both relations at the 5%-level.
- 1.3 Discard pairs that are not cointegrated in both step 1.1 and 1.2.
- 2.1 Filter out pairs based on transaction costs. Recall the “return like” expression in (1.26), where the magnitude of the spread could be interpreted as a measure of return to a pair. As suggested by Vidyamurthy (2007) we create a rule stating that for a given pair the following must hold true:

¹⁷ Note that by using orthogonal regression this will not be a problem. However, it is a good idea to estimate the parameters for later use in step 2.1 and 2.2.

$$\sigma > \theta_A + |\hat{\beta}|\theta_B \quad (2.3)$$

where σ is the standard deviation of the two year cointegration spread, θ_A is the average bid-ask spread of stock A over the last two years and $\hat{\beta}$ is the cointegration coefficient. Put differently, a trading threshold of one standard deviation must yield a positive return after average trading slippage due to bid-ask prices. Discard pairs that do not satisfy equation (2.3).

- 2.2 Filter out pairs based on trading possibilities by discarding the relation in equation (1.20) that has the *weakest* trading possibility, i.e. the lowest value of $\sigma - (\theta_A + |\hat{\beta}|\theta_B)$.
- 3.1 For all remaining pairs, assess historical cointegration coefficients. Compute the mean of the 1 year rolling coefficients updated every five days. Discard pairs not cointegrated over the last one year using the average coefficient at the 10%-level.

We propose that step 3.1 serves as an adequate test for the historical stability of the cointegration coefficient. A stable coefficient is important so that there is a higher chance that the estimated parameters will hold during the trading period. An R code to the complete estimation routine above is given in Appendix 2.1.

5.1.2 Trading Period

Pairs that survive the testing procedure of the formation period will be used for trading. The spread is constructed as the residual series of equation (2.1) with a one year cointegration coefficient, with the CCW-rule discussed in part 1 used as the relative weights in each stock. Trades are triggered in the same way as equation (1.7), with the exception that positions are opened when the spread crosses down/up towards equilibrium, i.e. the second crossing. This might help alleviate risk as the spread is believed to be on a path towards equilibrium. We also pose the restriction that the spread must be at least larger/smaller than φ in order to cover transaction costs, where φ is set so the magnitude of the spread is at least one standard deviation.

$$\begin{aligned} \hat{\varepsilon}_{t-1} > q & \rightarrow \text{short position} \\ \text{and } \varphi \leq \hat{\varepsilon}_t \leq q & \\ \hat{\varepsilon}_{t-1} < q & \rightarrow \text{long position} \\ \text{and } -\varphi \geq \hat{\varepsilon}_t \geq -q & \end{aligned} \quad (2.4)$$

The threshold is defined as before: $q = \delta\sigma_{spread}$. We will vary the value of q according to table 2.1 in order to test for sensitivity of parameters. Values for entry- and exit points are shown for upper-trades only – inverse values will be used for lower-trades. The parameters for the

δ	<i>entry</i>	1,5	1,5	2	2	2,5	2,5
	<i>exit</i>	0	-1,5	0	-2	0	-2,5
φ		1.0	0	1.0	0	1.0	0
σ	Constant	6M	4M	2M			
	(125 days)	(125 days)	(80 days)	(40 days)			

Table 2.1. Values for entry- and exit points, and window size of standard deviation calculations.

standard deviation indicate the size of the rolling window, whereas *constant* refers to a constant standard deviation defined as the 125 last days of the formation period. Pairs may open multiple times during trading. At the end of the period, all open pairs are forced closed, possibly with a loss. Additionally, we place a restriction on the timeframe for when a trade can be opened. Denote n_{avg} the average holding time per trade for the last 125 days of in-sample trading during the formation period, and $n_{min} = \max(n_{avg}, 14)$. If the trading period has n remaining days, a position can only be opened if $n \geq n_{min}$. For simplicity we set no restrictions on short sales, although this may be unrealistic for some stocks and/or time periods¹⁸, and assume that fractional shareholding is possible.

A *stop/loss*-rule is also incorporated, stating that any position is unwound if a loss of 25% or greater occurs. This is done so that we may exit an unprofitable position that is believed to continue to generate losses, that is, the cointegration relationship has broken down so the spread will not revert to equilibrium but instead keep widening. If a trade is stopped the pair cannot be reopened during the remainder of the trading period. Return calculations are done according to equation (2.5) following Hong and Susmel (2004). Consider a long position in stock A and a short position in stock B:

$$\frac{(P_{t+1}^A - P_t^A) + (P_t^B - P_{t+1}^B)}{P_t^A + \tau P_t^B} = \frac{\text{profit}}{\text{capital exposed}} \quad (2.5)$$

For simplicity we have dropped the notation of bid-ask prices. The variable τ represents a scale of the capital needed to trade on a margin account. In Norway, this requirement may vary daily from asset to asset based on its volatility. We therefore set $\tau = 1$ and compute a more conservative, but consistent return on all pairs – the *return on overall capital exposed*. Note that when $\tau = 1$ equation (2.5) can be written in the familiar sense of

$$\begin{aligned} \frac{(P_{t+1}^A - P_t^A) + (P_t^B - P_{t+1}^B)}{P_t^A + P_t^B} &= \left[\frac{P_t^A (P_{t+1}^A - P_t^A)}{P_t^A (P_t^A + P_t^B)} \right] + \left[\frac{P_t^B (P_t^B - P_{t+1}^B)}{P_t^B (P_t^A + P_t^B)} \right] \\ &= \left[\frac{P_{t+1}^A - P_t^A}{P_t^A} \cdot \frac{P_t^A}{P_t^A + P_t^B} \right] + \left[\frac{P_t^B - P_{t+1}^B}{P_t^B} \cdot \frac{P_t^B}{P_t^A + P_t^B} \right] \end{aligned} \quad (2.6)$$

¹⁸ During the late financial crisis some financial stocks were restricted from short selling.

which is simply the weighted return of each leg of the pair. The return to the overall portfolio is computed as a simple mean of the return to each individual pair, thereby assuming that all pairs are equally weighted. Positions are marked-to-market daily, with all profits being reinvested during the trading period. The strategy may therefore be interpreted as a buy-and-hold strategy in terms of returns.

Corrections are made for any missing data for each stock in a pair during the trading period. For example, if the ask price of stock A is missing, it is estimated using the corresponding bid price together with the average bid-ask-spread from the formation period. Correcting for missing data is only done so that we can always have an estimate of the midprice. If both the bid and ask price is missing, the midprice is assumed unchanged from the previous day. Obviously, we do no trade using estimated prices, but *actual* prices. If the ask price of stock A is missing, we cannot enter a long position in stock A (and consequently we will not enter a short position in stock B) and the trade is put on hold until a price is available (and the spread still signals a trade).

5.2 Results

5.2.1 Unrestricted Portfolio

Below we present the results from a portfolio where pairs are allowed to be formed both inside and outside of industry sectors. Every trading period consists of the 20 pairs that were found to yield the most stationary spread series from step 3.1 in the preceding formation period, i.e. the pairs with the lowest ADF test-statistic using the mean of the one year rolling coefficients. However, in order to maintain a certain degree of diversification, we restrict all stocks to only be included a total of three times in each leg of a pair. Table 2.2 gives the statistics (see Appendix 2.2 for a short description). Note that panel a, column three, is the parameterisation used by Gatev *et al.* (2006) which we call the *unrestricted baseline* case.

Even though the strategy performs exceptionally well in-sample, it makes significant losses out-of-sample with highly negative mean returns. This occurs not because of high transaction costs, but because the estimated cointegration relationships break down rather quickly out-of-sample, as is shown by the descriptive statistics. Firstly, the low number of trades per pair indicates a weak mean reverting behaviour of the estimated spreads, as the number of relevant zero-crossings will be low. Secondly, we see a low percentage of completed trades, signalling that the relationships break down and the spreads wander away from the estimated equilibrium. Also notice that the parameter values for entry and exit thresholds seem to have little impact on the profitability. Figure 2.2 shows monthly cumulative return of the unrestricted baseline case together with an index of Oslo Stock Exchange¹⁹, as well as a Kernel density estimate.

¹⁹ See Appendix 2.3 for a short description of pricing factors.

δ	entry		1.5		2		2.5	
	exit	0	-1.5	0	-2	0	-2.5	
Panel a): constant standard deviation								
Total return		-51.9 %	-50.7 %	-48.4 %	-41.6 %	-55.8 %	-46.5 %	
Monthly returns	<i>Average</i>	-0.60 %	-0.56 %	-0.54 %	-0.44 %	-0.67 %	-0.51 %	
	<i>Std. deviation</i>	1.40 %	2.20 %	1.49 %	1.47 %	1.37 %	1.40 %	
	<i>Skewness</i>	-0.144	3.108	0.109	0.305	-0.150	0.650	
	<i>Kurtosis</i>	1.116	22.071	0.139	0.774	0.255	1.481	
	<i>Minimum</i>	-5.25%	-5.64%	-4.76%	-4.94%	-5.17%	-3.83%	
	<i>Maximum</i>	3.86%	15.31%	3.00%	3.69%	2.62%	4.72%	
Annualised SR		-2.13	-1.30	-1.86	-1.64	-2.36	-1.90	
Avg. % open pairs		83.67 %	83.42 %	81.36 %	69.83 %	78.44 %	68.19 %	
Avg. # trades per pair		1.61	1.37	1.48	1.20	1.33	1.13	
Avg. % completed trades		62.05 %	44.85 %	55.40 %	34.63 %	44.40 %	21.54 %	
Panel b): rolling 6M standard deviation								
Total return		-54.4 %	-47.0 %	-61.5 %	-34.1 %	-61.9 %	-53.5 %	
Monthly returns	<i>Average</i>	-0.64 %	-0.50 %	-0.78 %	-0.33 %	-0.79 %	-0.62 %	
	<i>Std. deviation</i>	1.57 %	2.19 %	1.69 %	1.98 %	1.71 %	1.62 %	
	<i>Skewness</i>	0.334	2.331	0.061	2.400	0.165	0.110	
	<i>Kurtosis</i>	0.476	12.881	0.576	12.327	0.136	0.593	
	<i>Minimum</i>	-3.93%	-4.90%	-5.99%	-4.40%	-5.00%	-5.24%	
	<i>Maximum</i>	4.82%	13.42%	3.94%	11.95%	3.87%	4.01%	
Annualised SR		-1.98	-1.21	-2.12	-1.03	-2.12	-1.88	
Avg. % open pairs		88.2 %	89.2 %	92.2 %	76.7 %	88.9 %	75.1 %	
Avg. # trades per pair		1.63	1.38	1.41	1.14	1.21	1.07	
Avg. % completed trades		57.6 %	40.9 %	48.2 %	26.2 %	37.3 %	10.1 %	
Panel c): rolling 4M standard deviation								
Total return		-60.47 %	-58.56 %	-63.05 %	-55.66 %	-61.60 %	-49.21 %	
Monthly returns	<i>Average</i>	-0.76 %	-0.71 %	-0.81 %	-0.66 %	-0.78 %	-0.54 %	
	<i>Std. deviation</i>	1.50 %	1.87 %	1.84 %	2.04 %	1.74 %	2.07 %	
	<i>Skewness</i>	-0.071	0.541	-0.170	1.348	-0.085	1.249	
	<i>Kurtosis</i>	-0.237	1.696	0.943	6.917	0.592	6.297	
	<i>Minimum</i>	-4.73%	-4.07%	-7.28%	-6.30%	-6.71%	-5.55%	
	<i>Maximum</i>	2.91%	7.18%	4.17%	10.33%	3.71%	10.44%	
Annualised SR		-2.35	-1.81	-2.02	-1.56	-2.07	-1.35	
Avg. % open pairs		88.72 %	89.72 %	94.22 %	83.47 %	89.94 %	79.44 %	
Avg. # trades per pair		1.64	1.35	1.50	1.16	1.24	1.09	
Avg. % completed trades		59.28 %	40.12 %	51.66 %	27.12 %	40.38 %	14.01 %	
Panel d): rolling 2M standard deviation								
Total return		-65.71 %	-63.13 %	-64.62 %	-51.89 %	-63.75 %	-55.84 %	
Monthly returns	<i>Average</i>	-0.87 %	-0.81 %	-0.85 %	-0.59 %	-0.82 %	-0.66 %	
	<i>Std. deviation</i>	1.61 %	1.87 %	1.73 %	1.89 %	1.87 %	1.73 %	
	<i>Skewness</i>	-0.189	-0.169	-0.468	0.020	-0.363	0.025	
	<i>Kurtosis</i>	0.154	0.929	1.505	1.653	0.670	0.574	
	<i>Minimum</i>	-5.47%	-6.66%	-7.32%	-6.53%	-6.71%	-5.18%	
	<i>Maximum</i>	3.14%	4.27%	3.79%	5.69%	3.98%	4.21%	
Annualised SR		-2.44	-1.98	-2.22	-1.56	-2.01	-1.86	
Avg. % open pairs		87.47 %	90.47 %	94.75 %	86.75 %	91.47 %	82.69 %	
Avg. # trades per pair		1.82	1.52	1.61	1.22	1.41	1.14	
Avg. % completed trades		62.62 %	46.87 %	56.06 %	36.76 %	48.62 %	24.62 %	

Table 2.2. Unrestricted pairs portfolio, descriptive statistics.

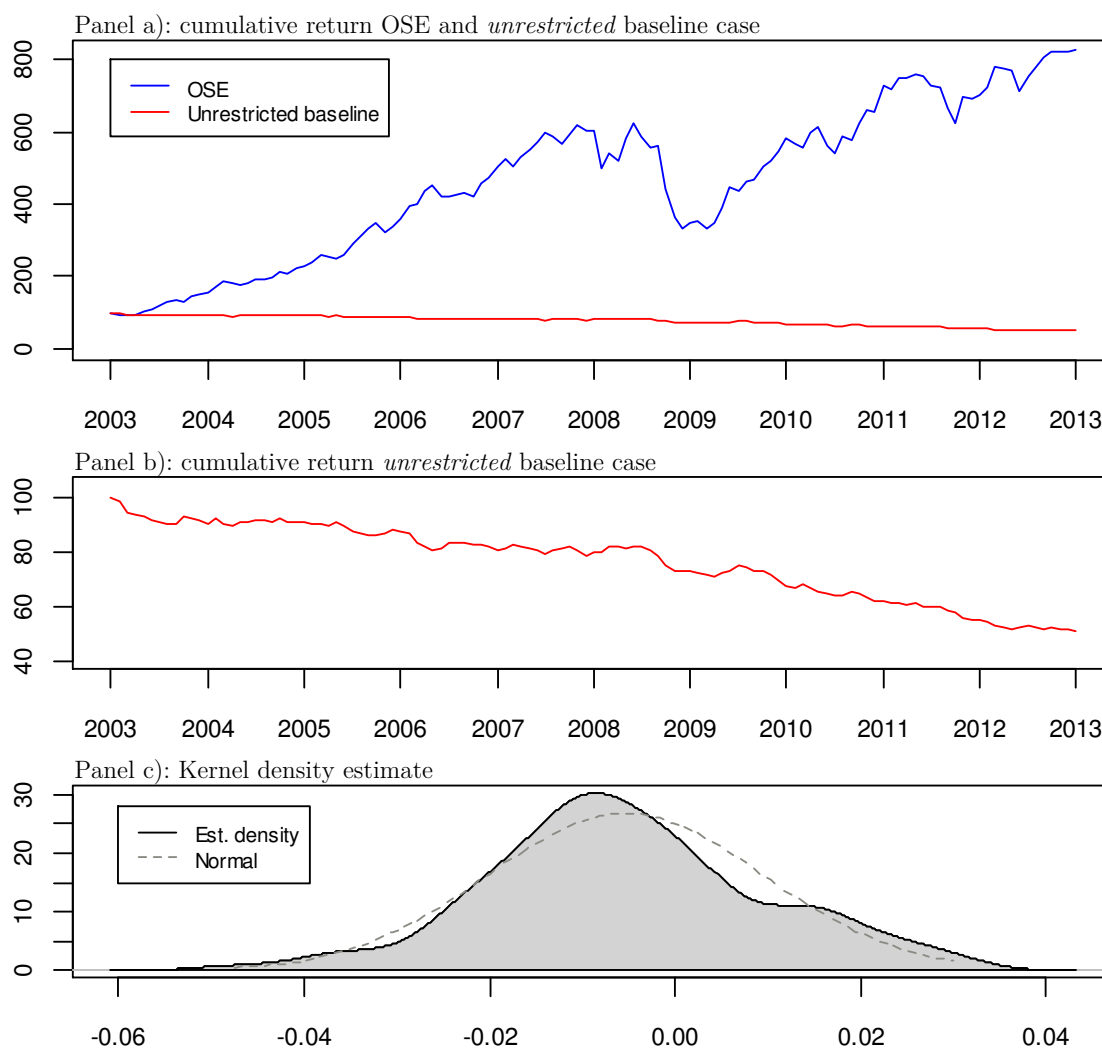


Figure 2.2. Return and distribution of unrestricted baseline portfolio.

As mentioned, Do and Faff (2010) finds that an increasing proportion of pairs do not converge upon divergence. Appendix 2.4 shows the development in the percentage of trades that are complete roundtrips, for each period. In contrast to their findings, we see no specific pattern in the data, indicating that the share of cointegration relationships that break down is moderately stable. An important question is therefore why the estimated relationships deteriorate so quickly, when they historically have been strong. Numerous reasons may exist. Firstly, two stocks that are found to be close substitutes in one period may cease to be so in forthcoming periods due to changes in fundamental factors affecting e.g. profit margins. The vast majority of pairs are formed with stocks belonging to different industrial sectors, so that any industry specific shocks might render a spread non-stationary, and consequently generate losses as the spread widens. Secondly, the cointegration coefficients may change; two stocks could still be cointegrated, but with different parameter values, so that the spread will oscillate

around a new mean value. This is called a *structural break*, and may generate losses if a position has been opened at divergence from the ‘old’ equilibrium value.

5.2.2 Restricted Portfolio

In order to reduce the risk of cointegration break-down we construct a new portfolio where pairs are restricted to belong to the same broad industry group as defined by the OSE: *consumer discretionary, consumer staples, energy, financials, healthcare, industrials, IT, materials, telecommunications* and *utilities*. Unlike the unrestricted portfolio, we trade all significant pairs from the formation period: on average we consider ten pairs per trading period, although this number greatly varies as shown in Appendix 2.5. Furthermore, about 90% of the pairs belong to the energy, industrials or IT sector – which is natural as these are the three major sectors on the stock exchange. Table 2.3 presents the descriptive statistics of the *restricted portfolio*. Again, panel a, column three, is the parameterisation used by Gatev *et al.* (2006) and is referred to as the *restricted baseline* case.

We see no specific improvement over the unrestricted portfolio. There are still relatively few trades per pair, together with a low percentage of completed trades. The restricted portfolio also suffers from the fact that most of the estimated cointegration relationships break down out-of-sample, thereby producing losses as the spreads keep widening throughout the trading period. The standard deviation of returns have increased slightly, which is natural seeing that we trade fewer pairs on average per period than in the unrestricted case, thereby reducing the effect of diversification. Although profitability seems to be sensitive to entry and exit points, this is in fact not true. Some cases outperform others simply because of one single period with extremely high return, which is to some extent evident by the higher standard deviations. The ‘slope’ of the cumulative return indices is still clearly negative for all cases. By trading a restricted pair portfolio we would have thought that there would be a lower fraction of pairs that would break down, due to the fact that industry specific risk would be reduced in each pair of stocks. When cointegration then do break down it is most likely due to other fundamental factors affecting only one stock in a pair, or affecting one stock heavier than the other.

Interestingly, we would be better off trading *against* the estimated relationships in the sense that the spread value could be seen as a momentum signal, so that one stock would continue to outperform the other. Put simply, profits would be positive if we were to bet against mean reversion – something that would be completely contradictory to the formation period.

Figure 2.3 shows monthly cumulative return of the restricted baseline case together with an index of Oslo Stock Exchange, and a Kernel estimate of the density of returns.

δ	Entry	1.5		2		2.5	
	Exit	0	-1.5	0	-2	0	-2.5
Panel a): constant standard deviation							
Total return		-51.7 %	-14.3 %	-49.7 %	-36.2 %	-45.6 %	-29.8 %
Monthly returns	<i>Average</i>	-0.58 %	0.02 % ^a	-0.55 %	-0.35 %	-0.49 %	-0.27 %
	<i>Std. deviation</i>	2.08 %	6.40 %	2.05 %	2.30 %	2.01 %	2.21 %
	<i>Skewness</i>	-0.432	8.871	0.381	-0.053	0.855	1.439
	<i>Kurtosis</i>	1.431	90.819	2.327	8.210	3.944	6.147
	<i>Minimum</i>	-6.79%	-10.75%	-6.18%	-11.65%	-6.09%	-5.39%
	<i>Maximum</i>	4.85%	64.95%	7.48%	10.40%	8.75%	10.85%
Annualised SR		-1.41	-0.13	-1.37	-0.91	-1.29	-0.83
Avg. % open pairs		90.14 %	85.71 %	89.10 %	73.84 %	81.92 %	70.09 %
Avg. # trades per pair		1.57	1.38	1.48	1.20	1.37	1.20
Avg. % completed trades		58.20 %	43.82 %	50.74 %	32.35 %	46.55 %	25.39 %
Panel b): rolling 6M standard deviation							
Total return		-49.9 %	-9.0 %	-54.0 %	-1.7 %	-38.0 %	-38.9 %
Monthly returns	<i>Average</i>	-0.55 %	0.08 % ^a	-0.61 %	0.15 % ^a	-0.38 %	-0.39 %
	<i>Std. deviation</i>	2.09 %	6.47 %	2.64 %	6.64 %	1.93 %	2.00 %
	<i>Skewness</i>	-0.132	8.534	0.959	8.805	0.135	-0.183
	<i>Kurtosis</i>	0.609	86.344	5.207	89.668	2.735	2.252
	<i>Minimum</i>	-6.35%	-11.80%	-6.40%	-11.13	-7.26%	-8.86%
	<i>Maximum</i>	4.90%	64.94%	12.58%	67.34%	7.28%	4.62%
Annualised SR		-1.35	-0.10	-1.14	-0.06	-1.15	-1.12
Avg. % open pairs		95.9 %	93.7 %	95.5 %	78.1 %	89.1 %	74.2 %
Avg. # trades per pair		1.63	1.37	1.45	1.18	1.23	1.08
Avg. % completed trades		57.3 %	42.0 %	48.3 %	25.0 %	42.3 %	10.1 %
Panel c): rolling 4M standard deviation							
Total return		-63.09 %	-55.11 %	-52.61 %	-12.57 %	-32.68 %	31.62 % ^a
Monthly returns	<i>Average</i>	-0.80 %	-0.63 %	-0.59 %	0.03 % ^a	-0.30 %	0.40 % ^a
	<i>Std. deviation</i>	2.28 %	2.57 %	2.39 %	6.10 %	2.66 %	6.89 %
	<i>Skewness</i>	-0.611	-0.518	-0.147	8.634	2.024	8.803
	<i>Kurtosis</i>	2.004	3.971	1.871	87.810	9.871	88.746
	<i>Minimum</i>	-9.87%	-11.08%	-8.29%	-12.67%	-6.61%	-9.57%
	<i>Maximum</i>	5.45%	8.71%	7.31%	61.41%	14.73%	69.97%
Annualised SR		-1.61	-1.20	-1.23	-0.13	-0.72	0.07
Avg. % open pairs		95.16 %	94.60 %	96.36 %	83.96 %	94.75 %	81.77 %
Avg. # trades per pair		1.66	1.39	1.51	1.17	1.33	1.12
Avg. % completed trades		57.19 %	44.18 %	51.26 %	28.38 %	46.18 %	22.54 %
Panel d): rolling 2M standard deviation							
Total return		-60.19 %	-57.62 %	-56.26 %	-43.04 %	-50.85 %	-31.99 %
Monthly returns	<i>Average</i>	-0.74 %	-0.68 %	-0.66 %	-0.44 %	-0.55 %	-0.27 %
	<i>Std. deviation</i>	2.06 %	2.63 %	2.34 %	2.29 %	2.73 %	3.21 %
	<i>Skewness</i>	-0.619	-0.059	-0.197	0.862	0.679	1.437
	<i>Kurtosis</i>	0.919	3.588	1.269	3.640	5.091	7.503
	<i>Minimum</i>	-6.96%	-11.02%	-7.39%	-7.29%	-10.55%	-10.55%
	<i>Maximum</i>	3.75%	10.11%	6.16%	9.37%	10.38%	16.23%
Annualised SR		-1.68	-1.23	-1.36	-1.06	-1.04	-0.57
Avg. % open pairs		93.12 %	93.86 %	97.26 %	87.10 %	95.11 %	85.33 %
Avg. # trades per pair		1.88	1.47	1.63	1.20	1.42	1.16
Avg. % completed trades		62.02 %	47.37 %	57.05 %	35.21 %	48.74 %	24.50 %

Table 2.3. Restricted pairs portfolio, descriptive statistics.

a) Return is positive due to a single period with extremely high return.




Figure 2.3. Return and distribution of restricted baseline portfolio.

5.2.3 Risk Decomposition

In order to provide insight on the risk and characteristics of the trading strategies, we decompose the return series by estimating the following regression equation:

$$\bar{r}_{P_t} = \alpha + \bar{r}_{M_t} + \bar{r}_{HML_t} + \bar{r}_{SMB_t} + \bar{r}_{PR1YR_t} + \varepsilon_t \quad (2.7)$$

where \bar{r}_{P_t} is the excess return on the pairs portfolio, \bar{r}_{M_t} is the excess market return, \bar{r}_{HML_t} and \bar{r}_{SMB_t} is the excess return on the two Fama-French factors *High Minus Low* and *Small Minus Big*, and \bar{r}_{PR1YR_t} is the excess return on Carhart's momentum factor. See Appendix 2.3 for more details on the pricing factors.

Panel a): unrestricted baseline case

	SS	df	MS		Number of obs.	120
Model	0.00263	4	0.00066		F(4, 115)	3.08
Residual	0.02450	115	0.00021		Prob. > F	0.0188
Total	0.02713	119	0.00023		R-squared	0.0968
					Adj. R-squared	0.0654
					Root MSE	0.0146
	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Market	0.03655	0.02910	1.26	0.212	-0.02110	0.09420
HML	0.00441	0.03482	0.13	0.899	-0.06456	0.07339
SMB	0.06790	0.04288	1.58	0.116	-0.01705	0.15284
PR1YR	-0.09917	0.03027	-3.28	0.001	-0.15913	-0.03922
Intercept	-0.00805	0.00140	-5.77	0.000	-0.01081	-0.00528

Panel b): restricted baseline case

	SS	df	MS		Number of obs.	120
Model	0.00150	4	0.00037		F(4, 115)	0.87
Residual	0.04951	115	0.00043		Prob. > F	0.4844
Total	0.05101	119	0.00043		R-squared	0.0294
					Adj. R-squared	-0.0044
					Root MSE	0.0208
	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Market	0.05648	0.04137	1.37	0.175	-0.02547	0.13843
HML	-0.01152	0.04950	-0.23	0.816	-0.10956	0.08653
SMB	0.01703	0.06096	0.28	0.780	-0.10372	0.13778
PR1YR	-0.03326	0.04303	-0.77	0.441	-0.11849	0.05197
Intercept	-0.00893	0.00198	-4.50	0.000	-0.01285	-0.00500

Table 2.4. Monthly risk exposure for unrestricted and restricted portfolios of pairs, *baseline case*.

Table 2.4 demonstrates why pairs trading can be considered market neutral. Whereas the restricted portfolio loads insignificantly on all five risk factors, the unrestricted portfolio only loads on the momentum factor with a negative sign, although the magnitude of the coefficient is severely small. Both portfolios have a significant negative alpha, as we would expect from eyeballing the time series plots.

As is commonly known, hedge ratios, or simply *betas*, tend to vary over time. We tried to capture the variability of the cointegration coefficients by requiring all pairs to cointegrate at a certain significance level using the average of historical rolling estimates (see point 3.2 under the estimation routine for the formation period). Even so, the relationships may significantly change out-of-sample, and the longer the trading period the more likely this is to happen. A shorter trading period, and consequently a higher frequency of parameter re-estimation, may help alleviate this problem. In figure 2.4 below we have constructed two portfolios for both the unrestricted and restricted case: one portfolio consists of the first two months of returns for each trading period (2F) whereas the second portfolio consists of the last

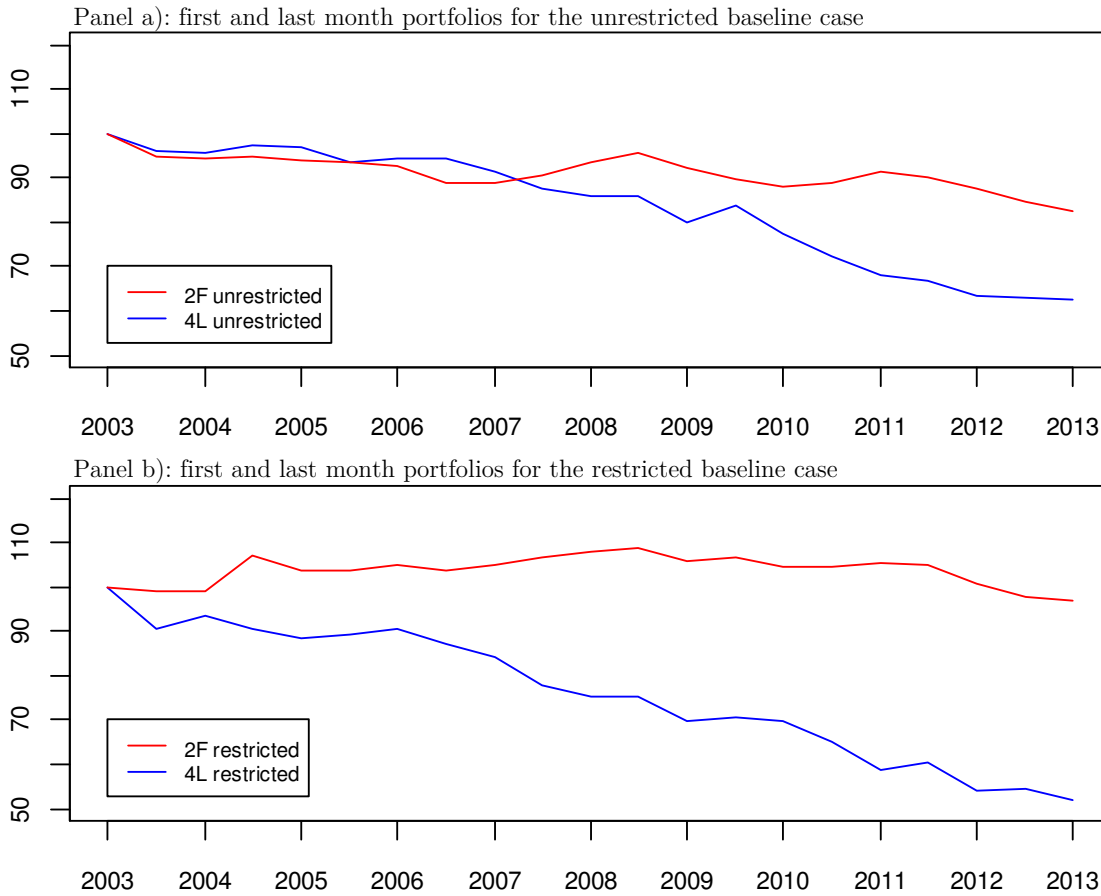


Figure 2.4. Synthetic portfolios of first and last month returns from each trading period.

four months (4L). If there are any gains to be made from a more frequent model re-estimation, we would expect that portfolio 2F would outperform portfolio 4L. Evidently, the 2F portfolio outperforms 4L, more so for the restricted baseline case. This indicates that a high share of the negative returns to the strategy occurs in the later months of trading, so that shorter trading periods could produce better results. Moreover, it serves to show how the forecasting ability of historical data diminishes as the time horizon increases.

Table 2.5 shows the share of returns smaller than zero, in addition to the maximum drawdown measured in both per cent and number of months. Maximum drawdown is defined as the peak-to-trough decline over the sample period, measured on monthly returns. This gives us a good identification of the downside risk associated with the portfolio, especially since the mean portfolio return could be biased upwards due to extreme outliers (as is the case with some of the restricted portfolio configurations). As suspected based on the negative alpha values, maximum drawdown is high for both portfolios regardless of entry and exit thresholds, clearly illustrating the unprofitability of the pairs trading strategy.

δ	Entry	1,5		2		2,5	
	Exit	0	-1,5	0	-2	0	-2,5
Panel a) Unrestricted baseline case							
Constant standard deviation							
Percentage of returns < 0		70,00 %	69,17 %	67,50 %	70,83 %	72,50 %	70,83 %
Maximum drawdown		-51,87 %	-52,01 %	-48,44 %	-41,60 %	-55,84 %	-46,52 %
Months of drawdown		120	100	120	112	120	120
Rolling 6M standard deviation							
Percentage of returns < 0		65,8 %	69,2 %	74,2 %	66,7 %	70,8 %	68,3 %
Maximum drawdown		-54,88 %	-47,53 %	-61,47 %	-39,21 %	-61,87 %	-53,61 %
Months of drawdown		118	92	120	97	120	119
Rolling 4M standard deviation							
Percentage of returns < 0		68,33 %	67,50 %	72,50 %	73,33 %	68,33 %	65,00 %
Maximum drawdown		-60,47 %	-58,56 %	-63,05 %	-55,96 %	-61,60 %	-49,63 %
Months of drawdown		120	120	120	97	120	119
Rolling 2M standard deviation							
Percentage of returns < 0		71,67 %	70,83 %	69,17 %	64,17 %	64,17 %	67,50 %
Maximum drawdown		-65,71 %	-63,13 %	-64,62 %	-52,07 %	-63,75 %	-55,84 %
Months of drawdown		120	120	120	119	120	120
Panel b) Restricted baseline case							
Constant standard deviation							
Percentage of returns < 0		60,83 %	65,00 %	60,00 %	59,17 %	62,50 %	61,67 %
Maximum drawdown		-52,13 %	-40,75 %	-50,06 %	-36,98 %	-45,57 %	-32,11 %
Months of drawdown		105	106	105	118	120	82
Rolling 6M standard deviation							
Percentage of returns < 0		57,5 %	61,7 %	60,8 %	59,2 %	63,3 %	61,7 %
Maximum drawdown		-51,57 %	-41,68 %	-56,20 %	-38,56 %	-38,03 %	-39,05 %
Months of drawdown		105	84	110	106	120	120
Rolling 4M standard deviation							
Percentage of returns < 0		60,00 %	61,67 %	62,50 %	60,00 %	57,50 %	59,17 %
Maximum drawdown		-63,27 %	-55,65 %	-53,03 %	-40,40 %	-43,13 %	-28,14 %
Months of drawdown		119	119	119	106	106	106
Rolling 2M standard deviation							
Percentage of returns < 0		59,17 %	65,00 %	65,00 %	62,50 %	62,50 %	60,83 %
Maximum drawdown		-60,41 %	-58,07 %	-56,82 %	-43,63 %	-50,85 %	-36,47 %
Months of drawdown		119	119	116	119	120	106

Table 2.5. Percentage of returns smaller than zero and maximum drawdown, for unrestricted and restricted portfolio of pairs.

5.3 Discussion of Findings and Concluding Remarks

Our findings coincide well with Do and Faff (2012) who conclude that pairs trading is largely unprofitable after 2002, seeing that our cointegration approach produces consistent negative returns independent of how entry and exit parameters are specified. While Gatev *et al.* (2006) claim that the significant decline in profits since the early 1990's might be due to increased activity among arbitrageurs, and therefore tighter spreads and less trading opportunities, our results suggest that it might be attributed to changes in fundamental factors governing the relationships between stocks. We seem to be unable to exploit historical statistical properties going forward, precisely as these properties change over time and are unpredictable. This is exactly in line with Do and Faff (2010), namely that a fairly high fraction of pairs will not converge upon divergence, thereby eroding profits completely. This may to some extent give support for the efficient-market hypothesis (weak-form efficiency), namely that we cannot achieve excess risk-adjusted market returns given past information.

However, because the back-test is done using daily data, we will not write off pairs trading entirely; there may still be potential for profit using higher frequency data, e.g. tick-by-tick, together with drastically shorter formation and trading periods – as practised by hedge funds. Any statistical relationships that are uncovered must quickly be exploited before they cease to exist, because, as we have shown, these relationships are not necessarily stable.

Moreover, as we have employed a purely statistical framework for trading pairs, we have not taken any fundamental aspects of the stocks into consideration during trading. In reality, a portfolio manager would continuously analyse company specific news so that she would have an idea of whether or not the spread will converge. The results in this thesis may consequently suffer from the fact that assumed trading behaviour is not fully realistic.

For future research it would be interesting to back-test a long-run strategy based on intraday data to compare with our results.

References

- Andrade, C. Sandro, Vadim di Pietro, and Mark S. Seasholes. "Understanding the Profitability of Pairs Trading." Working paper, UC Berkeley Haas School, 2005.
- Bowen, David, Mark C. Hutchinson, and Niall O'Sullivan. "High-Frequency Equity Pairs Trading: Transaction Costs, Speed of Execution, and Patterns in Returns." *The Journal of Trading* 5, no. 3 (2010): 31-38.
- Chiu, Mei Choi, and Hoi Ying Wong. "Dynamic Cointegrated Pairs Trading: Time-Consistent Mean-Variance Strategies." Working paper, The Chinese University of Hong Kong, 2012.
- Do, Binh, Robert Faff, and Kais Hamza. "A New Approach to Modeling and Estimation for Pairs Trading." Working paper, Monash University, 2006.
- Do, Binh, and Robert Faff. "Are Pairs Trading Profits Robust to Trading Costs?" *Journal of Financial Research* 35, no. 2 (2012): 261-287.
- Do, Binh, and Robert Faff. "Does Simple Pairs Trading Still Work?" *Financial Analyst Journal* 66, no. 4 (2010): 83-95.
- Elliot, J. Robert, John van der Hoek, and William P. Malcolm. "Pairs Trading." *Quantitative Finance* 5, no. 3 (2005): 271-276.
- Enders, Walter. *Applied Econometric Time series*. 3rd. New Jersey: John Wiley & Sons, 2010.
- Engelberg, Joseph, Pengjie Gao, and Ravi Jagannathan. "An Anatomy of Pairs Trading: The Role of Idiosyncratic News, Common Information and Liquidity." Working paper, University of California, University of Notre Dame, Northwestern University, 2009.
- Engle, Robert F., and C. W. J. Granger. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55, no. 2 (1987): 251-276.
- Galenko, Alexander, Elmira Popova, and Ivilnia Popova. "Trading in the Presence of Cointegration." *The Journal of Alternative Investments* 15, no. 1 (2012): 85-97.
- Gatev, Evan, William M. Goetzmann, and K. Geert Rouwenhorst. "Pairs Trading: performance of a Relative Value Arbitrage Rule." *The Review of Financial Studies* 19, no. 3 (2006): 797-827.
- Granger, C. W. J. "Some Properties of Time Series Data and Their Use In Econometric Model Specification." *Journal of Econometrics* 16, no. 1 (1981): 121-130.
- Gregory, Ian, Pieter Know, and Christian Oliver-Edwald. "Analytical Pairs Trading Under Different Assumptions on the Spread and Ratio Dynamics." Working paper, The University of Sydney, 2011.
- Hogan, S., R. Jarrow, M. Teo, and M. Warachka. "Testing Market Efficiency Using Statistical Arbitrage With Applications to Momentum and Value Strategies." *Journal of Financial Economics* 73, no. 3 (2004): 525-565.

- Hong, Gwangheon, and Raul Susmel. "Pairs-Trading in the Asian ADR Market." Working paper, University of Houston, 2004.
- Huberman, Gur. "A Simple Approach to Arbitrage Pricing Theory." *Journal of Economic Theory* 28, no. 1 (1982): 183-191.
- Huck, Nicolas. "Pairs Selection and Outranking: An Application to the S&P 100 Index." *European Journal of Operational Research* 196 (2009): 819-825.
- Huck, Nicolas. "Pairs Trading and Outranking: the Multi-step Ahead Forecasting Case." *European Journal of Operational Research* 207, no. 3 (2010): 1702-1716.
- Lanne, Markku. "Testing the Predictability of Stock Returns." *The Review of Economics and Statistics* 84, no. 3 (2002): 407-415.
- Lin, Yan-Xia, Michael McCrae, and Chandra Gulati. "Loss Protection in Pairs Trading Through Minimum Profit Bounds: A Cointegration Approach." *Journal of Applied Mathematics and Decision Sciences*, 2006: 1-14.
- Lo, Andrew W. "Long Term Memory in Stock Market Prices." *Econometrica* 59, no. 5 (1991): 1279-1313.
- MacKinnon, James G. "Critical Values for Cointegration Tests." *University of California San Diego Discussion Paper 90-4*, 1991: 267-276.
- Mudchanatongsuk, Supakorn, James A. Primbs, and Wilfred Wong. "Optimal Pairs Trading: A Stochastic Control Approach." *American Control Conference*. Washington, 11-13 June, 2008.
- Ødegaard, Bernt Arne. "Empirics of the Oslo Stock Exchange: Asset Pricing Results." Working paper, continuously updated: University of Stavanger. URL: http://finance.bi.no/~bernt/financial_data/ose_asset_pricing_data/index.html, 2013.
- Perlin, Marcelo Scherer. "Evaluation of Pairs Trading Strategy at the Brazilian Financial Market." *Journal of Derivatives & Hedge Funds* 15 (2009): 122-136.
- Pole, Andrew. *Statistical Arbitrage: Algorithmic Trading Insights and Techniques*. New Jersey: John Wiley & Sons, 2007.
- Said, Said E., and David A. Dickey. "Testing For Unit Roots in Autoregressive-Moving Average Models of Unknown Order." *Biometrika* 71, no. 3 (1984): 599-607.
- Teetor, Paul. *Better Hedge Ratios for Spread Trading*. 2011. <http://quanttrader.info/public/betterHedgeRatios.pdf> (accessed February 12, 2013).
- Vidyamurthy, Ganapathy. *Pairs Trading: Quantitative Methods and Analysis*. New Jersey: John Wiley & Sons, 2007.
- Wilmott, Paul. *The Best of Wilmott: Volume 1*. New Jersey: John Wiley & Sons, 2005.
- Wooldridge, Jeffrey. *Introductory Econometrics: A Modern Approach*. 4th. Canada: South-Western Cengage Learning, 2009.

Appendices

Appendix 1.1 – OLS vs. Orthogonal Regression (TLS)

To illustrate the difference between OLS and TLS we use the stock prices of two companies listed on the Oslo Stock Exchange: Deep Sea Supply (DESSC) and Solstad Offshore (SOFF). The data is from the period 04.07.2012-28.12.2012 – consisting of 125 trading days – and constitute the average of bid and ask prices, i.e. the midprice.

Figure 1.1 a) and 1.1 b) below depict OLS regressions between the logarithm of prices. In a) the dependent variable is DESSC, whereas in b) it is SOFF. The optimisation algorithm is illustrated by vertical lines connecting the residuals to the regression equation.

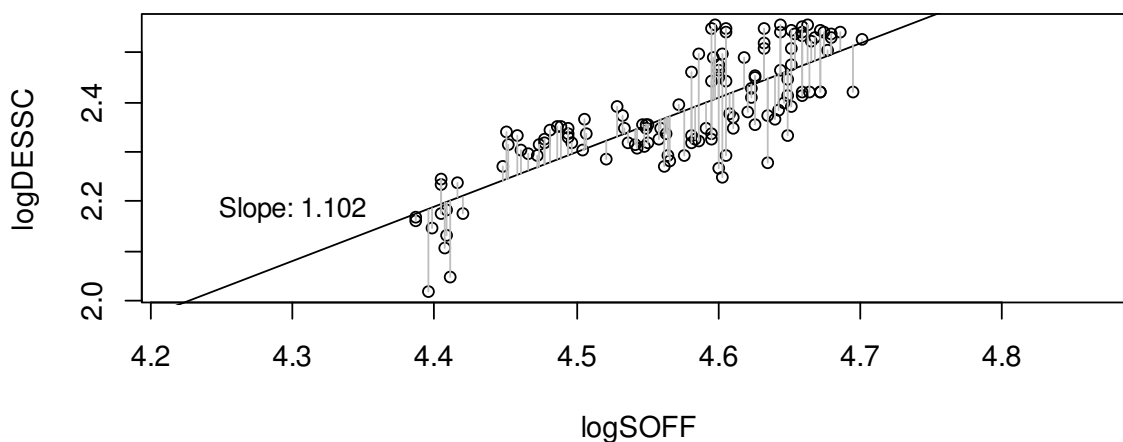


Figure 1.1 a). OLS regression: $\text{DESSC} \sim \text{SOFF}$.




Figure 1.1 b). OLS regression: $\text{SOFF} \sim \text{DESSC}$.

As can be seen, the two slopes are not the inverse of the other. This occurs because the OLS optimization algorithm only considers the variability of the dependent variable. Consider now the output from the same two regressions, this time using a TLS optimization procedure.




Figure 1.1 c). TLS regression: $\text{DESSC} \sim \text{SOFF}$.

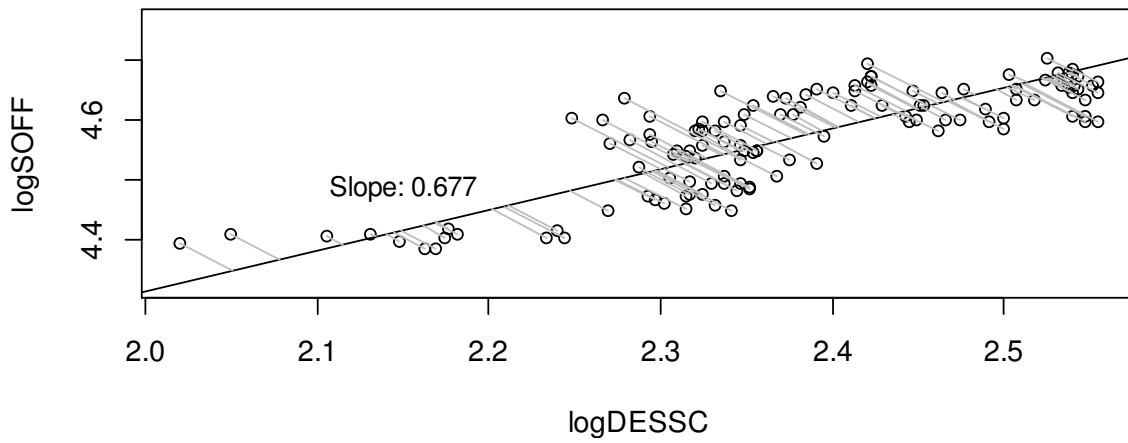


Figure 1.1 d). TLS regression: $\text{SOFF} \sim \text{DESSC}$.

The TLS optimisation algorithm minimises the *perpendicular* residuals of the regression line²⁰. The slopes are now invertible: $1/1.477 = 0.677$, producing consistent hedge ratios. Figure 1.1 e) and f) below show the estimated spreads and their absolute differences.




Figure 1.1 e). OLS regression residuals (spread).

²⁰ The aspect ratios of the plots are set to 0.35, so that the lines between the observations and the regression lines will not visually be right angles.




Figure 1.1 f). TLS regression residuals (spread).

The differences in the absolute values of the spreads show how TLS produce a more consistent hedge ratio. Note that for the purpose of this example the OLS and TLS spreads have been scaled, in the sense that the number of units of SOFF is equal in both regression equations.

Appendix 1.2 – Trading Thresholds

To illustrate how trading thresholds impact the number of trades and holding time we use two simulated cointegrated stock price series. We vary both entry- and exit-thresholds.

Figure 1.2 a) illustrates entry-thresholds of 0.20 and exit upon mean reversion. We count a total of six trades and an average holding time of 23 days.




Figure 1.2 a). Entry points: 0.2. Exit point: 0.0.

In figure 1.2 b) we have reduced the entry points to 0.10. This increases the number of trades to 10 and lowers the average holding time to 18 days.




Figure 1.2 b). Entry points: 0.1. Exit point: 0.0.

Figure 1.2 c) illustrates the impact of changing the exit points. If a trade is entered at 0.20 it will be unwound at -0.10. Number of trades is reduced from 6 to 5, and the average holding period increases to 38 days.




Figure 1.2 c). Entry points: 0.2. Exit points: -0.10.

Lastly we illustrate the use of a rolling window threshold. Entry points are 1.25 times the 20 day rolling std.dev, exit at mean reversion. 10 trades and an average holding time of 16 days.




Figure 1.2 d). Entry points: $1.25 \times \sigma_{20days}$. Exit point: 0.0.

Appendix 2.1 – Selecting Pairs

Below is a code written in R that identifies pairs of stocks suitable for trading.

```
#####
# This script finds pairs of stocks that are cointegrated.
# Written by Christoffer Hoel
#####
# Install packages if necessary (remove # before the command)
#install.packages("mcr")
#install.packages("tseries")
#install.packages("zoo")
#install.packages("urca")
#install.packages("fUnitRoots")
#install.packages("gtools")
#install.packages("Matrix")

rm(list = ls(all = TRUE))

#####
### PART 0 - Preparations
# Load the libraries
library(mcr)
library(tseries)
library(zoo)
library(urca)
library(fUnitRoots)
library(gtools)
library(Matrix)

# Read sample files
read.csv2("midprice2y.csv", header = TRUE) -> stocks
stocks <- log(stocks)
stocks <- zoo(stocks)
stocks <- as.ts(stocks)
nrStocks = length(stocks[1,])
nDays <- length(stocks[,1])
half <- nDays/2
halfround <- round(nDays/2, digits = 0)

if (half > halfround) {
    half <- halfround
} else
half <- half

read.csv2("bidaskspread.csv", header = TRUE) -> bidaskspread
bidaskspread <- zoo(bidaskspread)
bidaskspread <- as.data.frame(bidaskspread)

# Prepare matrices to store output information
ht2 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
ht1 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
htt <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
beta2 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
const2 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
beta1 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
```

```

const1 <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
betas <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
consts <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
betar <- numeric()
betat <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
constt <- matrix(data = NA, ncol = nrStocks, nrow = nrStocks)
baspread <- matrix(data = NA, ncol = nrStocks, nrow = 1)
spread <- list()

# This determines the critical value for the ADF-tests that will be performed.
# The critical values are based on MacKinnon (1991), as presented by Enders (2010).
# Critical values at the 5-pct and 10-pct levels.
crit52 <- 0
crit102 <- 0
crit51 <- 0
crit101 <- 0
if (nDays <= 49){
    crit52 <- -3.461
    crit102 <- -3.130
}
if (74 < nDays & nDays < 149){
    crit52 <- -3.398
    crit102 <- -3.087
}
if (149 < nDays & nDays < 350){
    crit52 <- -3.368
    crit102 <- -3.067
}
if (349 < nDays){
    crit52 <- -3.350
    crit102 <- -3.054
}
if (half <= 49){
    crit51 <- -3.461
    crit101 <- -3.130
}
if (74 < half & half < 149){
    crit51 <- -3.398
    crit101 <- -3.087
}
if (149 < half & half < 350){
    crit51 <- -3.368
    crit101 <- -3.067
}
if (349 < half){
    crit51 <- -3.350
    crit101 <- -3.054
}

# We now supply input to the baspread matrix, namely the average bid-ask spread for each stock.
# Because the spread (in this thesis) was defined as percentage spread (ie. (A-B)/A), we
# recalculate it using logarithms (ie. log(A)-log(B)).
for (j in 1:(nrStocks)) {
    baspread[, j] <- mean(log(1)-log(1-bidaskspread[, j]))
}

```

```

#####
### PART 1 - Finding pairs
#
# Ready, set, go! We now try to find the cointegrated pairs using the Engle-Granger
# two step procedure with orthogonal regression and MacKinnon critical values.
#####
# 2 year horizon
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {
    deming <- mcreg(stocks[, i], stocks[, j], method.reg = "Deming")
    const2[j,i] <- MCRresult.getCoefficients(deming)[1,1]
    beta2[j,i] <- MCRresult.getCoefficients(deming)[2,1]

    spread <- stocks[, j] - beta2[j,i]*stocks[, i] - const2[j,i]
    ht2[j,i] <- adfTest(spread, type = "nc", lags = 1)@test$statistic
  }
}
for (i in 1:(nrStocks-1)) {
  for (j in (i+1):nrStocks) {
    deming <- mcreg(stocks[, i], stocks[, j], method.reg = "Deming")
    const2[j,i] <- MCRresult.getCoefficients(deming)[1,1]
    beta2[j,i] <- MCRresult.getCoefficients(deming)[2,1]

    spread <- stocks[, j] - beta2[j,i]*stocks[, i] - const2[j,i]
    ht2[j,i] <- adfTest(spread, type = "nc", lags = 1)@test$statistic
  }
}
# We now filter out insignificant pairs.
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {
    if (ht2[j,i] > crit52 & ht2[i,j] > crit52 || ht2[j,i] > crit52 & ht2[i,j] <
        crit52 || ht2[j,i] < crit52 & ht2[i,j] > crit52){
      beta2[j,i] <- NA
      const2[j,i] <- NA
      beta2[i,j] <- NA
      const2[i,j] <- NA
      next
    }
  }
}
#####
# 1 year horizon
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {

    stock1 <- stocks[, i][half:nDays]
    stock2 <- stocks[, j][half:nDays]
    deming <- mcreg(stock1, stock2, method.reg = "Deming")
    const1[j,i] <- MCRresult.getCoefficients(deming)[1,1]
    beta1[j,i] <- MCRresult.getCoefficients(deming)[2,1]
    spread <- stock2 - beta1[j,i]*stock1 - const1[j,i]
    ht1[j,i] <- adfTest(spread, type = "nc", lags = 1)@test$statistic
  }
}

```

```

for (i in 1:(nrStocks-1)) {
  for (j in (i+1):nrStocks) {

    stock1 <- stocks[, i][half:nDays]
    stock2 <- stocks[, j][half:nDays]
    deming <- mcreg(stock1, stock2, method.reg = "Deming")
    const1[j, i] <- MCRresult.getCoefficients(deming)[1, 1]
    beta1[j, i] <- MCRresult.getCoefficients(deming)[2, 1]
    spread <- stock2 - beta1[j, i]*stock1 - const1[j, i]
    ht1[j, i] <- adfTest(spread, type = "nc", lags = 1)@test$statistic
  }
}

# We now filter out insignificant pairs.
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {

    if (ht1[j, i] > crit51 & ht1[i, j] > crit51 || ht1[j, i] > crit51 & ht1[i, j] <
    crit51 || ht1[j, i] < crit51 & ht1[i, j] > crit51){
      beta1[j, i] <- NA
      const1[j, i] <- NA
      beta1[i, j] <- NA
      const1[i, j] <- NA
      next
    }
  }
}

#####
# 2 year vs. 1 year
# We now filter out pairs that are not cointegrated in both periods.

for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {

    if (invalid(beta2[j, i]) || invalid(beta1[j, i])){
      next
    }else
      betas[j, i] <- beta2[j, i]
      consts[j, i] <- const2[j, i]
  }
}

for (i in 1:(nrStocks-1)) {
  for (j in (i+1):nrStocks) {

    if (invalid(beta2[j, i]) || invalid(beta1[j, i])){
      next
    }else
      betas[j, i] <- beta2[j, i]
      consts[j, i] <- const2[j, i]
  }
}

```

```

#####
### PART 2.1 - Selecting pairs - cointegration spread and bid-ask prices
# We now filter stock pairs based on the average bid-ask spread for each stock pair,
# in relation to a trading threshold of one standard deviation. In other words, if
# the average trading slippage is greater than the standard deviation of the spread,
# we will not trade the pair. We use the two year data as a proxy.
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {
    if (invalid(betas[j,i])) {
      next
    }else
    if (sd(stocks[, j] - stocks[, i]*betas[j,i] - consts[j,i]) < (baspread[, j] +
abs(betas[j,i]*baspread[, i]))){
      betas[j,i] <- NA
      consts[j,i] <- NA
      next
    }
  }
}
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {
    if (invalid(betas[i,j])) {
      next
    }else
    if (sd(stocks[, i] - stocks[, j]*betas[i,j] - consts[i,j]) < (baspread[, i] +
abs(betas[i,j]*baspread[, j]))){
      betas[i,j] <- NA
      consts[i,j] <- NA
      next
    }
  }
}
#####
### PART 2.2 - Selecting pairs - trading possibilities
# We now filter symmetrical pairs based on the spread with the highest standard
# deviation ABOVE the average trading slippage, i.e. A^B vs. B^A. The relationship
# with the highest standard deviation above the average trading slippage yields the
# best trading possibilities. We use the two year data as a proxy.
for (j in 1:(nrStocks-1)) {
  for (i in (j+1):nrStocks) {
    if (invalid(betas[j,i]) || invalid(betas[i,j])) {
      next
    }else
    if ((sd(stocks[, j] - stocks[, i]*betas[j,i] - consts[j,i]) - (baspread[, j] +
abs(betas[j,i]*baspread[, i])) > (sd(stocks[, i] - stocks[, j]*betas[i,j] -
consts[i,j] - (baspread[, i] + abs(betas[i,j]*baspread[, j]))))){
      betas[i,j] <- NA
      consts[i,j] <- NA
      next
    }else {
      betas[j,i] <- NA
      consts[j,i] <- NA
      next
    }
  }
}
}

```



```

#####
### PART 3 - Assessing historical hedge ratios (coefficients)
# We now assess the historical hedge ratios for the stocks in a given pair.
# Specifically, we calculate a rolling window of hedge ratios between the stocks.
# The calculation horizon is set to one year, with ratios updated every 5th day.
# The ratios are stored in a list, upon we then take the mean of the ratios. This
# mean is used as a hedge ratio (coefficient) between the two stocks. Using the
# hedge ratio and the stock' s price series a residual series is formed, i.e.
# residual = Y - b*X. After the series has been demeaned, an ADF test is
# performed to validate its stationarity, with a rejection level at 10-pct. If
# the pair passes, the coefficient and its corresponding intercept (computed by
# taking the mean of the residual series) is stored in new separate matrices,
# namely betat and constt.

# Calculating rolling coefficients
startpoint <- 0
endpoint <- half
for (j in 1:(nrStocks)) {
  for (i in (1):nrStocks) {
    if (invalid(betas[j,i])) {
      next
    }else
      betar <- numeric()
      for (l in seq(0, half, 5)) {
        start <- startpoint+l
        end <- endpoint+l
        rolling1 <- stocks[, j][start:end]
        rolling2 <- stocks[, i][start:end]

        deming <- mcreg(rolling2, rolling1, method.reg = "Deming")
        beta <- MCRResult.getCoefficients(deming)[2,1]
        betar <- append(betar, beta, after = length(betar))
      }
      betat[j,i] <- mean(betar)
    }
  }
}
next
}
# Constructing spreads and testing for stationarity

for (j in 1:(nrStocks)) {
  for (i in (1):nrStocks) {
    if (invalid(betat[j,i])) {
      next
    }else
      stock1 <- stocks[, j][half:nDays]
      stock2 <- stocks[, i][half:nDays]
      spreadt <- stock1 - betat[j,i]*stock2
      constt[j,i] <- mean(stock1 - betat[j,i]*stock2)
      spreadt <- spreadt - constt[j,i]
      htt[j,i] <- adfTest(spreadt, type = "nc", lags = 1)@test$statistic
    }
  }
}
}

```

```

# Filtering out stock pairs that are not stationary at the 10-pct level
for (j in 1:(nrStocks)) {
  for (i in (1):nrStocks) {
    if (invalid(betat[j,i])) {
      next
    }else
      if (htt[j,i] < crit101) {
        next
      }else
        betas[j,i] <- NA
        consts[j,i] <- NA
        betat[j,i] <- NA
        constt[j,i] <- NA
        next
      }
    }
}

#####
### PART 3 - Output
# We now write the output to two files that can be imported to Excel. We write the
# coefficients and intercepts from the 1 year horizon.

for (j in 1:(nrStocks)) {
  for (i in (1):nrStocks) {
    if (invalid(betas[j,i])) {
      betal[j,i] <- NA
      const1[j,i] <- NA
      next
    }
  }
}

write.csv2(const1, file = "const.csv")
write.csv2(betal, file = "beta.csv")
write.csv2(ht1, file = "t-statistic.csv")
write.csv2(constt, file = "const-avgr.csv")
write.csv2(betat, file = "beta-avgr.csv")
write.csv2(htt, file = "t-statistic-avgr.csv")

#####
### END OF SCRIPT
#####

```

Appendix 2.2 – Descriptive Statistics

Total return: the cumulative total return of the strategy.

Monthly returns:

Average: arithmetic average monthly return.

Standard deviation: standard deviation of monthly returns.

Skewness: skewness of monthly returns.

Kurtosis: kurtosis of monthly returns.

Minimum: minimum of monthly returns.

Maximum: maximum of monthly returns.

Annualised Sharpe Ratio: monthly Sharpe Ratio multiplied by the square root of twelve.

Average % open pairs: the percentage of pairs that open during a trading period.

Average # trades per pair: the average number of trades for a pair during a trading period.

Average % complete trades: the average percentage of trades that are complete roundtrips i.e. trades that are unwound when the spread crosses the exit threshold.

Appendix 2.3 – Pricing Factors

Market: A value weighted index of all listed equities on Oslo Stock Exchange is used as a proxy for the market. Illiquid and penny stocks are removed from the sample.

HML – High Minus Low: Fama-French factor, capturing the return of a portfolio of stocks with a high book-to-market ratio in excess of a portfolio of stocks with a low book-to-market ratio.

SMB – Small Minus Big: Fama-French factor, capturing the return of a portfolio of small stocks in excess of a portfolio of large stocks.

PR1YR – Previous 1 Year: Carhart's momentum factor, capturing momentum in stock market returns. Constructed as the average of stocks with the highest 30% eleven-month returns lagged one month minus the average of stocks with the lowest 30% eleven-month returns lagged one month.

Risk free rate: An estimate of the monthly risk free rate is obtained using monthly NIBOR quotes.

For a more detailed description please Ødegaard (2013). All data obtained from Ødegaard (2013).

Appendix 2.4 – Historical Average Percentage Completed Trades

<i>Baseline case</i>		Unrestricted pair formation	Restricted pair formation
Year	Subperiod	% completed trades	% completed trades
2003	1	53.8 %	50.0%
	2	50.0 %	50.0%
2004	1	59.1 %	85.7%
	2	50.0 %	25.0%
2005	1	58.1 %	69.2%
	2	69.0 %	75.0%
2006	1	31.6 %	0.0%
	2	47.6 %	37.5%
2007	1	60.9 %	44.4%
	2	71.4 %	52.9%
2008	1	79.4 %	85.2%
	2	64.3 %	62.5%
2009	1	58.3 %	62.5%
	2	30.0 %	31.3%
2010	1	45.5 %	32.0%
	2	52.4 %	33.3%
2011	1	55.0 %	80.0%
	2	58.1 %	56.3%
2012	1	63.6 %	46.2%
	2	50.0 %	38.5%

Table 2.1. Historical development the percentage of completed trades, baseline case.

Appendix 2.5 – Restricted Portfolio Composition

Year	Subperiod	Number of pairs
2003	1	8
	2	3
2004	1	5
	2	4
2005	1	18
	2	9
2006	1	6
	2	6
2007	1	7
	2	13
2008	1	20
	2	9
2009	1	7
	2	19
2010	1	27
	2	4
2011	1	7
	2	8
2012	1	15
	2	15

Table 2.2. Number of pairs for each period for the restricted pairs portfolio.

Year	Sub-period	Consumer discr.	Consumer stapl.	Energy	Financials	Health-care	Industrials	IT	Materials	Telecom	Utilities
2003	1	0 %	0 %	0 %	0 %	0 %	88 %	13 %	0 %	0 %	0 %
	2	0 %	0 %	100 %	0 %	0 %	0 %	0 %	0 %	0 %	0 %
2004	1	0 %	0 %	60 %	0 %	0 %	20 %	20 %	0 %	0 %	0 %
	2	0 %	0 %	50 %	0 %	0 %	25 %	25 %	0 %	0 %	0 %
2005	1	0 %	11 %	44 %	6 %	0 %	11 %	28 %	0 %	0 %	0 %
	2	0 %	0 %	33 %	0 %	0 %	11 %	56 %	0 %	0 %	0 %
2006	1	0 %	0 %	33 %	0 %	0 %	33 %	33 %	0 %	0 %	0 %
	2	0 %	0 %	50 %	17 %	0 %	0 %	33 %	0 %	0 %	0 %
2007	1	0 %	0 %	57 %	29 %	0 %	0 %	14 %	0 %	0 %	0 %
	2	0 %	0 %	38 %	0 %	0 %	54 %	8 %	0 %	0 %	0 %
2008	1	0 %	0 %	45 %	0 %	0 %	25 %	30 %	0 %	0 %	0 %
	2	0 %	0 %	33 %	0 %	0 %	11 %	56 %	0 %	0 %	0 %
2009	1	0 %	0 %	29 %	0 %	14 %	43 %	14 %	0 %	0 %	0 %
	2	0 %	11 %	47 %	0 %	5 %	32 %	5 %	0 %	0 %	0 %
2010	1	0 %	0 %	52 %	0 %	0 %	41 %	4 %	0 %	4 %	0 %
	2	0 %	25 %	50 %	0 %	0 %	0 %	25 %	0 %	0 %	0 %
2011	1	0 %	14 %	43 %	0 %	0 %	0 %	43 %	0 %	0 %	0 %
	2	13 %	0 %	38 %	13 %	0 %	13 %	25 %	0 %	0 %	0 %
2012	1	7 %	27 %	47 %	0 %	0 %	20 %	0 %	0 %	0 %	0 %
	2	0 %	20 %	47 %	7 %	0 %	27 %	0 %	0 %	0 %	0 %
Average		1,0 %	5,4 %	44,8 %	3,5 %	1,0 %	22,6 %	21,6 %	0,0 %	0,2 %	0,0 %

Table 2.3. Composition of pairs on sector/industry for each period for the restricted pairs portfolio.