Stockholm School of Economics

Department of Finance

CAN COMPUTERS BEAT THE MARKET?

– testing pairs trading using a simulated trading strategy on the Stockholm Stock Exchange

Aaron Ibbotson and Johan Nygren

Abstract

The purpose of this study is to increase the understanding of quantitative trading and to contribute to the general knowledge relating to market efficiency. In current paper we develop and simulate a popular trading methodology known as pairs trading. We simulate trading of over 30000 possible pairs during 10 years; specifically testing whether the excess returns previously achieved applying a similar trading strategy on US shares by Gatev, Goetzmann and Rouwenhorst (1998), are possible to achieve on Swedish stocks. In the study we do not achieve positive returns using a basic trading simulation, predominantly due to large losses incurred during the IT bubble of 1999. However, even excluding 1999 leaves the total return only *just* positive and not significantly so. Our study indicates that the methodology used to find pairs to trade could be more important than the trading strategy applied, which is why we believe that the most interesting area for further research on pairs trading relates to the choice of pairs.

Tutor: Assistant Professor Joel Reneby

Table of contents

1 Introduction and objectives

"Fund teams put faith in quantitative investing" (Financial times, Front page, August 14th, 2006)

"Quantitative money management has grown by 20% in recent years, twice the speed of the rest of the industry. The world's biggest money manager is Barclays Global Investors, which edged past Fidelity and Capital Group with an entirely quantitative investing method, which uses computer models rather than people to make trading decisions. [...] Three of the 10 biggest hedge funds in the world are purely quantitative. [...] Bon Jones, who heads Goldman Sachs' quantitative division, says: 'Quantitative strategies have simply delivered better and more consistent performance'"

1.1 Quantitative investing

Over the last ten years, an ever-increasing pool of investment funds have been transferred from being managed by investment professionals, whom to a large extent take decisions based on their view of a stock or group of stocks, to being managed using purely quantitative methodologies. A computer program will give signals when to buy and sell given certain predetermined criteria, and will today often also execute the actual trade.

In this research paper we aim to investigate the characteristics of one particular and within the investment community common quantitative investment strategy known as *pairs trading* (henceforth PT). We use a widespread definition of the investment strategy PT as "trading one financial instrument or basket of financial instruments against a second financial instrument or basket of financial instruments - Long one and Short the other"[1].

[1] As defined by PhD Russell Wojcik in a presentation at the Illinois Institute of Technology, see page on http://www.futuresindustry.org/downloads/Audio/Companion/Three-812.pdf [2006-12-06]

We have chosen to analyse the characteristics of PT using a simulated trading of shares on the Stockholm Stock Exchange. We will further narrow down our research to one specific area within PT, namely self-funding and market neutral equity PT; i.e. in our simulated trading we will only trade pairs of shares, and our long and short positions will be of equal size, making the trading strategy close to market neutral[2] and self-funding, the investment is zero. We will hence only carry negligible systematic risk and commit no capital, making any statistically significant profits contrarian to the theory of efficient capital markets.[3] We trade using a straight forward trading rule based exclusively on historical price and volume information.

An example of a market neutral equity trading strategy would be to go long Skr.1000 of Ericson and fund it by a Skr.1000 short position in Nokia. The strategy is profitable when the stock you buy outperforms the one you sell, adjusted for transaction costs. PT does not require any initial investment from the trader/investor (apart from transaction costs) since the short position is funding the long position.

PT, either in its pure form or various closely related variations thereof, is today a commonly used quantitative investment strategy. However, traditional financial theory (e.g. Fama 1991; Fama and French 1992), teaches that markets are efficient, at least with regard to the weak form of market efficiency, and hence it should not be possible to use past returns as a predictor of future excess returns. This traditional view has been challenged in some empirical studies. Jegadeesh and Titman (1993) find that buying winners and selling loosers generates significant risk adjusted excess returns on a 3-12 months time horizon, to mention one of many similar studies.

In a study of PT based on US shares between the years 1962-1997, Gatev, Goetzmann and Rouwenhorst (1998) achieve annual excess returns of around 12%. Even after

---

[2] Our aim is to take market neutral positions. However, individual pairs will occasionally have significantly different Beta's, making the position exposed to market risk.

[3] See for instance Fama, E. F. and K. R. French. 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance* 47:427-465. for one definition of weak form of market efficiency.

adjusting for transaction costs and other possible institutional factors such as the bid ask bounce, they achieve significant positive excess returns.

1.2 Research objective

The purpose of this study is twofold: to increase the understanding of PT and to contribute to the general knowledge relating to market efficiency. In current paper we aim to test the same trading methodology as Gatev, Goetzmann and Rouwenhorst (ibid. and henceforth GGR) apply, but expand the dataset from 2000 pairs to, at the most, over 30000 possible pairs. Furthermore we have included data up to June 2006, accounting for a period when computer-based trading has become more widely adopted. Finally, in addition to testing the trading methodology applied by GGR, we aim to make some adjustments to GGR's trading strategy to test for what we believe to be a more accurate description of how investment professionals use PT.

Specifically we will test whether the excess returns previously achieved applying a quantitative trading rule on US shares, are possible to achieve on Swedish stocks. We apply a few straightforward trading rules based on our best understanding of how investment professionals make use of PT. Firstly we aim to see if excess returns are at all possible to achieve, secondly we want to test whether those excess returns have changed as the investment community has become more computerised, and whether there are certain segments of the stock market or certain time periods when PT works better.

2 Method

In this section of the paper we discuss choice of data in our study, preparation of data and extensively our trading methodology. We have developed a proprietary trading program in VBA (Visual Basic for Application) in Excel. The trading program has been developed to have maximum degree of flexibility with regard to both input and output factors. In its entireness the trading program consists of over 5000 lines of VBA code and the development of this program makes up a large part of the empirical work behind this study.

2.1 Choice of data

We have chosen to work with all stock listed on the Swedish stock exchange main list, OMX all share over the 10 year period 1996-2005. Most of the published studies (e.g. GGR in section 1.2) within this field have used US shares. We use closing prices, bid/ask spreads, turnover and total return for each stock. The data provider we are using is Bloomberg. We use Swedish data and this time period for the following three main reasons.

- **More interesting:** Having general knowledge about the Swedish market ourselves as well as access to practitioners trading on the Swedish stock exchange we felt that to do the study on Swedish stocks would be more interesting for us as well as more useful for the Industry practitioners who use Swedish stocks.
- **Contributing to research**: As far as we know no empirical study of this magnitude and scope has been done on the Swedish stock market.
- **Suitable market:** The Swedish stock exchange has a few characteristics that we believe make it suitable for conducting this type of study. Firstly, the Swedish stock

exchange is among the most liquid stock exchanges in the world[4]. Secondly, a wide range of industries are represented on the stock exchange, which guarantees that our study represents a variety of sectors and not only one sector / industry. Thirdly, all stocks that are listed on the Stockholm Stock Exchange are traded in the same currency which makes it easier to set up the pairs without having to hedge out currency effects. Finally, Sweden currently also has a significant presence of hedge funds and day traders using fully- or semi-quantitative trading rules to trade. We believe that the 10-year period from 1996-2006 covers the transition from little use of computer-based quantitative trading rules to extensive use of the same. In addition we wanted to include as up-to-date data as possible (June 2006) but due to the scope of this study had to limit ourselves to 10 years of data.

An alternative research design would have been to look at a longer time period but with fewer shares. Although we considered that, we chose to use a dataset as extensive as possible to avoid problems with data mining and survivorship bias in our analysis. We could also have chosen to do a cross-country comparison, something we considered but chose not to do due to time limits of this research paper. Finally there are several other stock exchanges with similar liquidity and computerisation to the Swedish Stock exchange, and while we could have chosen to work with any of them, we chose the Swedish stock exchange mainly based on our origin.

2.2 Preparation of data

Before running our trading strategies on the chosen dataset we made a number of alterations to it, in order to match our objectives. Firstly we extracted a list of all stocks at the Swedish stock exchange, from its webpage. Due to different spelling on some companies, we converted our list of stocks into Standard ISIN codes. From these codes we were then able to convert into Bloomberg shortcodes, which is the standard system used by Bloomberg when retrieving data from their servers.

---

[4] See for instance a study by Dey, M. K. and S. Flaherty. 2005. "Stock Exchange Liquidity, Bank Credit, and Economic Growth." where the Stockholm Stock Exchange screens as the most liquid after the London Stock exchange and the Switzerland Stock Exchange.

Having created a list of stocks in Bloomberg shortcodes, we modified our sample set according to the following rules.

- The universe was checked for stocks not listed in Skr., which would make our tests vulnerable to currency fluctuations.
- All stocks in our sample were assigned to an industry segment according to SIC standard industry classifications used by Bloomberg. We could have used some other classification, but SIC is widely available to most practitioners.
- For companies that have several classes of stocks listed at the Swedish stock exchange, we eliminated the share class with the least liquidity. This was done by simply summing total turnover on each class over the entire trading period 1996-2006. The share class with the least turnover was then removed from our stock sample.

These modifications gave us our final universe of stocks 267 stocks, making the maximum total numbers of pairs possible to form 35511. The number of stocks that qualified for each trading period increased over time as stocks got listed on the Stockholm Stock Exchange.

Unfortunately using Bloomberg we did not have access to historical price data of de-listed stocks. Our data-sample thus suffers from survivorship bias.[5] Considering that our simulated trading strategy is market neutral, the issue with survivorship bias should not be substantial but is never the less an important shortfall of our study. We have not tried to quantify the impact of survivorship bias on our results due to the scope of this research paper. We do note, however, that the impact on our returns due to the issue with survivorship bias, is more likely to be negative than positive. There should be an overrepresentation of consistent underperformers of the de-listed stocks. And since our simulated trading strategy will on average loose money on consistent underperformers (it will suggest buying them) the total impact of with regard to survivorship bias is likely to be negative.

[5] For a more detailed discussion on survivorship bias, please see Brown, S. J., W. Goetzmann, R. G. Ibbotson, and S. A. Ross. 1992. "Survivorship bias in performance studies." *Review of Financial Studies* 4:553-580.

From Bloomberg we obtained historical price data for each stock in our sample. We gathered the closing price, open price, bid (at close), ask (at close), high, low, VWAP (volume weighted average price) and volume for each stock. The data was stored in an SQL database, in our case ACCESS.

The data we use is daily data, historically adjusted for dividend, rights issues and stock splits. Due to the risk of data error in our dataset provided by Bloomberg, we conducted tests to the data sample checking that the dividend adjustment provided by Bloomberg was stated correctly at Exdiv date and that the amount of cash dividend stated was within a reasonable range of the price movements of the stock at the Exdiv dates. The final adjusted data sample was then recalculated so that we also stored adjusted daily returns on each stock which we needed for the formation of our pairs.

2.3 Trading Methodology

Above we discussed the choice and preparation of data. In the following section we will discuss our trading methodology in some detail. The trading is conducted in three different steps; firstly pairs formation when we form up to 30000 pairs, secondly the pairs selection, when we choose what pairs to enter into our trading system, and finally trading, when we simulate trading of the chosen pairs.

The PT strategy we have chosen to test is a pure quantitative trading strategy. We apply strict quantitative rules to choose the pairs as well as quantitative rules to trade the pairs. Generally practitioners use a combination of quantitative and qualitative rules such as accounting data, sentiments, momentum, macro-indicators, analysts' or their own views on the shares etc. Considering the scope of the dataset, over 250 shares and 10 years, (1996-2006), we have chosen *not* to include any other metrics than historical price and volume data. An alternative would have been to look at a smaller sample of stocks, and add qualitative data and/or more complex quantitative data such as valuation and accounting data. We believe though that the quality of our data sample based on price/volume data is higher than the valuation metrics, such as P/E or EV/EBITDA estimates which would have been our second choice of data to study.

Generally when discussing the trading of pairs, one refers to shares which do not only have high correlation historically, but that also have similar fundamental characteristics, i.e. belong to the same industry. Considering that this limitation of pair formation is so widespread, we have chosen to perform two tests, one when the pairs formation is completely free; i.e. any two shares can form a pair; and one when we limit the pairs formation to within the same industry.

In our simulation we measure return as the total return in Swedish kronor from our total position (recall that in our trading methodology we simulate one short and one long position of Skr1.). If our traded pair yields a profit of Skr0.2, we attribute a return of 20% to that trade. We have chosen this approach for model purposes as well as our best understanding of how practitioners view their exposure. An alternative would have been to look at the total gross exposure, in this case Skr2, making the total return in the above example 10%. Another possible approach would have been to look at the level of capital an investor would have to commit as margin for the short position. That approach would imply significantly higher returns, since the capital would be a fraction of the total short position. This later approach would have been too complicated to simulate within the scope of this research paper.

Further, we treat all returns above zero as excess returns. This could be viewed as aggressive considering that our trading strategy naturally will incur some degree of volatility. We have chosen this approach since in our simulated trading strategy is self funding, our short position fund our long position and hence we do not commit any capital. Using a different level than zero would be somewhat arbitrary and make the results less straight forward to interpret, in our view. An alternative would have been to benchmark the returns against the risk free rate. One problem when comparing with the risk free rate is that since our portfolios are self funding our simulation methodology does not have any obvious capital base from which to calculate the risk free benchmark return. However, we could have modelled the maximum negative position during each trading period and used that as our capital base (mirroring what in practise would have been the likely maximum capital a hedge fund manager would have to commit to the bank from which s/he had lent the stock s/he shorted). Another suggested capital base to consider would have been the total gross exposure (i.e. the total long position plus the total short position).

2.3.1 Step one - Pairs formation

As mentioned above, we do two sets of pairs formation, one which uses our full set of available shares and one which limits the number of pairs to those formed within an industry. The pairs are being formed under a 12 month period and then traded during a 6 month period. This process is being repeated every month. Our first 12 month formation period is from 1 Jan 1995 to 31 Dec 1995, and our final is from 1 Dec 2004 to 30 Nov 2005. We limit the share sample to those shares that were traded every trading day during the specific formation period (i.e. the traded volume had to be greater than zero).

After the formation period we rank the pairs based on correlation of total returns, i.e. price appreciation plus dividend. We have chosen to use simple correlation to rank our pairs, due to three main reasons.

- It is the most common "rule of thumb" metric, at least to our best knowledge. Even though some quantitative trading strategies today consist of a number of complex statistical screening processes, correlation captures the basic principle regarding pairs trading, i.e. to find stocks that "move together", and this was our aim with the study.
- The most similar study to the one we are conducting (Gatev, Goetzmann, and Rouwenhorst 1998), used correlation to rank their pairs when studying US shares. We believe that using the same criteria increased the comparability of the studies.
- Taking into consideration the number of possible pairs, (35511 in total) our restrictions with regard to computer power and programming time limited the choice of statistical criteria for each pair.

An alternative method of ranking pairs would have been time series stationarity which measures the strength if the relation between two stocks in a pair is stable over time. Alternatively we could have used sensitivity or correlation with a number of macro-variables such as interest rates, FX rates, GDP growth, industrial production etc., and ranked our pairs with regard to similarities in correlation with such metrics. We believe that we to some extent capture this element by limiting the pairs formation to pairs within one industry in one of our tests.

We rank our pairs from highest to lowest with regard to correlations during the 12 month formation period. We start a new 12 month period every month, hence making the total number of formation periods 120 (one for each month during the 10-year testing period).

When testing with the restriction that pairs need to be within the same industry group, we used the same approach as above. As industry groups we used the classifications by SIC into 6 large industry groups. Considering that pairs could only form within the 6 industry groups, our pairs formation period included substantially fewer possible pairs.

2.3.2 Step two – Pairs selection

In our empirical sample we select a total of two different portfolios of pairs: the 5 pairs with the highest correlation and the 20 pairs with the highest correlation. In addition we constructed one random portfolio consisting of 20 randomly selected pairs. We select each portfolio after every formation period. We will thus have three sets of 120 portfolios, each consisting of the top 5 pairs, the top 20 pairs and 20 randomly selected pairs. One stock can appear in several portfolios, i.e. SEB vs. Handelsbanken could form a pair as well as SEB vs. Swedbank in the same portfolio. Given that there is a large degree of overlap during the formation periods, 11 months of period two will have been included in the formation period for period one as well, one pair often appears in several consecutive portfolios.

As discussed above, we acknowledge that practitioners rarely would execute trades based solely on a ranking system similar to the one we apply in our simulation, but rather use it as a screening method. One alternative to our approach when ranking the highest pairs, would have been to apply a specific limit to our pairs selection process. We could for instance have formed a portfolio of pairs above 0.9 in correlation and another with correlation above 0.8 etc. That methodology would probably have mirrored the approach employed by practitioners better. We have chosen to use a ranking methodology for three reasons.

- Firstly, the ranking methodology was the one used by GGR (1998), making comparisons between the two studies easier.
- Secondly, were we to have used a fixed limit, our trading portfolios would have more pairs during periods when the stocks show a high degree of correlation and vice versa. This would imply that we would have different numbers of pairs actively

trading at different time periods, making the results somewhat more complicated to interpret.

2.3.3 Step three – Trading simulation

In our research design we simulate trading each portfolio of pairs for 6 months after the formation period. Thus, we will have 6 portfolios running parallel at all times for each of the 5 portfolio types (2 using all shares, top 5 and top 20, 2 using only industry pairs, top 5 and top 20, and 1 random portfolio). During the first 5 months and the last 5 months of the testing period we will have less than 6 portfolios active, as we build up our trading portfolios and shrink them respectively.

As a signal to activate a trade we apply a simple two standard deviation rule. When the total returns of the two stocks diverge with more than two standard deviations, as measured during the formation period, we open a position going long the stock that has underperformed and short the stock that has outperformed. For modelling and return calculation purposes we calculate our long and short position as 1 Skr. We keep the position open until either the share returns converge, i.e. the performance 'gap' of two standard deviations closes or the six months trading period closes. One pair can open several times during one trading period. This trading rule is the one used by GGR (1998). There are a number of alternative trading rules than could have been used, most noticeably 1 or 3 standard deviations (or 1.5 or 2.5); using 3 would leave us with very few trades, and similarly using 1 standard deviation would expose us to significant trading costs. While the use of exactly 2 standard deviations is somewhat arbitrary, we believe that using a round number is to be preferred given the substantial risk of data-snooping (further developed below in 2.4), were we to try several different trigger levels.

One serious weakness of our trading rule is that for each trade, the loss is unlimited, while the profits are capped at two standard deviations. Firstly this will cause our returns to be asymmetric, with many small profits and a few big losses, making the statistical analysis of the results more complex; secondly this is contrary to our understanding of how practitioners trade. Given the seriousness of this weakness we have chosen to run the dataset using a stop loss rule as well as trading without a stop loss. In the design of our study, we have chosen to use a stop-loss of 10% and 20% respectively.

We use the closing prices of the following day when we open and close our positions to calculate returns. This will cause our returns to differ slightly from two standard deviations. An alterative would be to use the closing prices of the same day as the 2 standard deviation rule is triggered, however that would in essence assume that the trade was put on before the trigger actually took place (the two standard deviation rule is based on closing prices). Another alternative would be to use VWAP (Volume Weighted Average Price), while this would probably resemble more closely the actual prices that practitioners would be able to expect; we compared the average difference between VWAP and closing price and found the difference to be insignificant. GGR's (1998) study was also conducted using closing prices.

2.4 Data-snooping

Data-snooping, also called data mining, refers to the risk of coming up with spurious results as a consequence of trying enough different strategies. (cf. Sullivan, Timmermann, and White 1999) The risk of data-snooping when testing a pure quantitative trading strategy is considerable. With the aid of computer power there are few restrictions as to the number of trading rules, such as when to open and close a position, and statistical methodologies, such as co-integration, correlation, time-stationary etc., that could be tested for. We have tried to minimize the risk of data-snooping by using as straightforward rules as possible. In addition we have tried to mirror what we believe is a proxy of how practitioners trade. This could be viewed as data-snooping since practitioners might have tried a number of different trading rules in order to come up with one that yields positive returns. However, since the aim of this paper is to test whether the observed trading behaviour could generate excess returns using a pure quantitative rule; data-snooping was not an issue.

3 Theory and Previous Research

In this section we discuss the theoretical background to our study. We begin by giving a brief overview of the market efficiency theory and the early tests of return predictability that first challenged the view of market efficiency. We then look into the more recent tests that have been done with regard to trading systems – tests that are more closely related to the study which we have conducted. We end the theory section by looking at the test conducted by GGR (1998) which is similar to the test that we have performed.

3.1 Theoretical overview

Since it was first introduced, the theory of the efficient market, in either its weak or semi-strong form, has been one of the dominating research fields within finance. The questions asked include: are the markets efficient, to what degree are they efficient and how we should measure whether they are efficient or not? Although this paper is an empirical study and simulation of a trading rule, our study relates to the topic of efficient capital markets.

Tests of market efficiency can be divided into the following three groups. Firstly, tests of the weak form (how well do past returns explain future returns). Secondly, semi-strong-form tests (how quickly do security prices reflect public information), and finally strong form tests (do any investors have private information that is not fully reflected in market prices? (see e.g. Fama 1991) As suggested by Fama (ibid.), tests of weak form of the Efficient Market Hypothesis could be widened to include not only the forecasting power of past returns, but all tests of return predictability. Return predictability tests can be divided into two groups: tests of fundamental indicators such as dividend yield, P/E and P/B as predictor of future returns, and tests based solely on previous stock returns.

Many tests of market efficiency suffer from the joint hypothesis problem, i.e. in order to test for market efficiency (whether stocks are efficiently/correctly priced), we need to identify what we regard as correctly priced. In this case we need a model for what the correct price of a stock is, either relatively or in absolute terms. This dilemma became clearly exposed by Fama and French (1992) as they contested the single factor CAPM, which for many years had dominated the way academics as well as practitioners thought about asset pricing. CAPM states that investors should only be compensated for the systematic risk of the asset; systematic risk is normally measured as Beta. Many tests had also been conducted with CAPM as the starting point. Campbell and Schiller (1989), for instance, find that P/E has reliable forecasting power, something that at the time could be viewed as criticism of market efficiency. Fama and French (1992), however, argue that using a multifactor model, including a size and a value parameter, rather than the single factor CAPM model, can explain the P/E-“anomaly”.

In our research and simulation, as discussed in the methods section above, we only use historic and publicly available stock price information. The only part of the study that is not based on historical data is the trading rules. Our research thus fits into the group of empirical studies that have been performed relating to the weak form of market efficiency.

3.2 Tests of return predictability

In this part we discuss previous tests of return predictability. We group them into 3 categories with an increasing order of relevance with regard to the study we have conducted. The first 2 are tests of fundamentals and tests of calendar effects. The third, and for our research most relevant group of empirical tests refers to tests based on historical price information such as momentum, mean reverting of share price returns, auto-correlation of share price returns etc. In each of the three categories we aim to discuss the empirical evidence that an “anomaly” exists and possible explanations of the same.

3.2.1 First category: return predictability of fundamentals

Basu (1977) demonstrates that buying a portfolio of low P/E shares yields a higher return than just holding the market portfolio. Fama and French (1992) find predictive power in size and price to book multiples. The evidence is so striking that it challenges the dominating single factor CAPM, suggesting a multifactor one including size and value. Campell and Shiller (1989) further show that both dividend yield and P/E multiples have explanatory power in predicting returns, albeit over longer terms, between four and five years. Mueller (2001) finds predictive power in a profit margin proxy when comparing prices of raw materials to the Producer Price Index.

Fama and French (1992; 1996) explain the “anomalies” of predictive powers of price to book multiples and size not by the fact that the market is inefficient, but that the single factor asset pricing models are incorrectly specified. Malkiel (2004) points out that a large part of the return predictability of P/E ratios could be explained by changes in long term interest rates and hence could just be viewed as a consequence of changing discount rates.

3.2.2 Second category: the calendar effects

One empirical finding within the second category relates to the January effect. Rozeff and Kinney (1976) found a statistically significant excess return in January compared to the average of the other 11 months on the New York stock exchange. Gültekin and Gultekin (1983) found in a cross market study that a majority of the markets showed evidence of seasonality. A more recent study by Schwert (2003) using data up to 2001 found that the January effect had become less pronounced during the period 1980-2001, but still existed. Another well-documented calendar effect is the Weekend effect, also called the Monday effect. This effect is the phenomenon that stock price performance seems to be weaker on Mondays. French (1980) finds in a study of share performance between the years of 1953-1977 that share price performance tends to be negative on Mondays whereas returns are positive for the other four weekdays. In a cross market study by Agrawal and Tandom (1994) they find significantly negative returns in a large majority of countries on Mondays. Rogalski (1984) suggests that the Monday effect is actually a weekend effect, i.e. that the weak performance is due to the lower opening price on Mondays. Lakonishok and Smidt (1988) show that US stock returns are significantly higher at the turn of the month and Ariel (1984) find a similar effect on the last trading day of every month.

Several theoretical reasons behind the January and Monday effects have been suggested. Haugen and Jorion (1996), for instance, gives a possible explanation of the January effect. They suggest that successful managers who are risk averse reduce their positions ahead of year end to lock in their bonuses. Similarly, unsuccessful managers who wanted to avoid a very bad year which could potentially cost them their job, also reduce their risk positions before year end. In January they invest again, contributing to the January effect. Keim and Stambaugh (1984) suggest that the Monday effect could be explained by measurement errors. Draper and Paudyal (2002) find in a study on UK stocks that, adjusting for various possible influences, Monday returns do not differ significantly from other days.

3.2.3 Third category: past returns as a predictor of future returns

The majority of the tests using past returns are centered on the concept of buying winners (relative strength strategies) or buying loosers (contrarian or mean reverting strategies) according to some pre-specified rule. These empirical findings are harder to explain without including some element of inefficiency in the market, at least during certain time periods.

Fama (1965) finds that for most stocks included in the Dow Jones Industrial index, daily share price returns are auto-correlated, which suggests that just buying whatever stock performed the best the day before would result in outperformance. A later study by Lo and MacKinlay (1988) suggests that portfolio returns of stocks grouped into portfolios according to size are also auto-correlated. Jegadeesh and Titman (1993) find that buying winners and selling loosers generates significant returns on a three to twelve months time horizon.

Another group of closely related tests focus on mean-reverting and relative strength. Mean reverting could be described as the opposite of auto-correlation. The idea is that a stock has a fixed, or slow moving fundamental value, and that the share price can drift away from that fundamental value during shorter time periods. In the long run, however, the share price should revert back to its fundamental value; the share price should mean revert. Shiller, Fischer and Friedman (1984) as well as Summers (1986) argue that the daily returns are so small, and the noise so considerable, that tests looking at daily or weekly returns fail to acknowledge the mean-reverting characteristics, but find that over longer time periods mean-reverting is significant. DeBondt and Thaler (1985) achieve abnormal returns on a 3-5 year time horizon by buying the stock that performed worst over the previous 3-5 years and selling the ones that performed well. Jegadeesh (1990) as well as Lehmann (1987) establish that contrarian strategies using the stock returns of the previous week and month generate significant positive returns.

This mean-reverting anomaly is closely related to the study we are conducting. In effect that is exactly what we are doing although we do it on a relative basis rather than an absolute. We sell a stock that has outperformed relative to another closely correlated stock and buy the one that has underperformed relative to its long term trend.

The above findings constitute the most viable criticism against the weak form of the market efficiency theory. If one can use historical returns to develop a trading pattern that generates excess returns, markets are not efficient even with regard to the weak form of the market efficiency hypothesis. Fama (1991) suggests that although many studies have found evidence of, for example, auto-correlation in stock returns, the magnitude of the returns has generally been too small to take advantage of in practical trading, taking for instance trading costs and bid-ask-bounce into consideration. Lakonishok, Shleifer and Vishny (1994) argue that the value effect could partly be explained by investors' tendency to overreact to new information. They argue that investors fail to take into account the tendency of company cash flows (as a measure of fundamental value) to mean-revert. Hence, investors overestimate both positive and negative news about a company.

Another group of explanations refer to purely psychological factors. Daniel, Hirshleifer, and Subrahmanyam (1998), for example, suggest that investors show overconfidence and biased self-attribution, which could explain both short term autocorrelation and longer term mean reverting.

3.3 Theoretical contribution of our study

In this final part of our theoretical section we specifically underline a subgroup within the category of using historical returns to predict future returns, namely market neutral investing. Below we discuss one study of particular relevance to our research; the GGR-article from 1998. In the end of the section we highlight what we believe to be the theoretical research gap within studies of pairs trading.

3.3.1 Market neutral investing

Tests of market neutral trading strategies relate closely to the traditional tests of mean reverting and auto-correlation with two main differences. Firstly the tests generally focus on achieving "double alfa", i.e. buying a security that according to some statistical metric is cheap and selling another security that according to the same statistical metric is expensive. Secondly, market neutral strategies are, as the name suggests, interested in the trading strategies correlation with the market in general or a specific benchmark from which the securities are selected. The aim is to achieve as low correlation as possible with the benchmark index, in the case of US shares, mostly the S&P 500.

Alexander and Dimitriu (2002) achieve statistically and significant outperformance using co-integration, a measure of mean-reverting, to select long and short portfolios. A more qualitative study was conducted by Barra RogersCasey (2000) who analyse the performance of market neutral long short equity funds in the US. The Barra institute found that from 1991 to 2000, the average fund outperforms the treasury bills; adjusted for risk, however, the benefits look less promising with a Sharpe ratio below 1. A Sharpe ratio, however, might not be the correct way to evaluate a market neutral strategy. Barra RogersCasey (ibid.) argues that returns of market neutral strategies have no correlation with market returns, and hence should not be evaluated using a Sharpe measure as such.

3.3.2 Theoretical research gap

The study of market efficiency is one of the most researched fields within financial theory. And a number of studies (see section 3.2 above) suggest that it is possible to predict future returns using historical returns as a way to achieve returns in excess of what can be explained by systematic risk. These results points in the direction that markets at least historically have been inefficient even with regard to the weak form of the efficient market hypothesis. It is, however, probably safe to state that the markets are becoming increasingly efficient, why continuous research within this field is, in our view, warranted. We mean that several of the main drivers behind increased efficiency have accelerated over the last ten years including; lower transaction costs, increasing transparency, increased number of market participants, faster dissemination of information and significantly increased use of computer power when analysing stock returns.

GGR (1998) achieve up to 12% annual excess returns trading "top" pairs on the US stock exchange. As mentioned in our methodology section, they use historic correlation to rank pairs, and then trade them using several different trading strategies. GGR conclude that even taking transaction costs into account, these strategies achieve excess returns. Hence, the study indicates that it was possible to achieve excess returns using a PT strategy based on historical correlations from 1962 through to 1997 in the US. Whether this still holds today and whether it holds on other stock exchanges is still unclear. There has, as far as we know, been no study on Swedish data using this methodology.

GGR (ibid.) tested a very simplistic trading rule. There are a number of more advanced, and probably more realistic trading strategies that could be tested. The most significant, in our view, would be the use of stop-losses, something which is widely used within the hedge fund industry. There are a number of alterations to that very basic trading rule which have not been tested.

3.4 Summary of theoretical section

Above we introduced the main issues with regard to the study of market efficiency and how our study fits into that broad research field. We discussed a numbers of tests conducted in that field and divided them into three categories; tests of fundamentals, tests of calendar day effects and tests of historical returns as a predictor of future returns. We specifically emphasized a group of tests which are closely related to the study we aim to conduct in this particular paper.

Although there are possible explanations for most of the studies made that challenge the weak form of the market efficiency hypothesis, as a group they highlight the fact that anomalies are likely to exist in the stock markets and that rightly formulated and executed, it should be possible in theory to achieve excess returns. The most damaging criticism to these tests is, in our view, whether they are possible to execute in practice. Broadly there are two main issues; the difficulty of assessing the costs related to trading, both direct and indirect, and competition - as soon as the anomaly becomes known, it gets competed away.

We ended the section with a brief discussion on our view of the research gap on PT.

4 Empirical Analysis and Results

In this section of the paper we present the main empirical results from our trading strategy developed in section 2. We divide the results of our trading simulation into 3 categories. Firstly we present the results achieved by using the basic trading rule, where we use free pairs formation and no stop-loss. Then we discuss the impact on the empirical results by using a stop loss, and finally we look into the impact of using industry formation, i.e. where we only form pairs where both stocks belong to the same industry. We present the results both for the top 5 pairs and the top 20 pairs.

To increase the readers' understanding of the steps in our test, we will start with a summary of the results from our simulation on PT. We did not achieve a positive return either by trading the top 5 or the top 20 pairs. The main driver behind the negative performance was significant losses during 1999, predominantly within the IT sector. When we included a stop-loss the returns became positive, but only just. Trading industry pairs yielded a similar result, and it was only within the Banks sector where our trading strategy consistently delivered significant excess returns. The results trading random pairs yielded negative performance, see appendix 1. The returns were similar to those when trading top pairs, indicating that we, within our sample data, had a tendency for returns to be auto-correlated (see for instance Jegadeesh and Titman (1993) but not have any significant tendency for general mean reverting. If returns are auto-correlated, buying losers and selling winners should yield a significant negative performance.

We have simulated the trading both with and without a suggested trading cost of 15 basis points per trade. We have chosen to mainly present the "raw" results, i.e. excluding any trading costs, given that we do not achieve positive returns and thus we do not achieve any excess returns even excluding the trading costs. On average our pairs open just below once per trading period, why including trading costs would lower the returns with on average of just below 4 x 15 basis points (0.6%), since each traded pair include buying and selling each of the two shares. See appendix 1 for returns including trading costs.

4.1 Free pairs formation – no stop loss

Table 1 below shows the total return of the top 20 and top 5 pairs achieved using free pairs formation and no stop loss. The return is calculated as the average return per portfolio per 6 months period (the period during which the portfolio was traded). Here we underline 2 implications; firstly, that the return on traded pairs show the average return on all the pairs that opened at least one time during the 6 months period. Secondly, that the return on total pairs show the average return calculated including also the pairs that did not open. The average return looking at all pairs is close to 4% negative and the standard deviation is 26%. Worth noting is that the largest loss is substantially higher than the largest profit. This is a consequence of our asymmetric trading rule, where we close out of the position if the trade works (the statistical outperformer underperforms and vice versa) but leave the position open if it doesn't.

Table 4.1

No stop loss - all pairs	**Return on total pairs**		**Return on traded pairs**	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**-1.54%**	**-3.93%**	**-0.99%**	**-3.43%**
Median	1.9%	0.9%	2.2%	1.2%
Stdev	15.3%	25.9%	18.1%	29.7%
Max Loss	-98.1%	-203.8%	-122.6%	-226.4%
Max Profit	27.7%	26.3%	46.2%	43.8%
Periods >0	61.7%	57.5%	60.8%	57.5%
Z	-1.10	-1.66	-0.60	-1.26
Significant	FALSE	FALSE	FALSE	FALSE

We examined the total numbers of pairs for 2 out of 3 factors Fama and French (1992) uses in their multifactor model; size and Beta. Unfortunately, we neither had sufficient resources nor data to look into P/B.

Our total traded pairs do not show any substantial tendency to go long or short high beta stocks, the difference in Beta between our long and short positions is less than 5% on average. We do, however, see a small tendency to short large cap stocks. The median ratio between our short position and our long position with regard to market capitalisations is 1.2. Fama and French (ibid.), however, show that there is a systematic out-performance of *small* cap stocks, why this tendency can not explain the negative performance that we achieve.

4.2 Free pairs formation – using a stop loss

To take into account the asymmetric returns as well as our best impression of common practise we have traded using the same rules as above but added a stop loss, i.e we close out of any given position automatically if we incur a loss of 10% and 20% respectively. As can be seen in table 2 below, applying a stop loss changes the return characteristics materially; our trading rules changes from being loss-making to yielding a positive return, and in the case of the 10% stop loss a significant positive return, however, adjusting for trading costs makes the returns insignificant. Unsurprisingly, the standard deviation is much lower using a stop loss. A closer analysis of the results reveals that the lion's share of the improvement in returns is a due to us avoiding large losses that occurred within the IT sector in the late 1990's. Please see appendix 1 for a more detailed analysis of these trades.

Table 4.2

Return on total pairs	**20% stop loss**		**10% stop loss**	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**0.59%**	**0.71%**	**1.27%**	**1.00%**
Median	1.3%	0.8%	1.2%	0.4%
Stdev	8.0%	5.0%	6.6%	3.9%
Max Loss	-17.1%	-11.4%	-10.0%	-7.7%
Max Profit	27.7%	21.3%	25.4%	13.8%
Periods >0	57.5%	57.5%	55.8%	55.8%
Z	0.81	1.56	2.12	2.83
Significant	FALSE	FALSE	TRUE	TRUE

4.3 Pairs formation within an industry only

Above we highlighted the results using free formation and a stop loss. Below we will present the results when we restricted the pair formation to include only pairs of shares within the same industry. We traded the industry pairs both with and without a stop loss. We use the Swedish SIC definition and separated the stocks into 6 main industry groups.

Trading industry pairs without a stop loss, just like the case with all shares, yields a negative return. Adding a stop loss again changes the returns from negative to (significantly) positive.

Table 4.3.1

Trading Industry pairs	Return on total pairs		20% stop loss		10% stop loss		10% stop loss incl. trading costs	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**-7.33%**	**-2.44%**	**1.27%**	**1.57%**	**1.54%**	**1.44%**	**0.76%**	**0.87%**
Median	1.9%	0.7%	1.4%	1.0%	0.8%	0.9%	-0.5%	0.6%
Stdev	52.4%	22.8%	7.6%	4.3%	6.4%	3.7%	6.4%	3.5%
Max Loss	-456.1%	-160.5%	-16.0%	-8.5%	-10.0%	-7.0%	-10.0%	-7.0%
Max Profit	30.7%	20.0%	30.7%	16.5%	25.4%	12.4%	24.9%	12.1%
Periods >0	57.5%	57.5%	59.2%	59.2%	64.2%	64.2%	57.5%	57.5%
Z	-1.53	-1.17	1.84	3.98	2.63	4.30	1.31	2.71
Significant	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE

Below we have separated the top 20 pairs that were formed into separate industries. Again it is the large losses within the IT sector, predominantly in 1999, which contribute to the negative performance. Please see appendix 1 for more detailed discussion on this topic.

Table 4.3.2

	Sector	N	Return	STD	Significant
No stop loss	Total	2400	-2.4%	54.8%	True
	Basic materials	270	2.5%	14.4%	True
	Communications	350	-3.0%	25.6%	True
	Consumer Cyclical	33	4.4%	19.8%	False
	Financial	718	2.5%	11.0%	True
	Industrial	693	1.7%	12.2%	True
	Technology	335	-25.7%	139.5%	True
20% stop loss	Total	2400	1.6%	16.4%	True
	Basic materials	270	2.7%	13.9%	True
	Communications	350	-1.4%	18.5%	False
	Consumer Cyclical	33	2.7%	19.6%	False
	Financial	718	2.4%	10.6%	True
	Industrial	693	1.3%	12.4%	True
	Technology	335	2.3%	28.6%	False

4.4 Special cases – Enea and the Swedish banks

We have made an attempt to go through the entire data sample with a more qualitative approach, trying to establish what the most important drivers of the total performance are. We want to draw you attention to two, in our view, interesting observations in the sample: firstly the large losses incurred during the second half of 1999 in the software company Enea and secondly the consistent performance when trading Swedish banks.

4.4.1 Enea – key driver of losses

Both amongst the industry pairs and amongst the freely formed pairs, there are a very high degree of very large losses involving the software company Enea. Each of the 20 largest losses for both formation strategies includes a position in Enea, all formed during the 7 months period from July 1999 to January 2000 (and traded up until July 2000). Please see appendix 1 for a complete list of Enea pairs. During this time period Enea appreciated with over 800% and the trading strategy repeatedly suggested going short Enea. If we exclude all trades involving Enea, the return changes from negative to positive for both strategies.

Table 4.4.1

Total return	Including Enea		Excluding Enea		Enea pairs	
Pair 1 - 20	All	Industry	All	Industry	All	Industry
Average	**-3.93%**	**-2.44%**	**0.61%**	**2.72%**	**-164.42%**	**-145.31%**
Stdev	54.7%	54.8%	23.0%	19.8%	251.5%	257.1%
Max Loss	-850.4%	-869.2%	-149.0%	-166.3%	-850.4%	-869.2%
Max Profit	210.7%	131.0%	210.7%	131.0%	46.5%	53.6%
Nr of pairs	2,400	2,400	2,334	2,327	66	73

The issue with Enea clearly highlights one of the problems using correlation as opposed to other measures for selecting pairs, in that correlation only measures the degree to which the stocks move together and hence do not take into account the size of this movement. In our sample the stock Enea made a very large upwards movement in 1999-2000, resulting in that it showed a strong correlation with many other stocks in our sample that also appreciated during this time period. One solution to this problem would be to introduce beta as a restriction in the formation, and for instance block pairs that do not have similar Beta's.

4.4.2 Swedish banks – consistently delivering out-performance

We achieved the best and most consistent performance within the Financials sector, see table 4.3.2 above. A closer look at the Financials sector reveals that it is predominantly the four large Swedish banks (Nordea, SEB, Handelsbanken (SHB) and Swedbank (FSPA)) that deliver the return.

Table 4.4.2

Swedish bank pairs		Average	N pairs	STD	Max loss	Max profit	Sig.1% level
Nordea	SEB	5.59%	58	9.5%	-6.4%	35.4%	yes
Nordea	SHB	2.90%	58	7.7%	-11.2%	22.2%	yes
Nordea	FSPA	2.41%	57	7.1%	-9.5%	18.8%	yes
SEB	SHB	5.18%	104	11.2%	-19.2%	51.2%	yes
SEB	FSPA	4.92%	87	8.2%	-12.5%	35.3%	yes
SHB	FSPA	3.60%	92	8.7%	-11.0%	38.8%	yes
Total		**4.23%**	**456**		**-19.2%**	**51.2%**	

Looking at the 6 possible pair formations within the four Swedish banks, all combinations occurred frequently and all delivered significant excess returns with low standard deviations. Although we should restrain from drawing too far fetched conclusions on back of these results due to the risk of data snooping, we note that close to 20% of all industry pairs consists of a combination of two Swedish banks.

4.5 Summary and comparison to GGR

Contrary to GGR (1998) we did not achieve any excess return using the above discussed straightforward trading rule. Whereas GGR (ibid.) achieved an average excess return of close to 12% annually we achieved a negative annual return of almost 4%. The lion's share of the negative return was generated within the IT sector in the late 1990s. But even looking at the returns excluding the IT bubble we did not yield results similar to those of GGR. One explanation for this could be that the shares of the Stockholm Stock Exchange are less homogeneous than those of New York Stock Exchange, leading to our study comprising of more heterogeneous pairs. Another possible explanation might be that our time period consists of three rather distinct periods of strong trend performance: up 5 years (1995-1999), down 3 years (mid 2000-to mid 2003) and then followed by up 3 years. This left us more exposed to the issue of selling high beta stocks in a rising market and buying low beta stocks in a falling market.

5 Summary and suggestions for further research

In this research paper we have applied a straight forward quantitative trading rule known as pairs trading (PT in this paper). In the trading simulation we sold one (outperforming) security and bought one (underperforming) security according to historic price correlations. We treat all positive returns as excess returns considering that our simulated strategy does not involve committing any capital and only have negliable systematic risk. The trading strategy turns out to be unprofitable looking at the top 20 pairs or top 5 pairs over the whole period, without using a stop loss. The underperformance is predominantly driven by substantial losses during the autumn of 1999, when our trading strategy suggested going short Enea as the stock outperformed other IT companies. When we exclude either Enea specifically or the year 1999 we achieve positive and, in some cases, significant returns.

The picture changes somewhat when we apply a stop loss of 10% and 20% respectively. The return is then positive and in one case (10% stop loss trading industry pairs) significant even after adjusting for trading costs, largely due to the avoidance of the larger than 100% losses that occurred in 1999, highlighted above. However, the returns are small, less than 1% per portfolio after adjusting for trading costs, and unlikely to be attractive enough for a fund manager or hedge funds, at least to our best knowledge. In addition an issue with using a stop loss, particularly as low as 10%, is that around 45% of all traded pairs close at the stop loss, likely incurring significant trading costs, since we make a forced exit.

Looking and specific industries, we found that trading pairs within the financial sector, predominantly the large Swedish banks, consistently achieves a significant excess return large enough to attract capital even after adjusting for trading costs. A possible explanation for this could be that the four large Swedish banks are much more similar in characteristics, than for instance the two electrical engineering companies Electrolux and Autoliv, which both belong to the consumer cyclical category. Another characteristic common to all the Swedish banks is that they are large market cap, and did not participate in the boom/bust of the dot.com-bubble in the late 1990s. Hence, our study indicates that some sort of qualitative overlay is necessary to achieve excess returns trading pairs at the Stockholm Stock Exchange. Simply buying the statistical underperformers and selling the statistical outperformers is likely not enough. Our study also suggests that the use of a stop loss increases the return significantly. We believe that this highlights a shortfall in the applied trading strategy, namely that it does not take into account beta, leaving us exposed to the situations where high beta names will be shorted in a rising market.

Another observation from our research is that the choice of pairs seems to be more important than the choice of trading rule. As an example of this we underline that pairs formed from the four Swedish banks yielded an average of 6% per 6-months period, and achieved positive excess returns all ten years of simulated trading, 8 of which was significant on the 1%-level.

One area which we set out in our purpose to study was the development of excess returns over time. Considering that we overall, when not using a stop loss, achieved negative returns it is hard to draw any clear conclusions from the results. However, we do note that our first two years of study, 1996 and 1997 yielded a positive return whilst the last two, 2004 and 2005, yielded a negative return.

We believe that the most interesting area for further research on PT relates to the choice of pairs, such as excluding small cap, companies with unclear business model, as well as only trading pairs belonging to certain industries. Another suggestion for further research would be to study application of more advanced statistical selection criteria. Although we used correlation in our empirical analysis, another method would be to look at time stationary and strength in mean reversion.

6 References

Agrawal, A. and K. Tandom. 1994. "Anomalies or the illusion? Evidence from stock markets in 18 countries." *Journal of International Money and Finance*:83-106.

Alexander, C. and A. Dimitriu. 2002. "The Cointegration Alpha: Enhanced Index Tracking and Long-Short Equity Market Neutral Strategies." *ISMA Discussions Papers in Finance, University of Reading*.

Ariel, R. A. 1984. *A Monthly Effect in Stock Returns*: Massachusetts Institute of Technology.

Barra RogersCasey, Research 2000. "Market Neutral Investing." www.hedgeworld.com.

Basu, S. 1977. "Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis." *The Journal of Finance* 32:663-682.

Brown, S. J., W. Goetzmann, R. G. Ibbotson, and S. A. Ross. 1992. "Survivorship bias in performance studies." *Review of Financial Studies* 4:553-580.

Campbell, J. Y. and R. J. Shiller. 1989. "Stock Prices, Earnings and Expected Dividends." NBER Working paper.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam. 1998. "Investor Psychology and Security Market under-and Overreactions." *The Journal of Finance* 53:1839-1885.

DeBondt, W. F. M. and R. Thaler. 1985. "Does the stock market overreact." *Journal of Finance* 40:793-805.

Dey, M. K. and S. Flaherty. 2005. "Stock Exchange Liquidity, Bank Credit, and Economic Growth."

Draper, P. and K. Paudyal. 2002. "Explaining Monday Returns." *Journal of Financial Research* 25:507-520.

Fama, E. F. 1965. "The Behavior of Stock-Market Prices." *Journal of Business* 38:34.

—. 1991. "Efficient Capital Markets: II." *The Journal of Finance* 46:1575-1617.

Fama, E. F. and K. R. French. 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance* 47:427-465.

—. 1996. "Multifactor Explanations of Asset Pricing Anomalies." *The Journal of Finance* 51:55-84.

French, K. 1980. "Stock Returns and the Weekend Effect." *Journal of Financial Economics* 8:55-69.

Gatev, E., W. N. Goetzmann, and K. G. Rouwenhorst. 1998. "Pairs Trading: Performance of a Relative-Value Arbitrage Rule." in *Working Paper Series*

Gültekin, M. N. and N. B. Gultekin. 1983. *Stock Market Seasonality: International Evidence*: Wharton School, University of Pennsylvania.

Haugen, R. A. and P. Jorion. 1996. "The January Effect: Still There after All These Years." *Financial Analysts Journal* 52:27–31.

Jegadeesh, N. 1990. "Evidence of Predictable Behavior of Security Returns." *The Journal of Finance* 45:881-898.

Jegadeesh, N. and S. Titman. 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance* 48:65-91.

Keim, D. B. and R. F. Stambaugh. 1984. "A Further Investigation of the Weekend Effect in Stock Returns." *The Journal of Finance* 39:819-835.

Lakonishok, J., A. Shleifer, and R. W. Vishny. 1994. "Contrarian Investment." *Extrapolation, and Risk, Journal of Finance* 49:1541-1578.

Lakonishok, J. and S. Smidt. 1988. "Are seasonal anomalies real? A ninety-year perspective." *Review of Financial Studies* 1:403-425.

Lehmann, B. 1987. "Fads, Martingales, and Market Efficiency." *The Quarterly Journal of Economics, February*:1-28.

Lo, A. W. and A. C. MacKinlay. 1988. "Stock market prices do not follow random walks: evidence from a simple specification test." *Review of Financial Studies*.

Malkiel, B. G. 2004. "Models of stock market predictability." *Journal of Financial Research* 27:449-459.

Mueller, J. 2001. "Driving the Market, Looking for a Rational Theory of P/E Ratios." *Barron's, June* 25.

Rogalski, R. J. 1984. "New Findings Regarding Day-of-the-Week Returns over Trading and Non-Trading Periods: A Note." *The Journal of Finance* 39:1603-1614.

Rozeff, M. S. and W. Kinney. 1976. "Capital Market Seasonality: The Case of Stock Returns." *Journal of Financial Economics* 3:379–402.

Schwert, G. 2003. "Anomalies and market efficiency, chapter 17 of the Handbook of Economics and Finance (G. Constantinides, M. Harris, and R. Stulz, eds.)." North Holland, Amsterdam, Netherlands.

Shiller, R. J., S. Fischer, and B. M. Friedman. 1984. "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity* 1984:457-510.

Sullivan, R., A. Timmermann, and H. White. 1999. "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *The Journal of Finance* 54:1647-1691.

Summers, L. H. 1986. "Does the Stock Market Rationally Reflect Fundamental Values?" *The Journal of Finance* 41:591-601.

7 Appendix 1

7.1 Trading randomly selected pairs

We simulated traded of two randomly chosen portfolios, one with 5 pairs and one with 20 pairs, mirroring the trading which we simulated using top pairs. We simply used the function rand in Excel to choose our pairs. As can bee seen in table 1 below we achieved a significant and negative return trading random portfolios without a stop loss, indicating that our sample period include tendencies for returns to have auto-correlation. If we apply a stop loss we do not achieve negative or positive returns significantly different from zero.

Table 7.1.1 – random pairs, no stop loss.

No stop loss - random pairs	**Return on total pairs** Pair 1 - 5	Pair 1 - 20	**Return on traded pairs** Pair 1 - 5	Pair 1 - 20
Average	**-4.24%**	**-2.84%**	**-5.44%**	**-3.86%**
Median	-3.2%	-2.4%	-4.8%	-3.5%
Stdev	13.7%	8.9%	19.2%	12.8%
Max Loss	-46.2%	-27.6%	-55.4%	-44.1%
Max Profit	47.9%	28.0%	47.9%	43.1%
Periods >0	37.5%	37.5%	35.8%	37.5%
Z	-3.40	-3.50	-3.10	-3.31
Significant	TRUE	TRUE	TRUE	TRUE

Table 7.1.2 – random pairs, including a stop loss.

	10% stop loss		20% stop loss	
Stop loss - random pairs	**Total pairs** Pair 1 - 20	**Traded pairs** Pair 1 - 20	**Total pairs** Pair 1 - 20	**Traded pairs** Pair 1 - 20
Average	**0.53%**	**0.88%**	**-0.49%**	**-0.55%**
Median	-0.2%	-0.2%	-0.6%	-0.7%
Stdev	4.3%	6.5%	5.4%	8.1%
Max Loss	-7.1%	-10.0%	-13.1%	-18.3%
Max Profit	16.4%	25.0%	15.1%	24.7%
Periods >0	48.3%	48.3%	44.2%	44.2%
Z	1.33	1.48	-1.00	-0.75
Significant	FALSE	FALSE	FALSE	FALSE

7.2 Free formation and no stop loss

Table 1 below shows the total return of the top 20 and top 5 pairs achieved using free formation i.e. without taking into consideration the industry category whilst forming pairs. As can be seen in the table the average return looking at all pairs is a negative 3.93%. The total number of pairs traded is 2400 (20pairs*12months*10years). The standard deviation shown in the table is calculated as period standard deviation; i.e. standard deviation of the return for the 120 periods. The return is the same for individual pairs as it is for a period since we divide the total absolute return with the total number of pairs in a portfolio (20 and 5 respectively). When we analyse only the traded pairs their will be a small deviation due to different pairs having different weighting in the averages.

Table 7.2.1 – Portfolio return

No stop loss - all pairs	**Return on total pairs**		**Return on traded pairs**	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**-1.54%**	**-3.93%**	**-0.99%**	**-3.43%**
Median	1.9%	0.9%	2.2%	1.2%
Stdev	15.3%	25.9%	18.1%	29.7%
Max Loss	-98.1%	-203.8%	-122.6%	-226.4%
Max Profit	27.7%	26.3%	46.2%	43.8%
Periods >0	61.7%	57.5%	60.8%	57.5%
Z	-1.10	-1.66	-0.60	-1.26
Significant	FALSE	FALSE	FALSE	FALSE

No Stop loss - all pairs	**Return on total pairs**		**Return on traded pairs**	
Including trading costs	Pair 1-5	Pair 1-20	Pair 1-5	Pair 1-20
Average	-2.04%	-4.40%	-1.59%	-4.03%
Median	1.5%	0.4%	1.6%	0.6%
Stdev	15.3%	25.9%	18.1%	29.7%
Max Loss	-98.5%	-204.3%	-123.2%	-227.0%
Max Profit	27.4%	25.9%	45.6%	43.2%
Periods >0	53.3%	53.3%	58.3%	53.3%
Z	-1.46	-1.86	-0.96	-1.48
Significant	FALSE	FALSE	FALSE	FALSE

Of the 2400 pairs that were traded 520 (28%) never opened and hence had a return of 0% and standard deviation of 0%. Looking at the reminding 1880 pairs (72%) the return becomes more negative and the standard deviation increases.

Table 7.2.2 – Individual pair return

No stop loss - all pairs	Return on total pairs		Return on traded pairs	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**-1.54%**	**-3.93%**	**-1.86%**	**-5.02%**
Stdev	28.8%	54.7%	31.6%	61.8%
Max Loss	-379.9%	-850.4%	-379.9%	-850.4%
Max Profit	68.6%	210.7%	68.6%	210.7%
Pairs >0	276	999	276	999
Pairs <0	141	881	141	881
Pairs =0	103	520	0	0

Top 5 pairs yield a similar result.

Of the 1880 pairs that opened 999 (53%) yielded a positive return. The overall negative return is hence a consequence of the larger negative return from the pairs that yielded a negative return. The return standard deviation amongst the negative yielding pairs is also substantially higher. This result is quite unsurprising given that the losses are unlimited whereas the profits are capped at around 2 standard deviations of the returns of the respective pairs. Recall that our first trading rule closed the position when we reached 2 standard deviations positive returns but did not include any stop loss.

Table 7.2.3

No stop loss - traded pairs	Pair 1 - 5	Pair 1 - 20
Total N, positive returns	276	999
as % of traded pairs	55.5%	53.1%
Average positive return	12.4%	15.8%
STD of pairs with positive returm	11.3%	20.6%
Total N, negative returns	221	881
as % of total pairs	44.5%	46.9%
Average positive return	-19.6%	-28.6%
STD of pairs with negative returm	39.0%	81.3%

7.3 Negative performance predominantly driven by losses within the IT sector

In table 4 below, we highlight the annual total performance from our trading strategy. As can be seen the simulated trading strategy delivers substantial negative returns during 1999 and 1998. 2002 also shows negative returns whereas the other years show a positive return.

Table 7.3.1

Free formation	Return on top 20		Return on top 5	
Year	All pairs	Traded pairs	All pairs	Traded pairs
1996	7.1%	9.5%	9.9%	15.6%
1997	3.5%	4.0%	4.5%	4.4%
1998	-17.1%	-18.0%	-21.5%	-22.8%
1999	-93.6%	-103.8%	-37.4%	-42.8%
2000	10.0%	14.7%	9.9%	16.1%
2001	20.5%	31.0%	10.3%	16.1%
2002	-4.0%	-3.3%	-15.7%	-19.8%
2003	1.7%	4.8%	4.6%	5.6%
2004	-0.4%	-0.1%	3.8%	5.5%
2005	-6.3%	-7.2%	0.7%	0.7%

We have plotted the 20 worst trades during the 10 year trading period. The 20 trades show a consistent pattern. All 20 consist of a short position in the software company Enea versus a long position in another IT/Telecom stock. All also appear in the portfolios that started to trade the autumn of 1999 (and hence was traded up until June 2000), when Enea appreciated by over 800%.

Table 7.3.2 – worst 20 trades, free formation, no stop loss

Stock 1	Stock 2	Year	Month	Return	Return rank
ENEA SS EQUITY	MOD1 SS EQUITY	1999	11	-850%	2400
ENEA SS EQUITY	KNOW SS EQUITY	1999	9	-828%	2399
ENEA SS EQUITY	MOD1 SS EQUITY	1999	10	-820%	2398
ENEA SS EQUITY	MOD1 SS EQUITY	1999	9	-804%	2397
ENEA SS EQUITY	WMB SS EQUITY	1999	10	-662%	2396
ENEA SS EQUITY	WMB SS EQUITY	1999	9	-655%	2395
ENEA SS EQUITY	KNOW SS EQUITY	1999	8	-516%	2394
ENEA SS EQUITY	TELCB SS EQUITY	1999	9	-496%	2393
ENEA SS EQUITY	TELCB SS EQUITY	1999	10	-493%	2392
ENEA SS EQUITY	INDUA SS EQUITY	1999	9	-444%	2391
ENEA SS EQUITY	MAND SS EQUITY	1999	11	-418%	2390
ENEA SS EQUITY	MOD1 SS EQUITY	1999	8	-409%	2389
ENEA SS EQUITY	TEL2B SS EQUITY	1999	10	-402%	2388
ENEA SS EQUITY	TEL2B SS EQUITY	1999	9	-398%	2387
ENEA SS EQUITY	MAND SS EQUITY	1999	10	-380%	2386
ENEA SS EQUITY	MAND SS EQUITY	1999	12	-369%	2385
ENEA SS EQUITY	MOD1 SS EQUITY	1999	12	-366%	2384
ENEA SS EQUITY	MAND SS EQUITY	1999	9	-332%	2383
ENEA SS EQUITY	MOD1 SS EQUITY	1999	7	-208%	2382
ENEA SS EQUITY	KNOW SS EQUITY	1999	7	-199%	2381
Average Return				**-503%**	

In total Enea took part in 66 trades. 36 of the times the trade was negative, 17 it was positive and 13 of the times it did not open. The average return was -164% with a standard deviation of 254%..

Table 7.3.3 – all Enea trades (next page)

Stock 1	Stock 2	Year	Month	Return	Return rank
ENEA SS EQUITY	MOD1 SS EQUITY	1999	11	-850%	2400
ENEA SS EQUITY	KNOW SS EQUITY	1999	9	-828%	2399
ENEA SS EQUITY	MOD1 SS EQUITY	1999	10	-820%	2398
ENEA SS EQUITY	MOD1 SS EQUITY	1999	9	-804%	2397
ENEA SS EQUITY	WMB SS EQUITY	1999	10	-662%	2396
ENEA SS EQUITY	WMB SS EQUITY	1999	9	-655%	2395
ENEA SS EQUITY	KNOW SS EQUITY	1999	8	-516%	2394
ENEA SS EQUITY	TELCB SS EQUITY	1999	9	-496%	2393
ENEA SS EQUITY	TELCB SS EQUITY	1999	10	-493%	2392
ENEA SS EQUITY	INDUA SS EQUITY	1999	9	-444%	2391
ENEA SS EQUITY	MAND SS EQUITY	1999	11	-418%	2390
ENEA SS EQUITY	MOD1 SS EQUITY	1999	8	-409%	2389
ENEA SS EQUITY	TEL2B SS EQUITY	1999	10	-402%	2388
ENEA SS EQUITY	TEL2B SS EQUITY	1999	9	-398%	2387
ENEA SS EQUITY	MAND SS EQUITY	1999	10	-380%	2386
ENEA SS EQUITY	MAND SS EQUITY	1999	12	-369%	2385
ENEA SS EQUITY	MOD1 SS EQUITY	1999	12	-366%	2384
ENEA SS EQUITY	MAND SS EQUITY	1999	9	-332%	2383
ENEA SS EQUITY	MOD1 SS EQUITY	1999	7	-208%	2382
ENEA SS EQUITY	KNOW SS EQUITY	1999	7	-199%	2381
ENEA SS EQUITY	WMB SS EQUITY	2000	1	-194%	2380
ENEA SS EQUITY	TEL2B SS EQUITY	1999	8	-189%	2379
ENEA SS EQUITY	TELCB SS EQUITY	1999	8	-165%	2378
ENEA SS EQUITY	MAND SS EQUITY	1999	8	-143%	2375
ENEA SS EQUITY	TEL2B SS EQUITY	1999	7	-115%	2374
ENEA SS EQUITY	WMB SS EQUITY	2000	2	-88%	2364
ENEA SS EQUITY	KNOW SS EQUITY	1999	6	-75%	2350
ENEA SS EQUITY	TELCB SS EQUITY	1999	7	-63%	2332
ENEA SS EQUITY	WMB SS EQUITY	2000	3	-42%	2296
ENEA SS EQUITY	MOD1 SS EQUITY	1999	6	-41%	2286
ENEA SS EQUITY	KNOW SS EQUITY	1999	5	-21%	2157
ENEA SS EQUITY	KNOW SS EQUITY	1999	4	-15%	2077
ENEA SS EQUITY	TELCB SS EQUITY	1999	6	-15%	2058
ENEA SS EQUITY	HMB SS EQUITY	1999	2	-8%	1883
ENEA SS EQUITY	TLOG SS EQUITY	2000	11	-5%	1789
ENEA SS EQUITY	HMB SS EQUITY	1999	3	-1%	1572
ENEA SS EQUITY	MOD1 SS EQUITY	1999	5	0%	1381
ENEA SS EQUITY	TELCB SS EQUITY	1999	5	0%	1357
ENEA SS EQUITY	TELCB SS EQUITY	1999	1	0%	1326
ENEA SS EQUITY	TELCB SS EQUITY	1999	2	0%	1325
ENEA SS EQUITY	TELCB SS EQUITY	1999	3	0%	1296
ENEA SS EQUITY	TELCB SS EQUITY	1999	4	0%	1295
ENEA SS EQUITY	TLOG SS EQUITY	2001	6	0%	1290
ENEA SS EQUITY	TELCB SS EQUITY	1998	11	0%	1214
ENEA SS EQUITY	ANODB SS EQUITY	2000	11	0%	1185
ENEA SS EQUITY	TELCB SS EQUITY	1998	12	0%	1169
ENEA SS EQUITY	TLOG SS EQUITY	2001	7	0%	1079
ENEA SS EQUITY	KNOW SS EQUITY	1999	1	0%	1026
ENEA SS EQUITY	TLOG SS EQUITY	2001	5	0%	1017
ENEA SS EQUITY	MAND SS EQUITY	1998	11	2%	886
ENEA SS EQUITY	TEL2B SS EQUITY	1999	3	3%	852
ENEA SS EQUITY	TLOG SS EQUITY	2000	10	3%	835
ENEA SS EQUITY	TEL2B SS EQUITY	1999	4	9%	574
ENEA SS EQUITY	MAND SS EQUITY	1999	1	12%	460
ENEA SS EQUITY	TEL2B SS EQUITY	1999	6	14%	390
ENEA SS EQUITY	HMB SS EQUITY	1999	4	14%	385
ENEA SS EQUITY	TEL2B SS EQUITY	1999	1	15%	384
ENEA SS EQUITY	HMB SS EQUITY	1999	1	20%	236
ENEA SS EQUITY	MAND SS EQUITY	1999	3	21%	216
ENEA SS EQUITY	TEL2B SS EQUITY	1999	2	32%	91
ENEA SS EQUITY	MAND SS EQUITY	1999	5	33%	87
ENEA SS EQUITY	TEL2B SS EQUITY	1998	12	36%	76
ENEA SS EQUITY	MAND SS EQUITY	1999	2	36%	75
ENEA SS EQUITY	MAND SS EQUITY	1999	6	40%	61
ENEA SS EQUITY	TEL2B SS EQUITY	1999	5	40%	60
ENEA SS EQUITY	MAND SS EQUITY	1999	4	46%	49
Average Return				**-164%**	
STD				**254%**	

Another interesting observation is that if we exclude the portfolios constructed in 1999, the return changes from negative to positive.

Table 7.3.4

No stop loss - all top 20 pairs excluding 1999	All pairs	Traded pairs
Average	**0.83%**	**1.07%**
Stdev	21.9%	24.8%
Max Loss	-194.4%	-194.4%
Max Profit	210.7%	210.7%
Pairs =0	478	0
Traded pairs	2160	1682
Z	1.77	1.77
Significant	True	True

7.4 Free formation using a stop loss

Free formation 20% stop loss	**Return on top 20**		**Return on top 5**	
Year	**All pairs**	**Traded pairs**	**All pairs**	**Traded pairs**
1996	6.7%	9.0%	9.9%	15.6%
1997	3.2%	3.7%	4.3%	4.2%
1998	-3.0%	-2.8%	-12.9%	-13.6%
1999	-2.9%	-2.6%	-2.7%	-4.0%
2000	6.6%	9.9%	10.1%	16.1%
2001	8.3%	13.3%	3.6%	9.0%
2002	-2.5%	-1.6%	-10.0%	-12.6%
2003	2.6%	5.8%	4.3%	5.2%
2004	-0.3%	0.1%	3.8%	5.5%
2005	-4.5%	-5.2%	1.5%	1.5%

In table 10 above we show that adding a stop loss of 20% improves the returns. Measuring period return only based on traded pairs, yields a significantly positive return. The other three categories yield a small and insignificant return.

Worth noting is that despite using a stop loss (which naturally would have a bias towards many small losses) we still achieve positive returns on a majority of the portfolios.

Table 7.4.1

10% stop loss - all pairs	Return on total pairs		Return on traded pairs	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**1.27%**	**1.00%**	**2.15%**	**1.75%**
Median	1.2%	0.4%	1.3%	0.6%
Stdev	6.6%	3.9%	8.5%	5.4%
Max Loss	-10.0%	-7.7%	-10.0%	-9.6%
Max Profit	25.4%	13.8%	42.4%	17.8%
Periods >0	55.8%	55.8%	55.0%	55.8%
Z	2.12	2.83	2.76	3.55
Significant	TRUE	TRUE	TRUE	TRUE

Above we have used a stop loss of 10%. It is interesting to note the significant difference using stop losses on 10% versus 20%. We found of particular interest the fact that despite the stop loss kicking in at half the loss, we only reduced the number of positive periods by 1-2%. Looking at all individual pairs that opened (1880) 50.2% had positive returns with a stop loss of 20% versus 42% using a stop loss of 10%. However, another important observation is that a substantial number of all trades did close out at the stop loss. When we applied a stop loss of 10%, close to 80% of the negative trades closed out with the stop loss. It is probably reasonable to assume that in those “forced” exits the trading costs would be substantially higher that just the commission costs on average, reducing the 1% average return on the top 20 pairs.

7.5 Trading industry pairs

Table 7.5.1

No stop loss - all industry pairs	Return on total pairs		Return on traded pairs	
	Pair 1 - 5	Pair 1 - 20	Pair 1 - 5	Pair 1 - 20
Average	**-7.33%**	**-2.44%**	**-8.34%**	**-2.70%**
Median	1.9%	0.7%	2.3%	1.1%
Stdev	52.4%	22.8%	65.2%	28.7%
Max Loss	-456.1%	-160.5%	-570.1%	-200.6%
Max Profit	30.7%	20.0%	51.2%	30.8%
Periods >0	59.2%	57.5%	58.3%	57.5%
Z	-1.53	-1.17	-1.40	-1.03
Significant	FALSE	FALSE	FALSE	FALSE

Table 7.5.2

Industry pairs, free formation	Return on top 20		Return on top 5	
Year	**All pairs**	**Traded pairs**	**All pairs**	**Traded pairs**
1996	2.1%	3.2%	7.6%	12.1%
1997	5.1%	6.1%	-3.3%	-4.4%
1998	2.7%	3.4%	-20.1%	-23.3%
1999	-79.5%	-98.2%	-170.9%	-205.3%
2000	13.0%	16.9%	29.2%	39.8%
2001	14.5%	21.7%	13.5%	18.8%
2002	-4.5%	-4.7%	-10.0%	-14.3%
2003	1.2%	1.8%	4.6%	5.6%
2004	0.3%	0.0%	4.1%	5.7%
2005	-3.6%	-4.2%	-1.4%	-1.4%

Again it is the losses within the IT sector in 1999 that contribute to the negative performance. Excluding 1999 the return increases from a negative 2.4% to a positive 1.7% (average return per trade in 1999 is a negative 40%). 73 pairs include a long position in Enea, with an average return of a negative 145%. Excluding either 1999, the technology sector as a whole or just Enea all yield significant and positive returns.

Table 7.5.3

No stop loss - top 20 pairs, excluding technology		
	All pairs	Traded pairs
Average	**1.34%**	**1.79%**
Stdev	15.5%	17.9%
Max Loss	-82.7%	-82.7%
Max Profit	68.6%	68.6%
Pairs >0	886	886
Pairs <0	656	656
Pairs =0	523	0
Traded pairs	2065	1542
Z	4	4
Significant	True	True

Adding a stop loss of 20% increases the returns from a negative 2.2% to a positive 1.6%, but these improvements are exclusively explained by the improvements within the technology sector, where returns improve from a negative 26% to a positive 2%. Excluding the technology sector the changes are small and insignificant.

Table 7.5.4 – Trading industry pairs including trading costs. No stop loss.

Industry, including trading costs, no stoploss	Total pairs Pair 1 - 5	Pair 1 - 20	Traded pairs Pair 1 - 5	Pair 1 - 20
Average	**-7.98%**	**-3.02%**	**-9.13%**	**-3.47%**
Median	1.4%	0.3%	1.7%	0.4%
Stdev	52.5%	22.9%	65.3%	28.8%
Max Loss	-457.5%	-161.5%	-571.8%	-201.9%
Max Profit	30.1%	19.5%	50.2%	30.0%
Periods >0	51.7%	51.7%	54.2%	51.7%
Z	-1.66	-1.44	-1.53	-1.32
Significant	FALSE	FALSE	FALSE	FALSE