Stockholm School of Economics
Department of Finance
Master's Thesis
Tutor: Peter Englund

Pairs Trading in European Equity Markets

Bill Sawarell & Can Inci

May 2014

ABSTRACT

In this paper we use a known pairs trading strategy and examine its performance in three separate European equity markets using daily data over the 20-year period between January 1994 and December 2013. We find that the strategy produces positive average excess returns in all three markets, and that alphas are significantly positive both when controlling for market exposure using the CAPM and when controlling for exposure to the European Fama-French factors. Returns are market-neutral using both specifications, and returns are generally positively correlated to volatility and negatively correlated to liquidity. The strategy produces results that are generally consistent across all three markets, and results are in accordance with previous research on the topic. The strategy was originally used in a 2006 paper, where the authors examined the strategy's performance in the U.S. equity market. To verify the validity of our algorithm we replicate a subset of the original authors' results and find that our algorithm produces results that compare favourably to theirs. This paper also serves as an out of sample test to the original authors' paper, while providing an up-to-date analysis of the strategy using data that among other developments includes the recent financial crisis.

Keywords: Pairs trading, convergence trading, statistical arbitrage, distance method, alpha

We want to thank our tutor Peter Englund for many helpful and encouraging meetings in his office at the Swedish House of Finance. We also want to thank our friends and families for their support throughout the process of writing this thesis. All remaining errors are our own.

TABLE OF CONTENTS

1. INTRODUCTION

The first title that we encountered when we started our research on pairs trading strategies was very discouraging: “Why I won’t teach pairs trading to my students”, an article by Lex van Dam (2012). The article was based on a disastrous experience of a billionaire with pairs trading which ended with a bankruptcy. Despite such unfortunate experiences and other discouraging news, pairs trading strategies remain popular in the investment world.

Pairs trading is a market neutral trading strategy which means that the expected returns from the trading strategy has no correlation with the market’s upward and downward movements. That is why students who first learn about it oftentimes are very enthusiastic, assuming this method can be used to realize “risk-free” returns. In reality, the first quant group that used the pairs trading strategy was disbanded from Morgan Stanley. They actually succeeded to realize gains in the beginning, but subsequent failure resulted in a $50 million loss and the group resolved. This quant group at Morgan Stanley, run by Nunzio Tartaglia, implemented and popularized the pairs trade even though their first experience was controversial.

After the mid-1980s hedge funds used this strategy widely and later studies were published in the academic world. The basic idea is to find a pair of stocks, sell the overvalued one and buy the undervalued one. The innovation of pairs trading is fundamentally based on the process of valuation, it relates to how an investor decides if a stock is overvalued or undervalued. In pairs trading, valuation is usually not focused on “true” values, rather “relative” values between securities are used. Stocks which share similar characteristics are broadly expected have a “fixed relative value” relative to each other. The idea behind a pairs trading strategy is to identify discrepancies from this relative value. If one of the stocks in a pair is overvalued relative to the other, the rule of trading is to sell the overvalued one and buy the undervalued one, expecting a convergence in relative prices. This expectation of convergence of relative values constitutes the basic and most important factor in pairs trading and this factor discriminates it from pure arbitrage. In literature this kind of trading is mostly defined as statistical arbitrage, which commonly involves using quantitative trading algorithms.

In 1999 a draft for a paper was circulated that documented high excess returns from a simple pairs trading strategy. The paper was later published in 2006 with somewhat updated data. This paper is the paper most influential to our analysis and it was written by Evan Gatev, William N. Goetzmann & K. Geert Rouwenhorst. The authors found that by employing a simple pairs trading strategy in the U.S. stock market they achieved historical excess returns of about 12% annually over a 40-year period. The paper was one of the first papers within the field of statistical arbitrage, and the fact that such high returns could be achieved with such a simple trading strategy made the paper somewhat groundbreaking and as such very influential. The authors did however notice a declining trend in profitability for the strategy towards the end of the sample period.

Pairs trading can be defined as a contrarian, or **convergent** trading strategy. Interestingly, trend following and **divergent** strategies have also been doing well for decades, but during the most recent years trend following and divergent strategies have performed poorly after the previous decades of high returns. In 2013 Futures Magazine published an article called 'Observations on the death of trend following', a title that well describes the recent shift to poor performance for the industry. It seems like something has changed with regard to divergent strategies, and it makes sense to ask whether a similar change has occurred for convergent strategies as well.

The first implication of positive returns sustained by pairs trading strategies is that they can be regarded as a sign of inefficiency of markets. This is because the basis of the strategy is built on the idea that prices will converge to their true relative values, which implies that by looking at historical prices of securities an investor can possibly make a profit. Gatev et al. (1999) describe this situation in their study, in which they examined pairs trading in the U.S. market, as "if the U.S. equity market were efficient, risk-adjusted returns from pairs trading should not be positive". However as previously mentioned they found that the strategy produced 12% annualized excess returns, over the period 1962 – 2002.

This paper is indented to add to the overall body of research within the field of statistical arbitrage and pairs trading. We accomplish this by using an already established trading strategy, but we use different markets and more recent data than previous research to make this paper unique and relevant. This study is based mainly on the Gatev et al. (2006) study, so the strategy we use is the same as the one used in their study. In our study we apply this pairs trading strategy to the top three biggest stock markets in Europe: the U.K. market, the German market and the French market, for the 20-year period of January 1994 – December 2013. The main idea of this paper is to examine how the strategy performs in these European markets, using very recent data. Some specific questions we seek to answer are:

i) Does the strategy produce significant excess returns?
ii) Are the results comparable across all three markets?
iii) Are the findings in accordance with the findings by Gatev et al. (2006)?

The paper is structured as follows: Part 2 provides an overview of pairs trading as a strategy and an overview of previous research on the topic. Part 3 describes the unique focus of this study and Part 4 describes the method used in the analysis. Part 5 presents the main results of the study, followed by a discussion in Part 6. In Part 7 we summarize our conclusions.

2. PAIRS TRADING AND LITERATURE REVIEW

Pairs trading became popular after the implementation of the strategy by Nunzio Tartaglia's quantitative group at Morgan Stanley in the 1980s. The group made a $50 million profit in 1987, but as previously mentioned the group was disbanded by Morgan Stanley after poor performance in the following years. Despite the unsuccessful ending to the experience of the group pairs trading became popular in the investment world.

Pairs trading is a market neutral investment strategy. Vidyamurthy (2004) explains pairs trading as a 'market neutral statistical arbitrage strategy'. In its most simplistic form this means a trading strategy where the returns from the strategy are uncorrelated with the returns of the market. Theoretically an investor who applies a market neutral strategy will not be affected by fluctuations in the market. Whether the market goes up or down the investor's return is expected to be unaffected. The basis of the strategy is the relative pricing of two securities. It is assumed that stocks with similar characteristics are priced similarly. Similar prices do not mean exactly the same nominal prices, but it means that the ratio of two stocks' prices is expected to stay more or less the same over long periods of time. A market neutral strategy is formed by taking positions in two securities: a long position and a short position. A spread between these two securities is calculated, which can be calculated using different methods which will be discussed later. The spread can be thought of as the degree of mispricing. The magnitude of the spread is linearly related to the expected return of the strategy. The greater the spread the greater the expected return. Pairs trading is basically forming positions when this spread is far away from its mean value expecting that the spread will revert back. The position is reversed when convergence occurs. The mean reverting mechanism of pairs trading identifies it as one of the well-known convergent trading strategies.

Several contrarian trading strategies have previously been documented to be profitable, see for example Jegadeesh & Titman (1995b), and several attempts have been made to explain what the reason for this profitability is. As explained by Gatev et al. (2006) pairs trading is in essence a contrarian trading strategy. As such, the conclusions that have been reached trying to explain the profitability of contrarian strategies can be extended to explain the profitability of pairs trading.

The profitability of contrarian strategies is commonly attributed to short-term overreactions in the market. The so called 'overreaction hypothesis' predicts that people tend to overreact to dramatic news and events, regardless of whether these events are positive or negative in nature (Mun et al., 2001). Overreactions lead to deviations from fundamental values, which in turn can be exploited by contrarian traders. This view is still a common explanation of contrarian returns, and was also one of the earliest explanations. See for example Lehmann (1990) and Jegadeesh (1990). Consistent with the predictions of the overreaction hypothesis, portfolios of prior "losers" are found to outperform prior "winners." (Bondt & Thaler, 1985). An explanation of the excess returns produced by pairs trading strategies specifically is mentioned in Gatev et al. (2006) as the human nature, or 'behavioural', explanation, which is related to overreaction. The phenomenon of overreaction can have several explanations. It was described simply as "Human beings don't like to trade against human nature, which wants to buy stocks after they go up not down" by David Shaw, who is head of a hedge fund and a follower of Tartaglia's methods from Morgan Stanley, as mentioned in Hansell (1989).

Another explanation for the profitability of contrarian trading strategies is that there are differences in reaction times between stocks. This is an explanation most notably put forth by Lo & MacKinlay (1990). They conclude that "If returns on some stocks systematically lead or lag those of others, a portfolio strategy that sells "winners" and buys "losers" can produce positive expected returns, even if no stock's returns are negatively autocorrelated as virtually all models of overreaction imply." Andrade et al. (2005) relate divergence from relative prices to uninformed price shocks.

A further explanation is related to the liquidity of the marketplace. Some argue that lack of short term liquidity pushes prices up more than what is justified by fundamentals in the short term, and that this can be exploited by assuming a return to fundamental prices over the longer term. In this theory liquidity is assumed to be sufficient over the long term for prices to return to fundamentals (Grossman & Miller, 1988; Jegadeesh & Titman, 1995a). Nagel (2012) concludes that short-term reversal strategy returns can be interpreted as compensation for liquidity provision, and documents a positive relationship between expected reversal returns and volatility.

A final possible explanation for contrarian returns is that stock prices react with lags to common factors but overreact to firm specific information (Jegadeesh & Titman, 1995b). In this study the authors found an overreaction by investors to company specific information shocks. According to this theory of overreaction contrarian investment strategies that are buying when prices are falling and selling when prices are rising rapidly can generate profits. This implies that firm-specific information leads to temporary mispricings between stocks which, again, can be exploited by assuming a return to fundamentals. This finding is relevant for pairs trading, which uses the historical relationship between individual stocks to form the strategy.

Do & Faff (2010) implemented the pairs trading strategy of Gatev et al. (2006) with an extended sample period up until June 2008. They still achieved significant excess returns after 2002 but also they stressed that the profitability from pairs trading is declining over time.

Types of Pairs Trading

An understanding of the concept of pairs trading is fundamental to this study. Given this, we provide the reader with a brief description of the different types of pairs trading, even though only one of these methods will be used in this particular study. Broadly, pairs trading strategies can be classified into four different methods, each of which will be depicted briefly below. If the reader is already familiar with the concept of pairs trading we suggest the reader skips directly to Part 3. If the reader is specifically interested in the method used in this particular study, we advise the reader to mainly focus on the so called 'Distance Method' below. The four types of pairs trading are:

The Distance Method

The famous Gatev et al. (2006) study is based on this method. Gatev et al. (2006) identify liquid stocks and calculate cumulative total return indices for all stocks in the sample. Pairs are formed by trying to find pairs of securities that historically have tended to 'move together', as they describe it. This basically means that if two stocks' price series are normalized to have the same starting value, their cumulative return series would look similar if plotted over time. They then choose pairs to trade by finding the pairs of securities that minimize the sum of squared daily differences between the normalized price series. This is why it is called the distance method, because it is based on the distance between the normalized price series. They open a trade when prices diverge by two historical standard deviations. A short position is taken in the higher priced stock and a long position in the lower priced stock. The position is open until the stocks converge or until the end of a predetermined trading period. In their study Gatev et al. (2006) formed pairs over twelve month periods and traded the pairs over the following six month periods. Again, they achieved up to 12% returns annually.

Andrade et al. (2005) used the distance method in the Taiwanese market using data from 1994 to 2002 and they found excess returns of more than 10% annually. In this study, as previously mentioned, they linked the profitability of pairs trading to uninformed trading shocks. In another study Perlin (2009) found profitable results for the Brazilian market after implementing a strategy based on the distance method. He concluded that the profitability of the strategy increases with the frequency of trading. He also found that daily trading strategies had superior returns over weekly and monthly strategies. Nath (2003) further proposed a simple risk management framework as an extension to the strategy which can be stated as a stop-loss corridor to keep the spread between securities below a predetermined level. The distance method will be explained further later as it is the method used in our study.

The Cointegration Method

Vidyamurthy's (2004) book is mainly based on the cointegration method. This brief description is mainly based on his book. Do et al. (2006) provide a comparison between the cointegration method and other methods.

In this method, parameterization of the pairs trading strategy is carried out by trying to find pairs by cointegration (Engle & Granger, 1987). If the time series of prices of two securities are both cointegrated of order d then these two time series can be linearly combined to form a single time series of order $d - b > 0$. In the most simple case $d = b = 1$. When the combined time series are stationary it is eligible to use the time series for forecasting. Alternatively co-integrated time series can be represented in an Error Correction Model (ECM). In this model the dynamics of a time series is based on the correction of the last period's deviation from the equilibrium. As such, forecasts can be done based on past information. To test for cointegration Vidyamurthy (2004) uses Engle and Granger's 2-step approach (Engle & Granger, 1987). The log price of stock A is first regressed against the log price of stock B which is defined as the cointegration equation (1.1).

$$Log\ P_t^A - \propto Log\ P_t^B = \mu + \varepsilon_t \qquad (1.1)$$

Where $\propto$ is the cointegration coefficient and μ captures the premium in stock A versus stock B. The residual of the above equation is then tested for stationarity using an Augmented Dickey Fuller test (ADF). From equation (1.1) a trader buys 1 unit of stock A and shorts $\propto$ units of stock B. This has a long run equilibrium value of μ. The deviations from the equilibrium value are temporary because the residuals ε_t are found to be stationary. Any deviations from the equilibrium value will be corrected, and this is the mean reverting element of the pairs trading strategy.

$$Log\ P_t^A \ -\propto Log\ P_t^B \ = \mu - \Delta \qquad (1.2)$$
$$Log\ P_{t+i}^A \ -\propto Log\ P_{t+i}^B \ = \mu + \Delta \qquad (1.3)$$

The trading strategy described by Vidyamurthy (2004) implies buying the portfolio (buy A, short B) when the time series is Δ below the mean and selling the portfolio (short A, buy B) when the time series is Δ above the mean in steps of i. The profit will be the incremental change in the spread 2Δ. The trading strategy is summarized by equations (1.2) and (1.3).

The Stochastic Spread Method

This method is described in detail in the paper of Elliot et al. (2005). The main topic of the paper is the modeling of the spread. The spread is defined as the difference between the prices of two securities. As a general rule, a pairs trading strategy is long in one security and short in another security at a predetermined ratio. The resulting portfolio can be market neutral based on the selection of this ratio. The authors state that it is possible to model the spread as a mean reverting process, which they calibrate from market data. This model allows an investor to make predictions for the evolution of the spread. If observations are larger (smaller) than the predicted value (by some threshold value), the investor takes a long (short) position in the portfolio and unwinds the position for a profit when the spread reverts.

Do et al. (2006) state that the model in Elliot et al. (2005) has three major advantages: Firstly, it captures mean reversion which is the basis of pairs trading and the crucial part of the mechanism to apply to the strategy safely. Secondly, the model in Elliot et al. (2005) is a continuous time model and this allows for forecasting. A trader can compute the expected time that the spread takes to converge back to its long term mean, so expected holding time periods and expected returns can be computed, which are very crucial for a pairs trader. Lastly, the model is tractable which means that the parameters of the model can be estimated.

Despite the above mentioned advantages Do et al. (2006) state that this approach has a fundamental issue in that the model restricts the long run relationship between two stocks to one of return parity. This implies that returns of stocks in the long run must be the same and that any departure is expected to be corrected in the future. Do et al. (2006) argue that this is a big restriction because in practice it is rare to find two stocks with identical returns. Given this, Do et al. (2006) asserts that the Stochastic Spread Method should preferably be applied to a dual listed company (DLC) structure. In such a structure shares are traded on different exchanges, and shareholders are entitled to the same cash flows. Given these characteristics, such structures tend to attract pairs traders widely. Examples of such companies are Unilever NV /PLC, Royal Dutch Petroleum /Shell (until 2005) and BHP Billiton Limited /PLC. Do et al. (2006) also show that companies that are cross listed are other candidates for this strategy, with the rationale being similar to that of the DLC structure.

The Stochastic Residual Spread Method

Do et al. (2006) formulize a new approach to model relative mispricings for pairs trading purposes in a continuous time setting, and the brief description here is solely based on this paper. The 'new' idea in this approach lies in the quantification of the mean reversion behavior, taking into account theoretical asset pricing relationships which are different than previous approaches. The model starts with an assumption that there exists some equilibrium in the relative valuation of two stocks measured by some spread. Deviations from this equilibrium are interpreted as mispricings, and a mispricing in turn is defined as a state of disequilibrium. This disequilibrium is quantified by a 'residual spread' function as

$$G(R_t^A, R_t^B, U_t)$$

where U denotes 'some exogenous vector potentially present in formulating the equilibrium'.

The idea of the residual spread function is to capture any deviations, 'residuals', from the long-term spread, which is the equilibrium, and assuming that the relative valuation of the securities should mean revert to the equilibrium in the long run. Trading takes place when the disequilibrium is large enough and the expected correction timing is short enough. The proposed method adopts the same modeling framework as in Elliot et al. (2005). The basic idea of this method is basically what defines pairs trading as a strategy, to look for temporary deviations from long-term trends. Their paper is however quantitative in nature, and we will not to go too far into the details since it is deemed to be beyond the scope of our study.

Do et al. (2006) state that unlike existing pairs trading strategies, which are based on mispricings at the price level, their model is based on mispricings at the return level. The existing methods open positions when prices drift apart and unwind when they converge. In the proposed model of Do et al. (2006) positions are opened when the accumulated residual spread in the returns is large enough and closed out when the accumulated spread is equal to the long run level of the spread.

The pairs trading strategies described above are all different with regard to the actual implementation, but the main underlying idea is always the same. The strategies always revolve around trying to identify when securities are mispriced relative to each other, and trade based on the belief that these mispricings will be corrected. Some of the techniques, such as the distance method, are fairly straightforward and easy to understand, whereas others are more quantitative in nature. This paper is based on the distance method, so this method will be described in more detail later in the paper. For more detailed explanations of the other methods, we kindly direct the reader to the references mentioned in the above descriptions.

3. FOCUS OF STUDY

In this paper we employ the same technique as Gatev et al. (2006) use in their study and trade according to the signals given by the strategy. Although we base our analysis on the Gatev et al. (2006) study, we want it to provide additional information to the body of research within the field of pairs trading and not mainly carry out a replication of their analysis. As such, our analysis will differ in some significant ways. The main differences will be that we will have a European focus, and that the analysis will be carried out with more recent data that among other things includes the recent financial crisis. The strategy is tested in three separate markets. As datasets we use the most liquid stocks in the British FTSE 100 index, the French CAC 40 index and the German DAX 30 index. The indices are described briefly below.

- The FTSE100 index is a stock index consisting of the 100 companies with the largest market capitalizations on the London Stock Exchange (LSE). The index is maintained by the LSE's subsidiary group the FTSE Group and is one of the most widely used European stock indices.
- The CAC (Cotation Assistée en Continu) 40 is a French stock index, consisting of the 40 stocks with the highest market capitalizations on the Euronext Paris (Paris Bourse). The index is almost exclusively composed of French companies and the index is operated by Euronext.
- The DAX (Deutscher Aktien IndeX) 30 is a German stock index that consists of 30 of the largest stocks on the Frankfurt Exchange. The index is operated by the Deutsche Börse.

These are likely the most prominent stock market indices in Europe, and the respective countries are the three largest economies in Europe as measured by GDP. These facts make these markets logical choices for a study with a European focus. The study is carried out for the 20 year time period January 1994 – December 2013. We choose a time period of 20 years because we want an up-to-date study, but at the same time have part of our sample overlap with the original Gatev et al. (2006) study for comparison. The Gatev et al. (2006) study used data from 1962 - 2002, meaning that roughly the first half of our sample overlaps with the latter part of their sample. As such, it could partly be interpreted as an out-of-sample test for the latter part of their study. The second half of our sample extends the analysis to the present.

4. METHOD

In our study we use the distance method which, as explained above, is the same method used in Gatev et al. (2006). The technique is based on the concept of statistical arbitrage. Statistical arbitrage is not arbitrage in the sense that it is risk-free, but rather it is based on probabilities. If something is highly likely to yield a profit and this can be exploited it can be viewed as a kind of arbitrage, although not in the purest meaning of the term. In this chapter we describe the strategy in detail and the underlying rationale.

The technique employed by Gatev et al. (2006) is based on trying to find securities that tend to 'move together' and exploit periods when they drift apart by taking positions based on the belief that they will converge again. A set of stocks are chosen, for example all liquid stocks in an index, and all possible pairs of stocks in the sample are analyzed over a period of one year, called the '**formation period**'. At the beginning of the formation period all stock prices are normalized to 1 such that they all have the same starting price. Then, pairs of securities that tend to 'move together' are identified by squaring the daily differences between stocks and finding the pairs that minimize the sum of squared differences over the full formation period. In essence, finding a pair can be thought of as simply finding two time series of normalized cumulative returns that look similar to each other. This method of finding pairs is, again, referred to as the 'distance method'.

The five pairs of stocks that have the smallest sums of squared differences over the one year formation period are chosen, and these five pairs are then traded over the following six month period which is called the '**trading period**'. The idea is that these pairs of securities have exhibited similar price movements during the formation period and because of this are expected to continue to exhibit similar price movements during the trading period. As such, if the stocks in a pair start to drift apart during the trading period, it is assumed that they will converge again and positions in the stocks are taken accordingly. Stocks are only traded during the six month trading period, and the trigger to open a position in a pair is if the normalized prices of a pair of stocks differ by more than two standard deviations, as measured over the formation period. To open a position in a pair means taking a short position in the stock with the highest normalized price, and taking a long position in the stock with the lowest normalized price. In essence, it is taking a position based on the belief that the stock prices will converge rather than diverge further. A position is closed out either if the pair actually converges or if the six month trading period ends. If a pair converges during the trading period it can be opened again during the same trading period if the stocks again drift apart by two standard deviations. Each long and short position is assumed to be taken on a nominal one currency unit basis.

This process of forming pairs is repeated at the beginning of every month on a rolling basis, such that during a full year twelve formation periods will be initiated. Equivalently, twelve trading periods will be initiated during a full year, except for during the first year of a sample which exclusively serves a formation period where no trading can take place. This means that we will end up with rolling six month trading periods overlapping by one month. This can be viewed as a fund having six different managers with two trading periods each per year, such that these managers each trade consistently during the year and where the beginning of a new trading period is exactly at the end of the previous one. Given this, the analysis could be carried out with just two trading periods each year, and a new formation period beginning every sixth month, which would be equivalent of having just one manager at the assumed fund. Using a one month rolling approach however yields more output from the analysis and can lend more credibility to the results since the technique is repeated more frequently and thus yields more observations.

The strategy is deliberately relatively simple. Part of the reason to keep the strategy simple is that the reader can see that the strategy is not optimized to yield high returns, but rather designed to test the concepts of convergent trading and statistical arbitrage broadly. It is easy to optimize strategies to fit past data to find high returns, but the simple twelve- and six month periods, alongside the simple entry/exit rules for the strategy should make it clear that this strategy is not optimized to yield high returns. It should also be noted that even though the strategy is relatively simple, the strategy is not necessarily any easier to implement than more elaborate strategies.

Illustration of the Strategy

Below we provide a visual representation of the strategy to further clarify how the technique works (**Graph 1**). It is a real example from our analysis of the FTSE 100 index, and it displays one pair of stocks over a full 18 month cycle, containing a 12 month formation period and a following 6 month trading period. The stocks are Sage Group and Aviva. The time period displayed is January 2007 to June 2008. As such, the formation period is all of 2007 and the trading period is the first 6 months of 2008. The formation period and the trading period are separated by the vertical red dotted line in the graph.

At the beginning of the formation period the stocks' prices are normalized to 1. During the formation period the daily differences in normalized prices are calculated, squared, and at the end of the full one-year period the sum of squared differences is calculated. This sum is then compared to all other pairs, and if the sum is found to be among the 5 smallest, the pair of stocks is chosen for trading during the following trading period alongside 4 other pairs. This particular pair was found to be one of the 5 pairs with the smallest sums of squared daily differences and thus is traded in the trading period. During the formation period the standard deviation of the difference in normalized prices is calculated, as it will later be used to signal when to enter into a trade. At the beginning of the trading period the stocks' prices are normalized to 1 again. During the trading period we can see that there are 4 position openings (and 4 position closings). The shaded areas indicate when a position is open. Again, a position opens when the stocks' normalized prices differ by more than 2 standard deviations, as calculated during the formation period, and closes when the pair converges or at the end of the trading period.

GRAPH 1. ILLUSTRATION OF THE STRATEGY

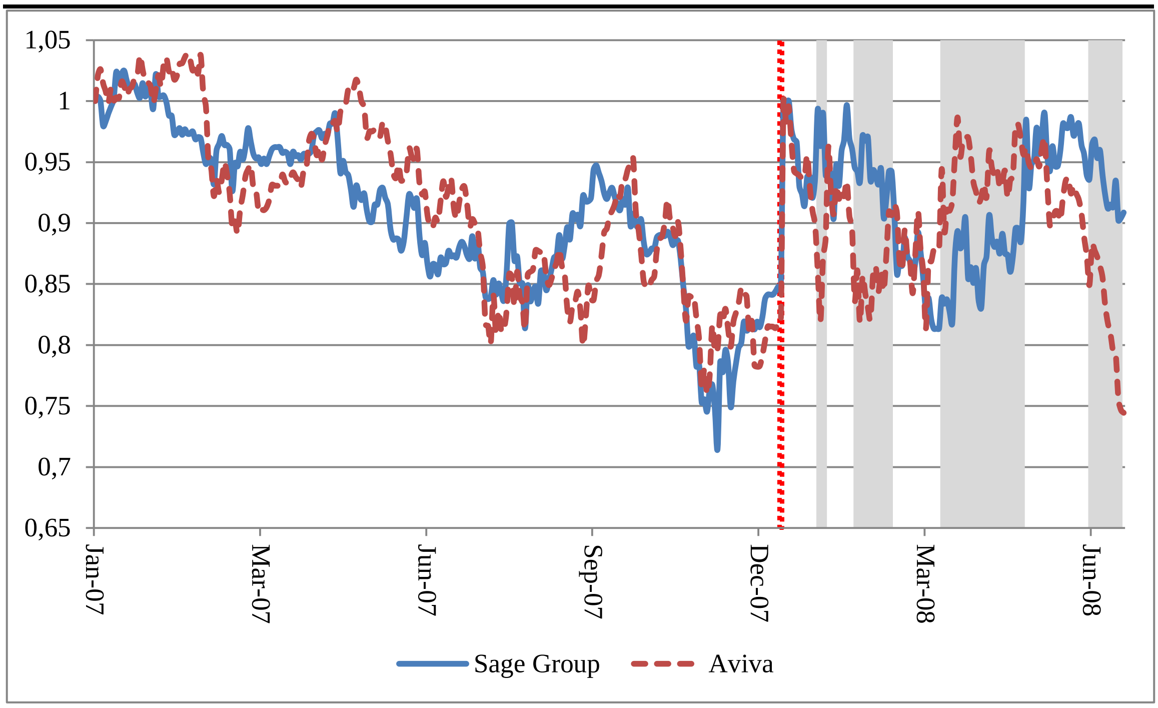

Graph provides a visual representation of how the strategy functions over a full 18 month cycle. The stocks in the graph are Sage Group and Aviva and the time period is January 2007 - June 2008. The red line indicates where the trading period starts and the formation period ends. Shaded areas indicate periods where trading takes place.

The following happens during the trading period for this particular pair:

- On January 18 2008 the normalized prices differ by more than 2 standard deviations and a trade is opened. A short position is taken in Sage Group because it has the highest price of the two, and a long position is taken in Aviva. 4 days later on 23 January 2008 the pair converges and the trade is closed out for a gain.
- On 7 February 2008 the normalized prices again differ by more than 2 standard deviations and a trade is opened. Again a short position is taken in Sage Group because it has the highest price of the two, and a long position is taken in Aviva. 14 days later on 27 February 2008 the pair converges and the trade is closed out for a gain.
- On 25 March 2008 the normalized prices again differ by more than 2 standard deviations and a trade is opened. A short position is taken in Aviva because it has the highest price of the two, and a long position is taken in Sage Group. 31 days later on 7 May 2008 the pair converges and the trade is closed out for a gain.
- On 11 June 2008 the normalized prices, again, differ by more than 2 standard deviations and a trade is opened. A short position is taken in Sage Group because it has the highest price of the two, and a long position is taken in Aviva. The pair never converges and the trade is closed out for a loss on 30 June 2008 since the pair has diverged even further than when the position was opened.

Replication of Previous Results

To make sure that our algorithm yields the same results as the Gatev et al. (2006) study, we replicate one of the key graphs from their paper using our own trading algorithm. Their graph shows how one pair of stocks behaves during the trading period ranging from August 1963 - February 1964 and can be found on page 8 of their study. The stocks constituting the pair are Kennecott and Uniroyal. For this replication only we download daily data on the same stocks from CRSP (via WRDS), which is the same data source they used. Applying our own algorithm on these two stocks over the same time period yields very favourable results that are almost identical to theirs, as can be seen in **Graph 2** below. Seeing this we can be certain that our trading algorithm is properly coded and as such we can be confident that our study is carried out accurately. The output from our algorithm is shown below and can favourably be compared to the graph in their paper. The two lines in the graph show the normalized price series for the two stocks and the shaded areas indicate when trading is taking place. As previously mentioned the trigger to enter a trade is when the normalized prices differ by more than 2 standard deviations, and the trigger to exit is either when the normalized prices converge or when the six month trading period ends.

GRAPH 2. REPLICATION

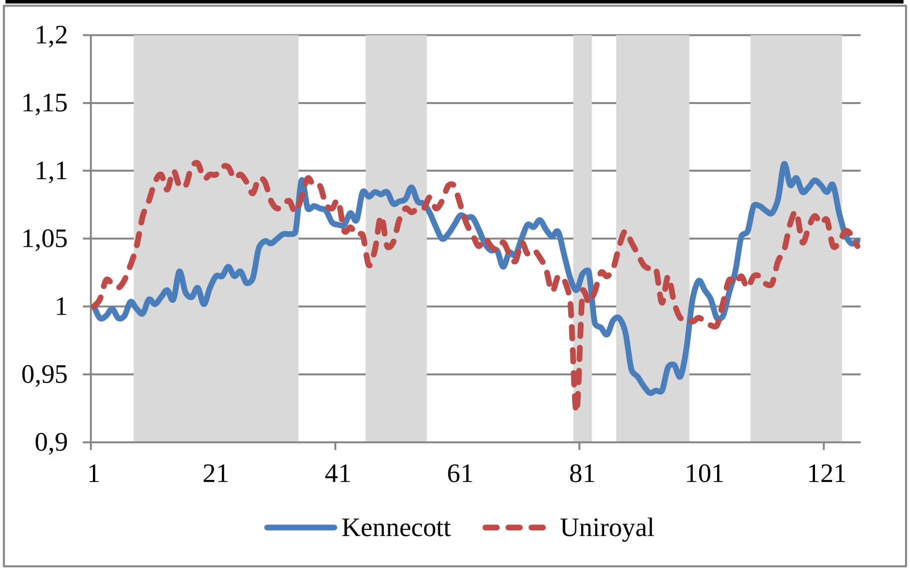

Graph shows the results of a replication of one of the graphs in the Gatev et al. (2006) study. Shaded areas indicate that trading is taking place. The stocks are Kennecott and Uniroyal, and the price series are normalized to 1 in the beginning of the period.

Excess Return Calculation and Transaction Costs

We use the same method as Gatev et al. (2006) when calculating the excess returns of the trading strategy. The trading periods are six months, and trades are always opened and closed solely within the trading period. The trades that are opened and closed out before the end of a trading period because of divergence will have positive cash flows. The trades that are closed out at the last day of a trading period may have either positive or negative cash flows. For each pair it is possible to have multiple cash flows during a trading period if the pairs diverge and later converge multiple times. It is possible to have no cash flows at all if the prices do not diverge by 2 standard deviations at any time during the trading period. The excess return is calculated as the sum of the cash flows during the trading period.

Gatev et al. (2006) define two versions of the excess return calculation:

a) **The excess return on committed capital**: In this version the excess return on committed capital takes the sum of the payoffs over all pairs during the trading period and divides by the total number of pairs in the portfolio. Gatev et al. (2006) state that this version of calculation is conservative because capital is committed to a pair even if it is not traded. They explain that this version takes into account the opportunity cost for hedge funds that has to commit capital to a strategy even if the strategy does not trade.
b) **The excess return on employed capital**: This is the sum of the cash flows divided by the number of pairs that actually open during the trading period. This version of calculation gives more realistic results in terms of hedge funds which use their funds more efficiently. This approach is however less conservative.

We mainly use the committed capital version ("**a**" above) of excess return calculation because it is the more conservative of the two. We do however report the percentages of pairs that actually never trade. Transaction costs are partly accounted for in the analysis, but not fully. Like many other studies, this study rests on an assumption of relatively efficient markets, which would imply for example no borrowing costs. Gatev et al. (2006) partly account for transaction costs by assuming that trades are postponed one day after pairs actually diverge or converge. They compare these implicit transaction costs to actual transaction costs as found in the studies by Petersen & Fialkowski (1994) and Keim & Madhavan (1997) and conclude that the implicit costs are higher than actual empirically observed costs in the same market. In this study, we also delay trading by one day to partly account for transaction costs.

In our analysis, we use daily stock price data from Thomson Reuters Datastream. The implementation is carried out using Stata, with the same code as for the above replication applied to our chosen markets. For a full list of stocks included in the analysis, please see Table A2 in the Appendix. Following Gatev et al. (2006) we only include the most liquid stocks in the analysis, meaning that not all stocks in the indices are used in the analysis. Liquid stocks are defined as stocks that have price data for every single day over the 20-year period January 1994 – December 2013. For further information about the data used in the study, we refer the reader to Table A1 in the

Appendix, which summarizes and briefly describes the data used in the study, alongside the sources from which the data has been retrieved.

5. RESULTS

In this part we report the main findings from the analyses of the three markets. **Table 1** summarizes key output statistics from the analysis. The table is followed by a market-by-market breakdown of the results.

TABLE 1. TRADING STATISTICS

	CAC 40	**DAX 30**	**FTSE 100**
Starting date		January 1, 1994	
Ending date		December 31, 2013	
Months in sample	252	252	252
Trading periods	240	240	240
Mean annualized return	7,61%	4,38%	5,13%
Mean monthly return	0,0063	0,0036	0,0043
Median monthly return	0,0042	0,0033	0,0039
Monthly standard deviation	0,0200	0,0193	0,0179
Sharpe ratio	0,32	0,19	0,24
Skewness	0,52	-0,09	0,39
Kurtosis	1,95	0,77	1,84
Minimum monthly return	-0,0671	-0,0581	-0,0456
Maximum monthly return	0,0793	0,0545	0,0652
Monthly Value at Risk percentiles			
1%	-0,0325	-0,0428	-0,0421
5%	-0,0203	-0,0287	-0,0237
10%	-0,0155	-0,0181	-0,0151
15%	-0,0108	-0,0140	-0,0111
20%	-0,0076	-0,0108	-0,0064
Statistics from full 6-month trading periods			
Number of full trading periods	234	234	234
Average number of openings	1,41	1,31	1,51
(Average number of closings)	1,41	1,31	1,51
Average number of trades over trading period	2,82	2,62	3,01
Number of pairs with excess return > 0	693	595	666
Number of pairs with excess return < 0	409	476	453
Number of pairs with excess return = 0	68	99	51
Winning/Losing ratio	1,69	1,25	1,47
Percent profitable trading periods	62,89%	55,56%	59,52%

Table shows summary statistics from the trading strategy implemented using the three indices CAC 40, DAX 30 and FTSE 100 over the time period January 1994 – December 2013.

CAC 40 Results

The strategy yields an average annualized excess return of 7.61% for the CAC 40, which is the highest return of the three markets. The average monthly return is 0.63% with a standard deviation of 2%, which yields a Sharpe ratio of 0.32. Returns are positively skewed. For a pair of stocks during a full six month trading period positions are opened (and closed) on average 1.41 times. Of trading periods with non-zero total returns 62.89% of trading periods are profitable, which is the highest percentage of the three markets. In 6.07% of the trading periods no trading is taking place because the pairs never diverge by more than 2 standard deviations.

DAX 30 Results

The strategy yields an average annualized excess return of 4.38% for the DAX 30. The average monthly return is 0.36% with a standard deviation of 1.93%, which yields a Sharpe ratio of 0.19. Returns are negatively skewed, something which is unique for this market. For a pair of stocks during a full six month trading period positions are opened (and closed) on average 1.31 times. Of trading periods with non-zero total returns 55.56% of trading periods are profitable, which is the lowest percentage of the three markets. In 8.69% of the trading periods no trading is taking place because the pairs never diverge by more than 2 standard deviations.

FTSE 100 Results

The strategy yields an average annualized excess return of 5.13% for the FTSE 100. The average monthly return is 0.43% with a standard deviation of 1.79%, which yields a Sharpe ratio of 0.24. Returns are positively skewed. For a pair of stocks during a full six month trading period positions are opened (and closed) on average 1.51 times. Of trading periods with non-zero total returns 59.52% of trading periods are profitable. In 4.47% of the trading periods no trading is taking place because the pairs never diverge by more than 2 standard deviations.

The results are generally similar across all three markets. The strategy produces positive excess returns in all markets, volatility is comparable, and trading volume is also similar. Later in the study we perform regression analyses of the returns on common factors to see if the excess returns can be explained by such exposures. A comparison between the excess returns produced by the strategy and the excess returns of the underlying indices over time is presented in **Graph A1** in the Appendix.

Value at Risk

Table 1 shows the monthly Value at risk (VaR) statistics that summarize the quantiles of the empirical distributions of monthly returns. The worst monthly loss during the full January 1994 – December 2013 sample period is 6.7%, 5.8% and 4.5% for the French, German and U.K. market respectively. We use the same interpretation of the VaR statistics as in Gatev et al. (2006). The VaR statistics show that on average, only once in every hundred months did these portfolios lose more than 3.25%, 4.28% and 4.21% for the French, German and U.K. market respectively. The full distributions of monthly returns are shown in **Graph 3**. The graphs are intended to give the reader a more detailed picture of the distributions of returns than what is possible via the summary statistics and the empirical VaR statistics.

GRAPH 3. MONTHLY RETURN DISTRIBUTIONS

CAC 40

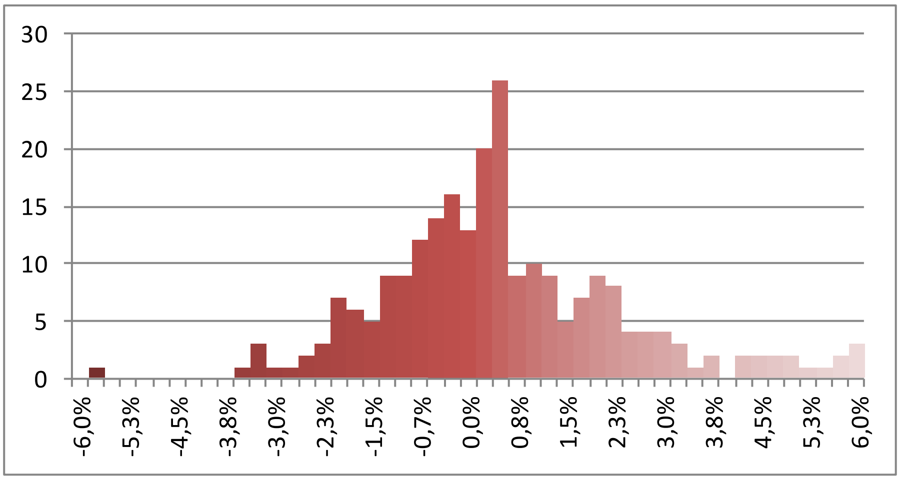

DAX 30

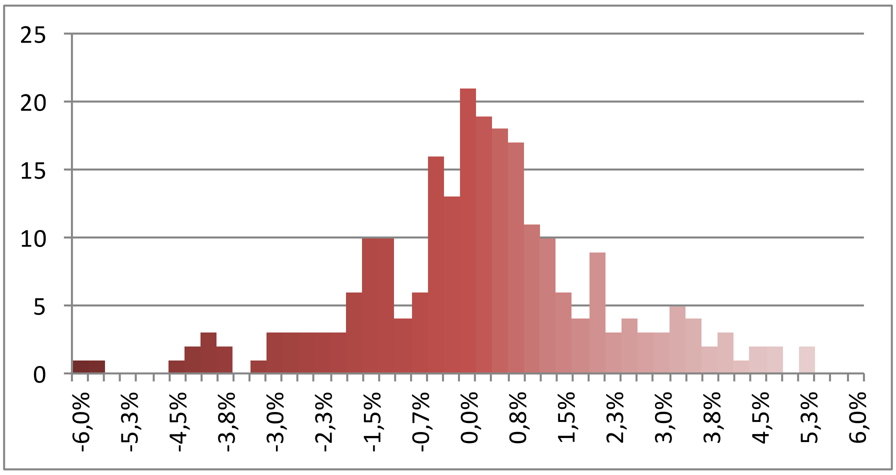

FTSE 100

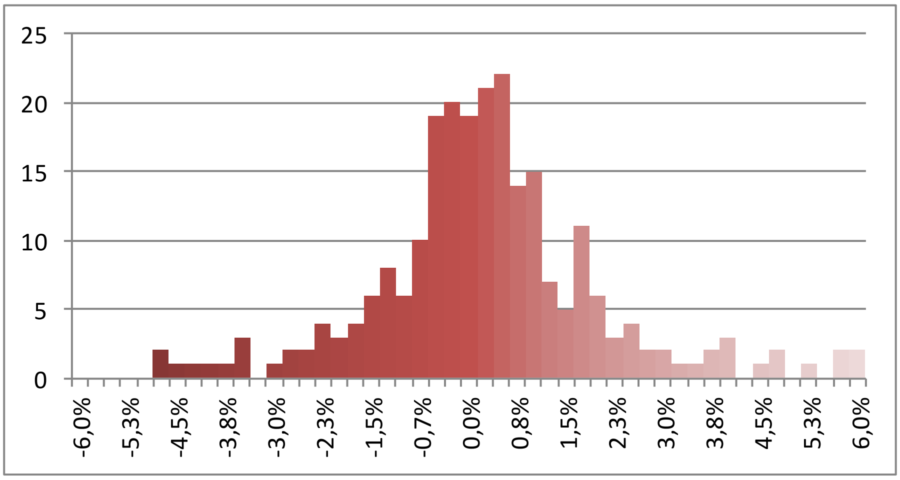

Graphs show the distributions of monthly returns from the strategy. Labels on the x-axis indicate the left hand (lower) side of the bins, the y-axis indicates the frequency.

Exposure to Common Factors

Above we have documented that the strategy produces excess returns in all three markets separately. In this part we perform regression analyses to see if the returns can be explained by exposure to common factors or if the returns are in fact 'alpha', according to common specifications.

CAPM

The first model we use to test for exposure to common factors is the Capital Asset Pricing Model (CAPM). We use the model to see if the returns produced by the strategy can be explained by market exposure. We perform regression analyses of the monthly returns of the strategy on the monthly excess return of the respective underlying index ('market'). The excess return of the market is calculated as the return of the index less the domestic 3-month Treasury bill rate. The output from these regressions is shown in **Table 2**. The results from the regressions show that the returns cannot be explained by market exposure alone, as defined in the CAPM. With the exception of a slightly positive market exposure for the CAC 40, the strategy is market neutral. The lack of market exposure is to be expected since the strategy implies always being equally long and short in the market, which isolates the performance from the direction of the broader market. The intercept (alpha) is significantly positive for all markets using the standard CAPM specification.

In this analysis we find results generally consistent with Gatev et al. (2006). We find that, with one small exception, most of the returns are explained by the intercept in this specification. This implies no significant market exposure and a return that is considered to be almost exclusively alpha according to the CAPM. Logically, the lack of market exposure yields intercepts (alphas) that are positive and highly significant in all three markets. Only in the French market can we document a statistically significant market exposure and only at a 5% significance level. Taken together, we believe that these findings strongly suggest that our pairs trading strategy is mostly market neutral in these three markets.

TABLE 2. CAPM FACTOR REGRESSIONS

	(STRATEGY RETURNS)		
	CAC 40	**DAX 30**	**FTSE 100**
CAC 40	0.0620*	-	-
	(0.0289)	-	-
DAX 30	-	0,0269	-
	-	(0.0221)	-
FTSE 100	-	-	-0,0441
	-	-	(0.0387)
Constant	0.0063***	0.0036**	0.0043***
	(0.0013)	(0.0013)	(0.0012)
R-squared	0,032	0,008	0,011
Dfres	238	238	238

* p<0.05, ** p<0.01, *** p<0.0001

Table shows output from regression analyses of monthly strategy excess returns on monthly excess returns of the respective underlying indices. Columns indicate excess strategy returns and rows indicate excess market returns of the respective indices. Excess market returns are calculated as the change in the index less the 3-month Treasury bill rate.

Fama-French

The second model we use to test for exposure to common factors is the Fama-French model. Again, in this regression analysis we follow the methodology of Gatev et al. (2006). In Gatev et al. (1999, 2006) excess returns are regressed on Fama-French factors to assess the exposure to these common factors. We employ the same type of Fama-French analysis as in Gatev et al. (2006), where two new factors were added to the earlier Gatev et al. (1999) study. The new factors that were added to the traditional Fama-French three-factor model are a reversal factor and a momentum factor. The factors used in our regressions are the European Fama-French factors, since we study European markets. Below we briefly describe the five factors in this extended model. The factors in the model are:

MKT: Market return in excess of the 3-month Treasury Bill rate specific to included countries.

SMB (Small Minus Big): Calculated as the average return on three 'small' portfolios less the average return on three 'big' portfolios, where

$$SMB = \frac{1}{3} * Small\,Value + Small\,Neutral + Small\,Growth - \frac{1}{3} * Big\,Value + Big\,Neutral + Big\,Growth$$

HML (High Minus Low): Average return on two value portfolios minus the average return on two growth portfolios, where

$$HML = \frac{1}{2} * Small\ Value + Big\ Value - \frac{1}{2} * Small\ Growth + Big\ Growth$$

WML (Winner Minus Loser): Equal weighted average of the returns for two winner portfolios minus the average of the returns for two loser portfolios.

$$WML = \frac{1}{2} * Small\ High + Big\ High - \frac{1}{2} * Small\ Low + Big\ Low$$

Reversal = Average return on two low prior return portfolios minus the average return on two high prior return portfolios.

$$Reversal = \frac{1}{2} * Small\ Low + Big\ Low - \frac{1}{2} * Small\ High + Big\ High$$

In Gatev et al. (1999) exposures of pairs trading portfolios to the market excess return are generally small and not significantly different from zero. Also in Gatev et al. (1999) returns are positively correlated with the difference between small and big stocks (SMB), and the difference between value and growth stocks (HML) for certain portfolios of pairs. However, these exposures are not sufficient to explain the total excess returns. In their study, exposure to Fama-French factors explains only around 100bp of the average annual performance of returns, which is a small fraction of the total returns.

As previously mentioned, in Gatev et al. (2006) five factors are used instead of the standard three factors commonly used in Fama-French analysis, and the additional factors are a momentum factor and a reversal factor. They find that a small portion of the excess returns exhibited by the strategy can be attributed to those five factors. Risk adjusted returns are significantly positive and lower than the excess returns by about 10-20bp per month. Pairs trading strategies are market neutral and exposure to the market is small and usually insignificant. Exposures to the two Fama-French factors SMB and HML are not significant and the estimates alternate in sign. Exposure to the momentum factor has a negative sign and exposure to the reversal factor has a positive sign, and these new exposures are statistically significant more than half of the time. To sum up their findings, exposures are not large enough to fully explain the average excess returns from the pairs trading strategy. Moreover, their results are generally in accordance with what is to be expected from a market neutral, contrarian trading strategy.

The Fama-French regression output from our study is presented in **Table 3**. In our Fama-French factor regressions all the intercepts (alphas) are positive and highly significant, which is in accordance with Gatev et al. (2006). We find that the size factor, SMB, is insignificant across all markets, with positive coefficients that are close to zero. The HLM factor alternates in sign, it is negative for the French and German markets but positive for the U.K. market. However, it is found to be insignificant across all markets. The reversal factor is generally insignificant, with one exception at the 5% significance level for the French market.

One difference between our results and those of Gatev et al. (2006) is the significance of the momentum factor, WML. We find all coefficients to be negative and highly significant. Gatev et al. (2006) also found a negative WML exposure, but as previously mentioned it was not always significant. The WML exposure was more consistently significant when Gatev et al. (2006) used the top 20 pairs as a portfolio instead of the usual top 5 pairs, whereas for us it is significant when using the standard top 5 pairs.

We believe that the significance of the momentum factor is in accordance with a contrarian trading strategy, as it implies convergence. It seems that some of the stocks that this particular pairs trading strategy shorts are short term "winners" and that some of the stocks that the pairs trading strategy buys are short term "losers", according to the WML definition. The negative sign of the momentum factor is consistent with our trading strategy, it shows that we are selling "winner" stocks and buying "loser" stocks. This is consistent with a convergent strategy such as pairs trading, as it implies selling short-term "winners" expecting that they will lose value in the near future and buying short-term "losers", expecting that the two will converge.

TABLE 3. FAMA-FRENCH FACTOR REGRESSIONS

	CAC 40	DAX 30	FTSE 100
MKT	-0,0001	-0,0001	-0.0008*
	(0.0004)	(0.0003)	(0.0003)
SMB	0,0008	0,0004	0,0002
	(0.0008)	(0.0006)	(0.0007)
HML	-0,0002	-0,0007	0,0006
	(0.0007)	(0.0007)	(0.0005)
WML	-0.0013***	-0.0012***	-0.0010***
	(0.0003)	(0.0003)	(0.0003)
Reversal	0.0009*	0,0007	0
	(0.0004)	(0.0004)	(0.0004)
Constant	0.0074***	0.0051***	0.0055***
	(0.0014)	(0.0013)	(0.0014)
R-squared	0,095	0,077	0,093
Dfres	235	235	235

* $p<0.05$, ** $p<0.01$, *** $p<0.0001$

Table shows output from regression analyses of monthly strategy excess returns on Fama-French monthly European factors (the reversal factor is global as no European version is available). MKT denotes the excess market return, SMB is a size factor, HML is a value factor, WML is a momentum factor and Reversal is a short-term reversal factor.

Volatility and Liquidity

Lastly we perform regression analyses of the monthly excess returns from the strategy on monthly changes in volatility of the respective underlying index and a liquidity factor. To keep the results comparable to Gatev et al. (2006) and given the general lack of explanatory power in the earlier regressions we perform this regression on a stand-alone basis. Using changes in volatility indices as a proxy for volatility exposure is common in similar studies, the choice of the liquidity factor is based mainly on Chen et al. (2009). The output can be found in **Table 4**. We find that the excess returns produced by the strategy are significantly positively related to the volatility of the underlying indices as indicated by the coefficients in the table. This significance holds true for all markets. For two of the three markets, the German and the U.K. markets, there are significant negative relationships between excess returns and liquidity. For the German market however, the coefficient is very close to zero. The results suggest that lower liquidity in these markets imply higher excess returns from the trading strategy. For the third market, the French market, there is no significance in the liquidity coefficient. Lack of liquidity has, as previously mentioned, been suggested as a possible explanation for the profitability of contrarian trading strategies. The results will be discussed further later in the study.

TABLE 4. VOLATILITY AND LIQUIDITY FACTOR REGRESSIONS

	CAC 40	**DAX 30**	**FTSE 100**
CACVOLI	0.0007***	-	-
	(0.0002)	-	-
VDAXNEW	-	0.0004*	-
	-	(0.0002)	-
VFTSEIX	-	-	0.0006**
	-	-	(0.0002)
LIQ	0.0223	-0.0000**	-0.0787*
	(0.0627)	(0.0000)	(0.0385)
Constant	-0.0088*	-0.0047	-0.0062
	(0.0042)	(0.0037)	(0.0037)
R-squared	0.091	0.031	0.122
Dfres	165	237	164

* p<0.05, ** p<0.01, *** p<0.0001

Table shows output from regression analyses of monthly strategy excess returns on monthly changes in volatility of the respective underlying stock index and a liquidity factor. Volatility data is retrieved from Datastream, and CACVOLI is the volatility of the CAC 40 index, VDAXNEW is the volatility of the DAX 30 index and VFTSEIX is the volatility of the FTSE 100 index. LIQ is the Pastor-Stambaugh traded liquidity factor as used in Chen et al. (2009). The differences in degrees of freedom reflect the fact that not all volatility indices have data all the way back to January 1994.

Return Characteristics over Time

Gatev et al. (1999, 2006) find that returns from pairs trading are decreasing over time in the U.S. market. Here we try to answer if the returns from pairs trading are decreasing over time in our study as well. Decreasing excess returns were mentioned both in Gatev et al. (1999, 2006) and Do & Faff (2010). **Graph 4** plots the monthly returns over time. The calculated cumulative excess returns for the first half of the period are 87%, 56% and 61% for France, the U.K. and Germany respectively. The cumulative returns for the second part of the period are 65%, 46% and 26% for the same respective markets. The decline of the excess returns can be further observed if we exclude periods of financial instability, such as the recent financial crisis, and calculate the cumulative returns for the first three years and for the last three years in the sample. The cumulative returns for the first three years of the sample are 10%, 26% and -12% for France, the U.K. and

Germany respectively, and the cumulative returns for the last three years are 0%, 1% and -4% for the same respective markets.

Gatev et al. (2006) also compare the first half of their sample to the second half of their sample and find similar results. They provide two possible explanations to this downward trend in excess returs. The first one is that transactions costs are assumed to be higher in the early part of the sample period and that declining transaction costs over time may have attracted more relative value equity arbitrage to the market. They further believe that this declining trend of excess returns likely has increased as a consequence of the introduction of more sophisticated trading technologies and networks. This could be a plausible explanation of the seemingly declining trend of excess returns in our case as well, even though the time periods studied differ. The other explanation they provide with regard to the downward trend is that there has been an increase over time of hedge funds with “market neutral”, “relative value” and “arbitrage” strategies. Also this is a possible explanation to our findings.

GRAPH 4. MONTHLY RETURNS OVER TIME

CAC 40

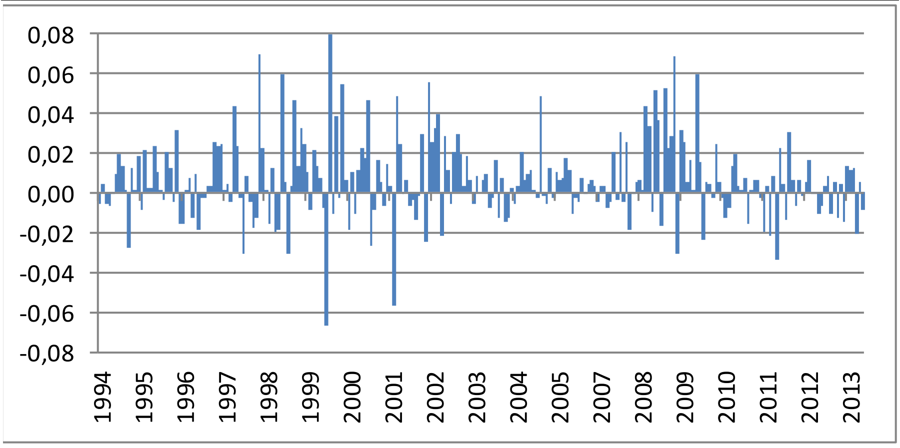

DAX 30

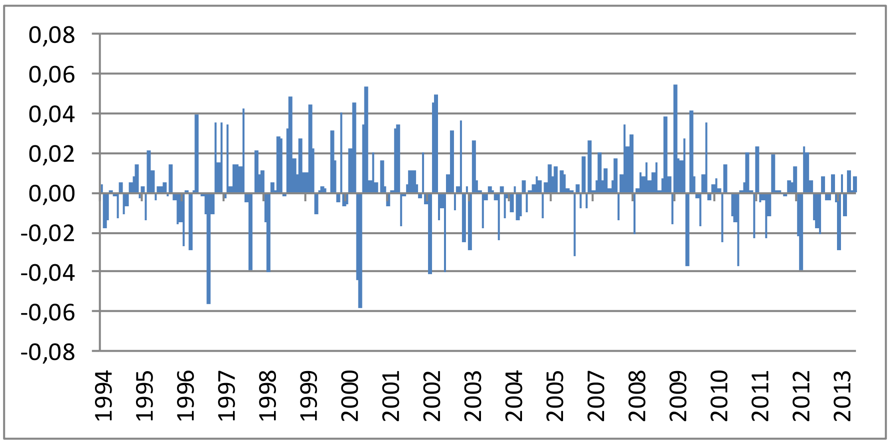

FTSE 100

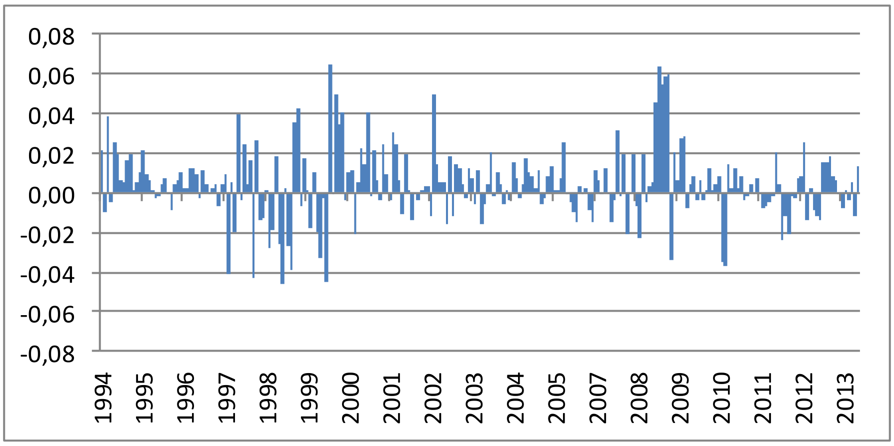

Grahps show monthly returns from the strategy over the full time period January 1994 - December 2013.

6. DISCUSSION AND INTERPRETATION

As we have shown, we find that the strategy ‘works’, in the sense that it produces positive excess returns. Moreover, the fact that we implemented the strategy in several markets means that we can provide out-of-sample tests for each market, and we have shown that the results are not market specific but rather general in nature. The results are generally in accordance with previous research on the topic. Most importantly, they are in line with the results of Gatev et al. (2006), which is the study our paper is fundamentally based on. Gatev et al. (2006) found significant excess returns in the U.S. market, which they later proved to be independent of exposure to common factors, mainly according to an extended version of the Fama-French three-factor model. The exposures to the various factors in the specification are not exactly alike ours in all cases, but broadly they are very much in accordance with each other.

We also find that our results are in line with what is to be expected from a contrarian trading strategy such as pairs trading. The returns are generally market neutral, and negatively correlated to momentum exposures such as the Fama-French WML factor. The implications of these findings are several. Perhaps most importantly it speaks to the accuracy of the implementation of the analysis; if the results are in line with expectations it lends credibility to the study. Moreover, the similarities with the Gatev et al. (2006) study of the U.S. market show that pairs trading as a strategy still works when using more recent data, and it also shows that their results can be extended to other markets. A European focus is more or less unique to this study, and we believe this focus, alongside the significance of the results, provides a valuable addition to existing research within the fields of convergent and contrarian trading broadly, and pairs trading specifically.

As we have shown, the returns of the strategy cannot be explained by common factor models, specifically using the CAPM and an extended version of the Fama-French model. As described in the literature review, contrarian strategies have been shown to be profitable in previous studies. However, there is no agreed upon explanation for as to why these persistent returns exist; each study provides more or less a unique explanation. When reviewing these previous explanations we find that some of them are similar, and when ignoring minor differences it is possible to reduce a large number of unique explanations to a few broader explanations. Generally, these relate to the overreaction hypothesis, differences in reactions between firms to firm-specific shocks and lack of liquidity.

As our study is based on Gatev et al. (2006), our main focus has been on creating a study in which the results can be compared to theirs. As such, we have first focused on the accuracy of the algorithm by favourably replicating a subset of their results. Then, we have applied the trading algorithm to our chosen European markets, and interpreted the results so as to be able to compare the outcome to theirs. As described above, our results are generally comparable to theirs. In their study they successfully prove that the strategy is profitable, and that the strategy produces significant alpha according to common factor models. They do not however find a single explanation to as to why these returns exist, or to why the returns are so persistent over time. As we have noted previously, this is a common outcome in previous research on pairs trading, and contrarian trading in general. Commonly, authors can show that a contrarian strategy produces excess returns over time, but the explanations for why these returns exist differ from study to study and are usually hard to prove.

This is broadly the outcome for our study as well; there are several possible explanations to why the strategy is profitable. Gatev et al. (2006) try to find a common factor that explains the performance of the strategy, but end their study with the phrase ‘A further examination of this common factor and its link to the profitability of pairs trading is an important question for future research’. Given that our study is based on the Gatev et al. (2006) study, we have as mentioned above mainly chosen to focus on results that can be compared to theirs, and thus we believe that a deeper explanation of what causes the returns is outside the scope of this paper. However, we fully agree that it is still an important question for future research, and we will provide a brief discussion with regard to this question.

Lo & MacKinlay’s (1990) theory that contrarian trading returns can be explained by differences in reaction times between stocks is one possible explanation. Partly, it is an explanation that theoretically fits the behavior of a pairs trading strategy. If there are differences in reaction times between stocks to firm-specific shocks, one would expect stocks to diverge more than what is justified in the long run over short time periods. This is exactly what a pairs trading strategy takes advantage of: temporary differences in security prices that are assumed to revert back to what has historically been normal. Theory aside, what is also relevant to this study is their quote ‘a portfolio strategy that sells "winners" and buys "losers" can produce positive expected returns, even if no stock's returns are negatively autocorrelated as virtually all models of overreaction imply'. As we have shown, we find a highly significant negative exposure to the Fama-French momentum factor, which is in accordance with their description of selling "winners" and buying "losers". This means that, although we cannot prove that the explanation of the returns in our study is explicitly because of differences in reaction times between stocks, it makes for a likely explanation, and one that fits both the theoretical framework and the empirical results of our study.

The fact that the explanation put forth by Lo & MacKinlay (1990) seems to fit our findings does not however rule out the possibility of other causes for the returns. As we have shown, the returns are positively correlated to changes in the volatility of the underlying indices. This result is consistent with several of the existing explanations of contrarian returns, and as such it is hard to tie the results to one explanation exclusively. For example, it is reasonable to assume that overreactions in the market will lead to higher volatility since an overreaction followed by a correction would create more volatility than a 'direct' move to the 'correct' value of a security. As such, the fact that returns are positively correlated to volatility could support the 'overreaction' explanation of the returns, and it would also be theoretically justified.

Furthermore, the notion that lack of short-term market liquidity causes short-term divergences between security prices could be a likely explanation for the returns. A lack of short-term liquidity implies that prices tend to move more than what is fundamentally justified in the short-term, to revert back over time as long-term liquidity is sufficient. This explanation fits the strategy as it is based on taking advantage of short-term divergences. Also this explanation can be justified by the results, most notably by the documented negative relationship between the excess returns from the strategy and the liquidity factor, but also by the negative correlation between the strategy's returns and the momentum factor. The findings are also in accordance with Nagel (2012), who found a positive relationship between reversal returns and volatility. Nagel (2012) interpreted the returns as compensations for liquidity provision, and this interpretation fits our results as well.

We cannot state for certain which explanation is the correct one. Like Gatev et al. (2006) we note that it is an important question, but leave it to future research.

Extensions to the Study

Finally before stating our conclusions we provide some suggestions for future research, and possible extensions to this study that have been deemed to be outside the scope of this paper. One relevant extension is a more detailed inclusion of the costs to trade. Generally, in previous research as in this paper, studies commonly rely on an assumption of relatively efficient markets. This implies that factors such as transaction costs commonly are excluded on a per-trade basis and sometimes accounted for by a more generalized assessment of total trading costs. As described previously, the study partly accounts for transaction costs by postponing trades by one day after either divergence or convergence, but a more detailed approach would likely provide some additional depth to the analysis. We have deemed this to be outside the scope of this paper partly because it is based on the Gatev et al. (2006) approach, and partly because our approach with three markets would make it challenging since transaction costs would be different in each market.

Another possible extension would be to extend the specification of the strategy to include more (or fewer) pairs than the standard five. It is clear that increasing the number of pairs reduces the volatility of the portfolio since the volatility of individual pair returns are averaged over a larger number of pairs. It is however likely that increasing the number of pairs would affect the return of the strategy, since less efficient pairs would have to be included in the analysis. This could lead to interesting risk-return discussions, and optimal strategies could be found and compared across different markets. A related extension is to change the frequency of trading, for example by moving from daily trading to intraday trading.

Finally, further research with regard to the causes of contrarian returns is still needed. As presented, there are a variety of opinions with regard to why contrarian strategies like pairs trading produce excess returns and significant alphas. The continued search for common factors that can explain the phenomenon is key to the understanding of the topic, and as such an important and logical extension to the current body of research.

7. CONCLUSION

In this study we have applied a known pairs trading strategy in three separate European equity markets over the 20-year time period January 1994 – December 2013. We conclude that the strategy produces excess returns and significant alpha in all three markets, and that results generally are consistent across all markets. The results are generally in accordance with previous studies on the topic, most importantly with the Gatev et al. (1999 & 2006) study that this study is based on. The results from the study are also in accordance with what is to be expected by a pairs trading strategy, as evidenced among other things by the proven market neutrality of the returns.

The performance of the strategy cannot fully be explained by exposure to common factors, but the study does prove significant exposure to a momentum factor. We cannot conclude with certainty what causes the profitability of the strategy. Several explanations from previous studies are possible, and it may well be the case that the results are caused by a mixture of these, and not one explanation exclusively. We document that excess returns from the strategy are significantly related to both volatility and liquidity, and these empirical findings are in accordance with some of the existing explanations for the profitability of contrarian trading strategies.

The study has a European focus, and the strategy has been implemented in markets that are unique to the current body of research. Given this, we believe that our study constitutes a valuable contribution to the existing research on pairs trading specifically and contrarian and convergent trading broadly. Moreover, we have documented several statistically significant results that generally are in accordance with previous research on the topic. The use of multiple markets shows that the results are not market specific but rather general in nature, and we believe that this adds credibility and robustness to the reported results.

REFERENCES

Andrade, S., Di Pietro, V., & Seasholes, M. (2005). Understanding the profitability of pairs trading. *Unpublished working paper, UC Berkeley, Northwestern University*.

Barinov, A. (2009). *Turnover: Liquidity or Uncertainty?*. Working Paper, University of Georgia.

Bondt, W. F., & Thaler, R. (1985). Does the stock market overreact?. *The Journal of finance*, *40*(3), 793-805.

Casselman, B. (2013) Observations on the death of trend following. *http://www.futuresmag.com/2013/06/10/observations-on-the-death-of-trend-following*

Chen, H. J., Chen, S., & Li, F. (2009). Empirical investigation of an equity pairs trading strategy. *Shaojun Jenny and Li, Feng, Empirical Investigation of an Equity Pairs Trading Strategy (February 29, 2012)*.

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

Do, B., Faff, R., & Hamza, K. (2006, May). A new approach to modeling and estimation for pairs trading. In *Proceedings of 2006 Financial Management Association European Conference*.

Elliott, R. J., Van Der Hoek, J., & Malcolm, W. P. (2005). Pairs trading.*Quantitative Finance*, *5*(3), 271-276.

Engle, R. F., & Granger, C. W. (1987). Co-integration and error correction: representation, estimation, and testing. *Econometrica: journal of the Econometric Society*, 251-276.

Gatev, E. G., Goetzmann, W. and Rouwenhorst, K.(1999) Pairs Trading: Performance of a Relative Value Arbitrage Rule. *Unpublished Working Paper, Yale School of Management*.

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

Grossman, S. J., & Miller, M. H. (1988). Liquidity and market structure. *the Journal of Finance*, *43*(3), 617-633.

Hansell, S., (1989) Inside Morgan Stanley's Black Box, *Institutional Investor*, May, 204.

Jegadeesh, N. (1990). Evidence of predictable behavior of security returns. *The Journal of Finance*, *45*(3), 881-898.

Jegadeesh, N., & Titman, S. (1995a). Short-horizon return reversals and the bid-ask spread. *Journal of Financial Intermediation*, *4*(2), 116-132.

Jegadeesh, N., & Titman, S. (1995b). Overreaction, delayed reaction, and contrarian *profits. Review of Financial Studies, 8(4), 973-993.*

Keim, D. B., & Madhavan, A. (1997). Transactions costs and investment style: an inter-exchange analysis of institutional equity trades. *Journal of Financial Economics*, *46*(3), 265-292.

Lehmann, B. N. (1990). Fads, martingales, and market efficiency. *The Quarterly Journal of Economics*, *105*(1), 1-28.

Lo, A. W., & MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction?. *Review of Financial studies*, *3*(2), 175-205.

Mun, J. C., Vasconcellos, G. M., & Kish, R. (2001). The contrarian/overreaction hypothesis: An analysis of the US and Canadian stock markets. *Global Finance Journal*, *11*(1), 53-72.

Nagel, S. (2012). Evaporating liquidity. *Review of Financial Studies*, *25*(7), 2005-2039.

Nath, P. (2003). High frequency pairs trading with us treasury securities: Risks and rewards for hedge funds. *SSRN eLibrary*.

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

Petersen, M. A., & Fialkowski, D. (1994). Posted versus effective spreads: Good prices or bad quotes?. *Journal of Financial Economics*, *35*(3), 269-292.

Van Dam, L. (2012). Why I won’t teach pairs trading to my students. *MarketWatch*, http://www.marketwatch.com/story/why-i-wont-teach-pair-trading-to-my-students-2012-10-01

Vidyamurthy, G. (2004). *Pairs Trading: quantitative methods and analysis* (Vol. 217). John Wiley & Sons.

APPENDIX

TABLE A1. DATA USED IN THE ANALYSIS

NAME	DESCRIPTION	SOURCE
CACVOLI	CAC 40 VOLATILITY INDEX	Datastream
DAXINDX	DAX 30 PERFORMANCE - PRICE INDEX	Datastream
FRCAC40	FRANCE CAC 40 - PRICE INDEX	Datastream
FTSE100	FTSE 100 - PRICE INDEX	Datastream
TRBD3MT	TR GERMANY T-BILLS BID YLD 3M (E) - RED. YIELD	Datastream
TRFR3MT	TR FRANCE T-BILLS BID YLD 3M (E) - RED. YIELD	Datastream
TRUK3MT	TR UK T-BILLS BID YLD 3M (£) - RED. YIELD	Datastream
VDAXNEW	VDAX-NEW VOLATILIT INDEX	Datastream
VFTSEIX	FTSE 100 VOLATILITY INDEX	Datastream
LIQ_V	Traded liquidity (Pastor & Stambaugh)	http://faculty.chicagobooth.edu/
Europe_Factors	Fama/French European Factors	http://mba.tuck.dartmouth.edu/
INTGSTDEM193N	Interest Rates, Government Securities, Treasury Bills for Germany	St. Louis Fed
KENNECOTT	Kennecott stock price data, permno = 12706	WRDS
UNIROYAL	Uniroyal stock price data, permno = 14912	WRDS

Table summarizes the data used in the analysis, alongside short descriptions and the sources from where the data was retrieved.

TABLE A2. STOCKS INCLUDED IN THE ANALYSIS

CAC 40	**DAX 30**
ACCOR	ALLIANZ
AIR LIQUIDE	BASF
ALCATEL-LUCENT	BMW
AXA	BAYER
BOUYGUES	BEIERSDORF
CAP GEMINI	COMMERZBANK
CARREFOUR	CONTINENTAL
DANONE	DEUTSCHE BANK
ESSILOR INTL.	E ON
KERING	FRESENIUS
L'OREAL	HEIDELBERGCEMENT
LAFARGE	HENKEL PREF.
LVMH	K + S
MICHELIN	LINDE
PERNOD-RICARD	DEUTSCHE LUFTHANSA
PUBLICIS GROUPE	MUENCHENER RUCK.
SAFRAN	RWE
SAINT GOBAIN	SAP
SANOFI	SIEMENS
SCHNEIDER ELECTRIC	THYSSENKRUPP
SOCIETE GENERALE	VOLKSWAGEN PREF.
SOLVAY	
TOTAL	
UNIBAIL-RODAMCO	
VALLOUREC	
VINCI	
VIVENDI	

FTSE 100

ABERDEEN ASSET MAN.	PEARSON
AMEC	PERSIMMON
ANGLO AMERICAN	PRUDENTIAL
ANTOFAGASTA	RECKITT BENCKISER GROUP
ASHTEAD GROUP	REED ELSEVIER
ASSOCIATED BRIT.FOODS	REXAM
AVIVA	RIO TINTO
BABCOCK INTL.	ROLLS-ROYCE HOLDINGS
BAE SYSTEMS	ROYAL BANK OF SCTL.GP.
BARCLAYS	ROYAL DUTCH SHELL B
BG GROUP	RSA INSURANCE GROUP
BP	SAGE GROUP
BRITISH AMERICAN TOBACCO	SAINSBURY (J)

BRITISH LAND	SCHRODERS
BT GROUP	SEVERN TRENT
BUNZL	SMITH & NEPHEW
CAPITA	SMITHS GROUP
DIAGEO	SSE
GKN	STANDARD CHARTERED
GLAXOSMITHKLINE	TATE & LYLE
HAMMERSON	TESCO
HSBC HDG. (ORD $0.50)	TRAVIS PERKINS
IMI	TUI TRAVEL
ITV	TULLOW OIL
JOHNSON MATTHEY	UNILEVER (UK)
KINGFISHER	UNITED UTILITIES GROUP
LAND SECURITIES GROUP	VODAFONE GROUP
LEGAL & GENERAL	WEIR GROUP
MARKS & SPENCER GROUP	WHITBREAD
MEGGITT	WOLSELEY
MORRISON(WM)SPMKTS.	WPP
NEXT	

Table shows all the companies that are included in the analyses of the respective stock indices. Only the most liquid companies in the indices are included in the analyses, which is why the total number of stocks used are fewer than the index totals.

GRAPH A1. PERFORMANCE CHARTS

CAC 40

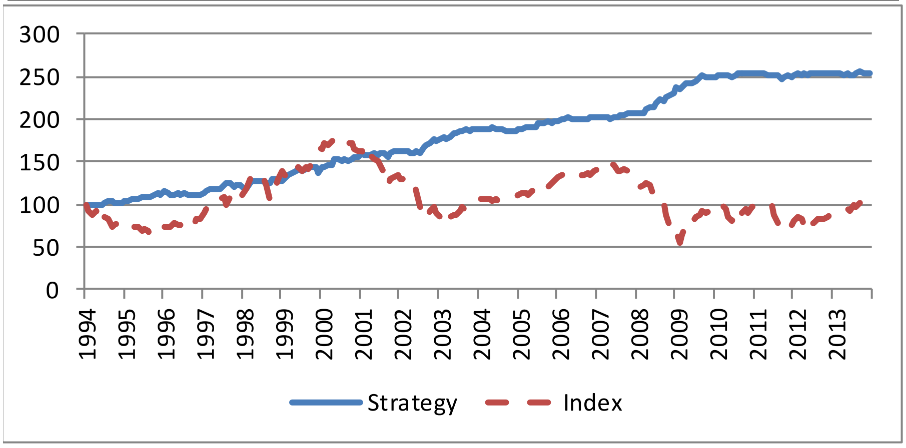

DAX 30

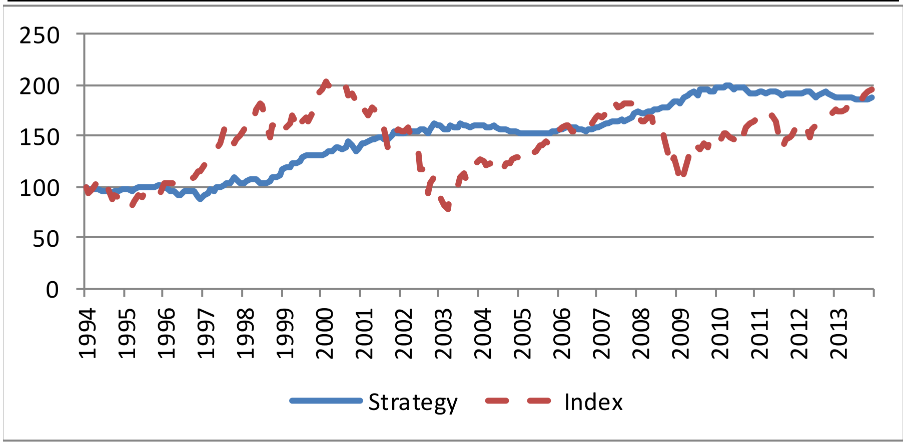

FSTE 100

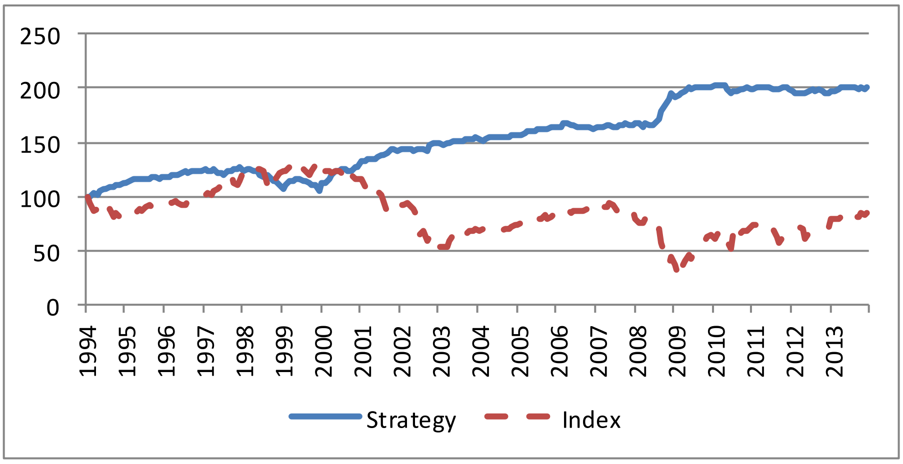

Graphs show performance of the strategy as compared to the underlying excess return of the respective index as measured by the change in the value of the index less the return on the domestic 3-month treasury bill.