

Beating the Benchmark - an Active Trading Strategy Based on Contrarian Trends and Jump Detection

Gustav Fernström and Linus Svensson

Bachelor Thesis

Lund University, Department of Economics

Supervisor: Hossein Asgharian

June 8, 2006

Abstract

Extreme price changes are identified in the stock market and related to jumps in the price process. After such extreme events, a reverting (also called contrarian) trend is examined. There is evidence that the initial price change is often an overreaction, caused by the psychology of market participants, and that the subsequent reversal corrects for the temporary inefficiencies in the market. Confirming the empirical findings by other studies, it is found that larger extreme price changes are followed by a stronger reversal effect.

From a statistical analysis of the extreme events, a systematic trading strategy is formed. After an extreme price change, losing stocks are bought with the hopes of following a reverting trend. In the long run, probability is assumed to be in favor of the strategy, generating significant positive returns.

Returns from trading are evaluated for both an estimation and validation period. Trading performance is excellent during the estimation period, beating the defined benchmark. During validation, the active trading also generates promising results, outperforming the benchmark. It is however suspected that the nature of the extreme events change depending on current market conditions. During strong positive market trends, the reverting trends are likely to be more significant.

Keywords: Contrarian trading, Extreme price changes, Jump detection, Black-Litterman

Acknowledgements

We would like to thank our supervisor, Associate Professor Hossein Asgharian at Lund University, Department of Economics, for supporting us in our work and helping us decide on empirical issues. We are grateful that we were allowed to use the data set originally prepared and cleaned by PhD candidates Mia Holmfeldt and Markus Larsson. The idea to study losing stocks in the Swedish market was initially sparked by Christer Andersson at Systematiska Fonder.

Lund, June 6th 2006

Gustav Fernström and Linus Svensson

Contents

1	Introduction	6
1.1	Background	6
1.2	Goal and Purpose	7
1.3	Delimitations of the Study	8
1.4	A Study of Quantitative Trading Based on Logics	8
1.5	Study Overview	9
2	Theoretical Concepts and Methods	10
2.1	Financial Modeling Through Diffusion Processes	10
2.2	Jump Diffusion Processes	10
2.3	Realized Variance and Bipower Variation for Jump Diffusion Processes	11
2.4	A Technique for Jump Detection	12
2.5	Sampling Issues for Relating Jumps to Extreme Events	12
2.6	Identifying Extreme Events Using Empirical Distributions of Returns	13
2.7	Mean-Variance Theory for Forming Portfolios	14
2.7.1	The Risk Aversion Coefficient	15
2.7.2	The Covariance of the Risky Assets	15
2.7.3	The Risk Premiums of the Assets	15
2.7.4	Finding Optimal Portfolio Weights	16
2.7.5	Implementation Issues	16
2.8	The Black-Litterman Model	17

2.8.1	Advantages of the Black-Litterman Model	18
2.8.2	Model Inputs	19
2.8.3	The Neutral View Risk Premiums	19
2.8.4	Formulating Views	20
2.8.5	Finding the BL Portfolio Weights	21
2.9	Defining a Benchmark	22
2.10	Obtaining Exceptional Returns	22
2.11	Evaluating Risk-Adjusted Returns	23
3	Data	24
3.1	Data Criticism	25
4	Methodology and Empirical Results	27
4.1	Identifying Extreme Events	27
4.2	Measuring the Reversal Effect	28
4.3	Examples of Extreme Events	31
4.4	Implementing a Portfolio Strategy	32
4.4.1	The Benchmark and the Simple Form of the View Portfolio	32
4.4.2	Estimating the Input Parameters	33
4.4.3	Calculating View Premium from the Expected Returns	34
4.4.4	Estimating the View Confidence	35
4.4.5	Setting the Risk Aversion Coefficient	36
4.5	Results from Portfolio Strategies	36

4.6	Implementing Trading with Fixed Bets	38
4.7	Details on Trades for Individual Stocks	42
5	Conclusion	44
5.1	Suggestions for Future Studies	44
A	Empirical Results Appendix	49
A.1	Regressions of Expected Return	49
A.2	Evaluating Trades for Individual Stock	50
A.3	Positions from Portfolio Implementation	52
B	Technical Appendix	54

1 Introduction

1.1 Background

Is buying recent losers profitable in the stock market? Popular belief seems to indicate that abnormal losses should recoil with a reverting trend where prices return to levels more in line with fundamental values. By identifying such events of extreme price changes, a strategy may be formed for systematic trading of losing stocks based on purely technical indicators.

The Efficient Market Hypothesis, as presented by Fama (1970), postulates that a price set by the market is fair, i.e. always represents the true value of an asset based on the available set of information. Studies of financial markets find that fair prices are rarely observed empirically. Instead, market participants frequently under- or overreact to news releases, thus making mispriced assets rather a rule than an exception. Chiarella and He (2002) state that it is likely to believe that not all investors have fully assessed a new piece of information immediately. Information slowly diffuses across investors and may cause predictability in time series of asset prices. Momentum investment is an example of such a strategy where one hypothesizes that the market under-reacts to news, thus making it profitable to follow the trend by buying recent winners. Jegadeesh and Titman (2001) is just one of many studies that provide empirical evidence of momentum profitability.

Contrarian investments are the opposite of momentum, following the philosophy that abnormal price changes are likely to be overreactions with a subsequent reversal effect. Especially over the short term, it is reasonable to believe that it is trading psychology, (like fear and herd behavior) which causes the market to overreact. Fehle and Zdorovtsov (2005) examine stocks in the US market that have experienced great losses and measure an over-night adjustment effect that suggests the opportunity of forming profitable contrarian strategies.

Since Fehle and Zdorovtsov (2005) evaluate events at fixed time intervals, only examining daily returns at the end of the trading day, they run the risk of entering the reverting trend at a late stage. In a previous study, Svensson (2006) explores similar patterns in the

foreign exchange market, where a statistical methodology is developed for continuously identifying extreme events at a high-frequency time scale. Svensson (2006) also relates events to jumps in the price process, detected by using formal tests developed by Tauchen and Zhou (2005). The jump detection technique relies on the availability of high-frequency intraday data and utilizes the volatility measures realized variance and bipower variation, as defined by Barndorff-Nielsen and Shephard (2003). Following the reasoning by Tauchen and Zhou (2005), it is reasonable to believe that jumps are rare and large, thus associating them with extreme price changes.

The theoretical conclusions from the Capital Asset Pricing Model (CAPM), as drawn by for example Mossin (1966), suggest that the best investment strategy is a passive one, holding the market portfolio. In practice, such a portfolio would be equivalent of investing in some index, for example the S&P 500 for US stocks. Numerous studies, e.g. Cumby and Glen (1990), Malkiel (1995) and Elton et al. (1996), provide evidence for actively managed funds underperforming their relevant indexes. It would thus be both academically and practically interesting to examine opportunities for beating the benchmark index.

For forming portfolios of financial assets, the mean-variance framework may be applied, formulated by Markowitz (1952). This model states the main objectives of asset management; maximizing return and minimizing risk, defined as the standard deviation of returns. Black and Litterman (1990) stress some issues that make the mean-variance theory hard to apply in practice and suggest their own modifications. Black-Litterman methodology is ideal for studying how active investment management adds value, since it is based on the benchmark portfolio and only adjusts portfolio weights when an investor has a unique view about an asset.

1.2 Goal and Purpose

The presence of profitable contrarian trends is examined in the Swedish stock market. By taking the approach of Svensson (2006), statistical methods are utilized for identifying extreme events. The time window for the study is stretched from the intraday perspective in Svensson (2006) to also include trends over multiple days before and after the event.

Inspired by the philosophies of Grinold and Kahn (2000), the intention is to develop a quantitative investment strategy that relies on probability for making profits. Such a quantitative approach is disciplined, generating systematic rules for trading and taking out any feelings from the decision-making.

The goal is to beat a defined benchmark by adding value from an active strategy that follows short-term contrarian trends. For implementing the strategy in portfolios, Black-Litterman methodology is applied, where this paper also examines the necessary practical considerations. A simpler approach for implementation is also examined, which circumvents the many assumptions and estimations that are associated with the Black-Litterman model.

The study intends to generate valuable insights for traders, as well as adding to empirical research, collecting further evidence for overreactions and contrarian trends.

1.3 Delimitations of the Study

The available data is limited, relying on an already cleaned data set. Due to time restrictions, additional data collection has been left outside the scope of the study. The methods applied and the programs developed¹ are however fully compatible for handling larger sets of data, but it is left for future studies to do so.

The reader is assumed to be at a Bachelor level of financial economics. The study further assumes that the reader is well acquainted with basic statistical methods, not fully covering the theoretical foundations of such techniques. For full understanding, it is also preferable with some experience of basic matrix algebra.

1.4 A Study of Quantitative Trading Based on Logics

Technical trading explores patterns in the price process of financial assets and examines indicators that have proven profitable in the past. By relying on statistics from historical events, one runs the risk of over-fitting to data, i.e. optimizing a strategy for maximizing

¹All statistical analysis is performed in MathWorks Matlab 7.0.4.

profits on the current data set, while performance is poor for future data². However, if a pattern can be motivated by economic theory, it is reasonable to believe that it will also be present in out-of-sample data.

Over the short-term investment horizons considered in this study, the hypothesis is that the market is mostly affected by the behavior of the active economic agents and that prices do not follow theories about efficient markets. The contrarian behavior examined is strongly related to human psychology, which may be assumed to have similar properties over time. Any contrarian trends that are statistically significant are thus well motivated by logics and trading indicators derived from past events are likely to be useful for developing a quantitative strategy for the future.

1.5 Study Overview

The study is organized as follows. Section 2 describes the related theoretical concepts, where several mathematical and statistical methods are presented. It should be noted that full understanding of the technical background is not required for an overall understanding of the study. Section 3 describes the data set and Section 4 describes how theory is implemented in practice, certain considerations and adjustments that are made to the current case and the resulting empirical findings. Section 5 concludes and makes suggestions for further studies. An Appendix contains some extra details on the study.

²See for example Qi and Wu (2005) for an evaluation of technical trading rules and how they relate to over-fitting.

2 Theoretical Concepts and Methods

2.1 Financial Modeling Through Diffusion Processes

Stochastic differential equations (SDEs) are a standard approach in financial mathematics for describing how prices of financial assets evolve over time. A standard SDE typically has two major components, a drift term and a diffusion term. The drift represents the deterministic trend over the time interval considered, while the randomness is introduced by the so-called diffusion term. A Brownian Motion (BM) is the stochastic process that is most frequently utilized for introducing the random walk and would assume normally distributed noise added to the drift.

An SDE with drift and diffusion would be written as

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t, \quad (1)$$

where X_t is the price of a financial asset at time t , with dX_t denoting returns. $\mu(t, X_t)$ is the drift term, $\sigma(t, X_t)$ is the diffusion term and dW_t represents the BM.

For a more elaborate discussion on the theoretical foundations of diffusion processes, see for example Madsen et al. (2004) or Björk (2004).

2.2 Jump Diffusion Processes

The randomness introduced by the BM in (1) would suggest normally distributed returns, which is contradicted by empirical findings, e.g. in Tauchen and Zhou (2005). An appropriate model that fully explains stock returns would also need to include a jump process, thus leading to a jump diffusion process of the form

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t + \sum_{j=1}^{N(t)} \kappa_j, \quad (2)$$

where the process $N(t)$ counts the number of jumps occurring with jump size $\kappa(j)$.

2.3 Realized Variance and Bipower Variation for Jump Diffusion Processes

Recent derivation of new volatility measures for high-frequency data has made it possible to decompose the total variation of the observable price process into its continuous and jump part. The first measure utilized is the quadratic variation of returns, or realized variance, defined as follows for day t

$$RV_t \equiv \sum_{j=1}^m r_{t,j}^2 \quad (3)$$

where the trading day t is divided into $\Delta = 1/m$ sampling intervals and $r_{t,j}$ is the return for each interval during the day. m denotes the total number of samples taken for the time period, i.e. during one day. Barndorff-Nielsen and Shephard (2003) show that as sampling becomes highly frequent ($m = 1/\Delta \rightarrow \infty$), the realized variance converges to the total variation of the price process,

$$RV_t \rightarrow \int_{t-1}^t \sigma^2(s) ds + \sum_{j=N(t-1)+1}^{N(t)} \kappa^2(s_j). \quad (4)$$

The second measure utilized is the bipower variation, defined as

$$BV_t \equiv \frac{\pi}{2} \frac{m}{m-1} \sum_{j=2}^m |r_{t,j}| |r_{t,j-1}| \quad (5)$$

If log returns for prices follow the jump diffusion process (2), Barndorff-Nielsen and Shephard (2004) show that with some additional restrictions and the assumption that $\sigma(s)$ is independent of the Brownian Motion W_t , BV_t converges to the variation of the continuous part as $m = 1/\Delta \rightarrow \infty$

$$BV_t \rightarrow \int_{t-1}^t \sigma^2(s) ds \quad (6)$$

2.4 A Technique for Jump Detection

BV_t and RV_t may be utilized for detecting jumps in the price process, since the measures should be identical when there are no jumps, but differ when the jump process is active. However, since there are practical limitations of sampling frequency, measurement errors cause the difference between the volatility measures to almost always be different from zero. A statistical test is needed where jumps are filtered out at some level of significance and such a techniques are discussed in Tauchen and Zhou (2005). After deriving the normalized jump variable ZJ_t , described in the Appendix B, one can filter out the return associated with the jumps realized during the trading day t as

$$J_t = \text{sign}(r_t) \sqrt{(RV_t - BV_t) I_{t, (ZJ_t \geq \Phi_\alpha^{-1})}}, \quad (7)$$

where Φ is the cumulative distribution function of the standard normal, α is the significance level of the test and $I_{t, (ZJ_t \geq \Phi_\alpha^{-1})}$ is a binary variable that takes the value one when there is a jump during the day. $RV_t - BV_t$ may be negative due to the presence of measurement errors, but this is never a problem in (7) as the test is one-sided, only filtering out jump events using the right tail.

2.5 Sampling Issues for Relating Jumps to Extreme Events

From economic intuition it is reasonable to believe that a jump is rare and large; dominating total return during the time period it occurs (Tauchen and Zhou (2005)). Such events are highly likely to be related to extreme price changes. However, with the filtering technique in (7), one is only able to identify that a jump has occurred at some unknown point during a daily time interval. For accurately isolating the reversal effect, it is considered crucial to find the precise event time of the extreme price change.

One could imagine that very frequent sampling of data could allow jump detection on an hourly basis instead of daily. However, as sampling becomes more frequent, the BV_t and RV_t estimates are spoiled by microstructure biases in tick data, such as price discreteness and price reporting errors. Fleming and Paye (2005) examine different sampling frequen-

cies and find that five minute intervals are optimal, thus making it inappropriate to set the time interval, t , too short. BV_t and RV_t rely on asymptotic convergence and there is a significant difference between setting t to one hour, including twelve five minute intervals, and setting t to one day, including 288 five minute intervals.

2.6 Identifying Extreme Events Using Empirical Distributions of Returns

For identifying event times accurately, it is necessary to use other methods in addition to formal jump detection. Svensson (2006) explores techniques for comparing the current return to its empirical distribution and labeling it as extreme if the return exceeds some cut-off level, i.e. a high or low quantile of the distribution. Zawadowski et al. (2004) stress the importance of adjusting for the U-shaped curvature of volatility during a trading day, thus making it necessary to introduce different distributions and cut-off levels depending on the hour of the day. For finding extreme events on intraday data, it is then suggested to filter on both relative and absolute terms, where the relative comparison is made with the “normal” level for the given hour in addition to comparing with a fixed, absolute level that is constant over the current day.

Svensson (2006) finds that an adaptive approach is favorable for estimating the empirical distribution of returns, thus automatically dealing with any properties of non-stationarity (time variations in mean and variance of the distribution). Adaptation is achieved by utilizing a moving estimation window. Setting the window length to T means that all T days leading up to the current day t are used for deriving the empirical distribution applicable for day t . For day $t + 1$, the return for day t is included in the data set for estimation and the last observation is excluded on a first-in-first-out basis. In comparison, other studies of extreme price changes like Fehle and Zdorovtsov (2005) use constant extreme cut-off levels for the entire time period considered.

The cut-off levels for the empirical distribution will vary over time, but are determined at a fixed level of significance, i.e. a constant threshold for the level of extremity. In Figure 1, an example is presented for data simulated from a standard normal distribution to demonstrate the concept of finding cut-off levels for a data set.




Figure 1: Cut-off levels marked with triangles at the significance level 1%.

When estimating the empirical distribution, there is clearly a trade-off that has to be considered before settling for a length of the estimation period. On one hand, it is better to only include returns that are relevant for the current regime of the price process (making a small time frame preferable), while on the other hand, one needs enough observations for making reasonable estimations (making a long time frame preferable). Gustafsson and Hegerin (2005) examine trading opportunities on high-frequency data and use a one month estimation period for distribution parameters in the foreign exchange market.

2.7 Mean-Variance Theory for Forming Portfolios

In the mean-variance theory by Markowitz (1952), the investors choose their optimal portfolio weights based on their individual risk aversion, A , the covariance of the risky assets, Σ , and the risk premiums of the risky assets, ν . In addition to the risky assets, the portfolio constructed by mean-variance theory includes weight of a risk free asset, w_0 , which is used to sum the portfolio weights to one.

2.7.1 The Risk Aversion Coefficient

By definition, the risk aversion coefficient, A , is the rate at which an investor will forego expected return for less variance. That is, the higher the risk aversion of the investor, the more readily expected return is forsaken in exchange for less risk (Idzorek (2004)). Due to its subjective nature, it is hard to empirically quantify the risk aversion coefficient and it is usually set to a reasonable value depending on the context.

2.7.2 The Covariance of the Risky Assets

The covariance matrix, Σ , is a symmetrical $K * K$ matrix, where K is the number of risky assets available to the investor. The individual variances of the assets can be found along the diagonal of the matrix, while the other positions in the matrix are the covariances of the assets.

$$\Sigma = \begin{pmatrix} var_1 & cov_{1,2} & cov_{1,3} \\ cov_{2,1} & var_2 & cov_{2,3} \\ cov_{3,1} & cov_{3,2} & var_3 \end{pmatrix}$$

When estimating the covariance matrix, it is important to consider the length of the estimation period. Covariance is a noisy measure and the longer the estimation period, the less impact of noise on the estimates. Grinold and Kahn (2000), p. 54, discuss historical estimation of covariances and stress the drawback of a long estimation period, where the changing nature of assets due to mergers, spinoffs and other restructuring may be lost. Grinold and Kahn (2000) also note that mathematics demand the number of time periods, T , used for estimation to be greater than the number of assets, i.e. $T > K$.

2.7.3 The Risk Premiums of the Assets

The risk premiums, ν , are defined as the expected returns of the assets in excess of the risk free rate of return, i.e. the reward of holding risk bearing assets to holding a risk free asset. In the mean-variance framework, the investor has to provide expected returns for

all assets available. It is convenient to put all values of expected returns into a column vector of length K .

2.7.4 Finding Optimal Portfolio Weights

When the inputs from Sections 2.7.1-2.7.3 are attained, a simple matrix operation will give the optimal portfolio weights, w , as a column vector of length K .

$$w = \frac{1}{A} \Sigma^{-1} \nu \quad (8)$$

The resulting weights are a combination of long and short positions which, in combination with the risk free portfolio weight, must sum to one. The optimal weight of the risk free asset, w_0 , is thus defined as

$$w_0 = 1 - w_1 - w_2 - \dots - w_K. \quad (9)$$

The risk free weight is strictly dependent on the risky assets' portfolio weights and can be either negative or positive.

2.7.5 Implementation Issues

Academically, the mean-variance framework works nicely, but when implementing the theory in reality, some major problems occur, as noted by Black and Litterman (1990).

First, the model requires the investor to have a complete set of expected returns of all available assets. The selection of assets is usually wide, especially when constructing a portfolio of individual stocks. Normally, the investor has a relevant view of the expected returns of only a small number of assets, and has to estimate the expected returns of all other available assets. This is often approximated from historical returns which usually provide inaccurate guidance for future returns. The model is very sensitive to the input of expected returns, and even though the expected returns seem reasonable, the portfolio weight outputs are often wild.

Second, there are some major problems resulting from the often unreasonable portfolio weights. The unconstrained model often returns large short positions, but when introducing the constraint of no shorting, the model returns corner solutions in the form of zero-weights for many assets. No shorting is a common constraint and the resulting zero-weights lead a weakly diversified portfolio.

Third, there is no way of influencing the results by the confidence the investor has in each individual view. In the mean-variance framework a view with low certainty, or even an assumption of the expected return of an asset on which the investor has no view, has the same impact on the portfolio weights as one with high certainty. The effect of this lack of distinction is that even a view with very high certainty, which should generate a good return will be expressed along with many very uncertain expected returns.

In order to resolve these practical issues, Black and Litterman (1990) developed their own quantitative asset allocation model.

2.8 The Black-Litterman Model

The Black-Litterman (BL) model combines the mean-variance theory with the concepts of the capital asset pricing model, CAPM. The CAPM explains the excess return of asset A , r_A , defined as the total return less the risk free rate of return, by dividing it into two parts, defined for any given time interval as

$$r_A = \beta_A r_M + \theta_A, \tag{10}$$

where the first term is perfectly correlated with the excess market return, r_M , and θ_A is the component of r_A that has no correlation with the market³. β_A is proportional to the asset's covariance with the market as

$$\beta_A = \frac{Cov(r_M, r_A)}{Var(r_M)}. \tag{11}$$

³ θ_A is also referred to as the idiosyncratic return.

According to CAPM, the expected value of the uncorrelated component is zero, $E[\theta_A] = 0$. This means that the expected return of the asset can be derived solely from the excess market return and β_A .

For the market portfolio, the assumption of $E[\theta_A] = 0$ must hold since summing (10) for all stocks will always sum to the total market capitalization. The CAPM, however, takes the concept one step further as it assumes that the assumption holds for all individual assets and portfolios, thus implying that the β reflects all risk associated with an asset or a portfolio. A portfolio's β -risk can be eliminated by diversification, thereby leaving systematic or market risk as the only source of risk exposure. CAPM further implies that no rational investor should take on diversifiable risk, since only non-diversifiable risk, i.e. market risk, is rewarded at market equilibrium. When all diversifiable risk is eliminated, the portfolio β will be one. In a market at equilibrium, given that the underlying assumptions of the CAPM hold, this optimal portfolio is one with weights proportional to the assets' market capitalizations. In other words, the CAPM states that the best investment strategy is a passive one where the average investor holds the market portfolio (Grinold and Kahn (2000)).

The BL model uses the market portfolio as calculated by CAPM as a starting point for active investments. The model enables an investor to tilt the assumed original position in the market portfolio towards his or her own beliefs by the input of views. An investor's view is a piece of information about an asset that is not reflected in the asset price. A view can be either relative, e.g. asset A will outperform asset B by 3 %, or absolute, e.g. asset A will outperform the market expectation by 3 %. If an investor has no expectation of the performance of an asset, he is said to have a neutral view.

2.8.1 Advantages of the Black-Litterman Model

The virtues of the BL model are numerous. The investor no longer needs to make assumptions of expected returns for assets on which the investor has no view. When the investor has a relative view, the portfolio weights will only change for the assets in question. When the view is absolute, however, the portfolio weights of all assets will be decreased in order to favour the weight of the asset with a view. Additionally, the model takes into consid-

eration the degree of confidence the investor has in each view; depending on the excess expected return that can be derived from the investor's view, and the level of confidence the investor has in the excess expected return (i.e. the volatility of the excess expected return), the portfolio is rebalanced in a way that takes the uncertainty of information into account. If, for example, an investor has a view on an asset which will give a high expected return with high confidence, the portfolio will be weighed heavily in that asset to reap as high a reward as possible while still guarding the risk exposure. If the risk is higher, the increase in portfolio weight of the asset will be lowered accordingly.

2.8.2 Model Inputs

To attain the Black-Litterman portfolio weights, a number of inputs are required; the neutral view risk premiums and the views -consisting of the view portfolio, the view premiums and the view confidence. These two rather complex inputs will generate the Black-Litterman risk premiums which are transformed into the Black-Litterman portfolio weights by using regular mean-variance theory.

2.8.3 The Neutral View Risk Premiums

The neutral view risk premiums, ν , are the expected risk premiums by the average investor, implied by the market portfolio weights (therefore also known as the implied excess equilibrium return vector) (Idzorek (2004)). A neutral view means that the investor has no opinion about the returns of an asset, and therefore follows the crowd. The portfolio of an investor with neutral views on all assets on the market, assumed to be the average investor, will hence have the same weights as the market portfolio.

By the inverse of the operation that gave the portfolio weights in mean-variance theory, a vector with the implied neutral view risk premiums can be obtained as

$$\nu_n = A\Sigma w_n, \tag{12}$$

where ν_n is a column vector with neutral view risk premiums, A is the risk aversion

coefficient, Σ is the covariance matrix of the assets and w_n is a column vector of market portfolio weights.

2.8.4 Formulating Views

For each view, the investor should construct a view portfolio, P , which is a vector with a coefficient of plus one for the asset favoured by the view and coefficients summing to minus one for the assets disfavoured by the view. The coefficients of disfavoured assets should be relative to the assets' weights in the market portfolio. For assets not affected by the view, the coefficients are zero. In other words, the view portfolio is a long position in the asset favoured by the view and short positions in assets disfavoured by the view, while taking no position at all in the assets that are unaffected by the view. For multiple views, P is a matrix with rows equal to the number of views and column equal to the number of assets, K .

For example, there may be four available assets for an investor, denoted A, B, C and D. The investor has two views, one absolute view that asset B will outperform the market (the other three assets) and a relative view that asset C will outperform asset D. The resulting view portfolio would be

$$P = \begin{pmatrix} -\frac{1}{3} & 1 & -\frac{1}{3} & -\frac{1}{3} \\ 0 & 0 & 1 & -1 \end{pmatrix}$$

where column one represents the weight in asset A, column two represents the weight in asset B and so on.

The second component of a view is the view premium, Q , which is the expected excess return resulting from the view portfolio compared with the neutral view portfolio. This is also known as the view portfolio alpha.

Continuing the example above, the investor may believe that asset B will outperform the market by 10% in the absolute view and that asset C will outperform asset D by 5% in the relative view. The resulting vector of view premiums would be

$$Q = \begin{pmatrix} 0.10 \\ 0.05 \end{pmatrix}$$

The third component is the view confidence, Ω , which is the degree of certainty the investor feels about the view. Ω is equal to the variance or volatility squared of the view premium. This means that Ω is zero if the investor is fully confident in the view and infinite if the investor has no confidence in the view.

Completing the example, the investor may believe that asset B will outperform the market by 10% with an expected volatility of the return of 20%, rendering a view confidence of 0.20^2 . For the relative view the volatility is expected to be 10%, resulting in the following view confidence matrix

$$\Omega = \begin{pmatrix} 0.20^2 & 0 \\ 0 & 0.10^2 \end{pmatrix}$$

2.8.5 Finding the BL Portfolio Weights

To obtain the Black-Litterman risk premiums, ν_{bl} , some matrix algebra is necessary. The result is a weighted average between the neutral view risk premiums and the view premiums. The weights are inversely proportional to the uncertainty of the views, as represented by the view confidence and the covariance matrix.

$$\nu_{BL} = (\Sigma^{-1} + P'\Omega^{-1}P)^{-1}(\Sigma^{-1}\nu_n + P'\Omega^{-1}Q) \quad (13)$$

With the same operation as in the mean-variance theory, the optimal portfolio weights with respect to the views can be obtained as

$$w_{BL} = \frac{1}{A}\Sigma^{-1}\nu_{BL} \quad (14)$$

Since the view portfolio coefficients for assets unaffected by the view is zero, the Black-Litterman portfolio weights will not deviate from the neutral view portfolio weights for these assets.

2.9 Defining a Benchmark

The market portfolio has previously been mentioned as a benchmark for measuring the performance of an active strategy. In theory, such a portfolio would contain all available assets over the world, which in practice becomes unwieldy and difficult to handle.

A predefined benchmark portfolio is often used instead, where the assets included make up a relevant reference of performance for the individual strategy. The benchmark can just as well be a simple standard strategy, such as a buy and hold, or an old strategy previously developed if the scope of the study is to further develop that strategy (Grinold and Kahn (2000)).

2.10 Obtaining Exceptional Returns

α is a measure of the excess return from the active strategy over a passive strategy. Once a benchmark portfolio has been defined, one can derive exceptional returns of an asset from the expected returns, its correlation with the benchmark, the expected benchmark return, the exceptional benchmark return and the risk free rate of return as presented by Grinold and Kahn (2000), p. 91,

$$\alpha_A = E(R_A) - R_f - \beta_A \mu_B - \beta_A f_B, \quad (15)$$

where μ_B is the expected benchmark return and is estimated in various ways depending on the scope of the study; often as a long-term historical average with time frames ranging up to several decades. f_B is the exceptional benchmark return and represents a short-term trend that is assumed to hold over the future time period considered. R_f is the risk free rate over the time period. Note that β_A is the asset's correlation with the benchmark B and not the same as presented in the discussion of CAPM

$$\beta_A = \frac{Cov(r_A, \mu_B)}{Var \mu_B}. \quad (16)$$

2.11 Evaluating Risk-Adjusted Returns

In Sharpe (1966), a ratio of return compared to risk (defined as volatility of realized returns, σ) is introduced. This so-called Sharpe ratio is a commonly used measure for assessing the compensation for the risk taken and defined as

$$SR = \frac{R_t - R_f}{\sigma}, \quad (17)$$

where R_f is the risk free rate deducted from the average realized return R_t . When calculating the measure on an annual basis, it is reasonable to calculate the average monthly/weekly/daily returns over the full period examined. Returns are then multiplied by the number of time periods for the full period (i.e. 12 for months) and the σ of these returns is multiplied by the square root of the number of time periods (i.e. $\sqrt{12}$ for months). The resulting Sharpe ratio may easily be compared with other investments as the annualized version is the one used by convention.

3 Data

The study is based on stock prices of thirteen major stocks traded at the Stockholm Stock Exchange. The stocks included are major in terms of both market capitalization (i.e. the total amount of equity outstanding, defined as stock price times the number of stocks issued) and turnover. Stocks have been selected based on data availability, but may be considered ideal for the purpose of this study since high liquidity may be necessary for reverting trends, as shown by Fehle and Zdorovtsov (2005). The stocks are (stock name tickers in parentheses) ABB (ABB), Ericsson B (ERICB), Stora Enso (STER), Svenska Handelsbanken A (SHBA), Swedish Match (SWMA), Astra Zeneca (AZN), Tele2 B (TEL2B), Electrolux (ELUXB), TeliaSonera (TLSN), Volvo B (VOLVB), Nordea (NDASEK), Autoliv (ALIV) and Hennes & Mauritz B (HMB).

Tick data is sampled at one-minute intervals from January 3 2003 through November 25 2004 during trading hours 09.00 to 17.30. The total number of trading days over the time period is 477, where the first half of the data (238 days) is used for estimating the model and tuning the parameters (called in-sample or estimation period), while the second half (239 days) is used for validation of the model (called out-of-sample or validation period). When returns for the individual stocks were calculated, data was resampled at five-minute intervals for dealing with the market microstructure issues mentioned in Section (2.5). Using five-minute sampling frequency is a standard approach for balancing between capturing the distribution and mitigating measurement errors, see for example Andersen et al. (2001), Bollerslev et al. (2005) and Fleming and Paye (2005).

The returns were calculated as

$$r_t = \log\left(\frac{p_t}{p_{t-5}}\right), \quad (18)$$

where r_t is the return at time t (with time being measured on a one-minute basis) and p_t is the price at time t .

Stock data was originally extracted from the STORQ database and SixTrust. Descriptive statistics for stock data is summarized in Table 1.

Stock ticker	Price, day 1	Change over in-sample (%)	Change over out-of-sample (%)	Change over full sample (%)
ABB	21.91	50.2	15.1	65.3
ERICB	6.75	56.7	64.6	121.3
STER	96.00	3.6	6.8	10.4
SHBA	122.50	15.1	14.7	29.8
SWMA	69.50	2.8	10.3	13.1
AZN	304.50	10.9	-24.4	-13.5
TEL2B	230.00	52.8	-41.3	11.5
ELUXB	143.00	9.3	-7.6	1.7
TLSN	33.30	8.4	12.9	21.3
VOLVB	152.00	37.2	21.7	58.9
NDASEK	40.50	21.2	26.2	47.3
ALIV	187.50	37.9	15.2	53.1
HMB	174.50	-1.1	20.9	19.7

Table 1: Descriptive statistics for stocks in the studied time periods, where prices are denoted in SEK.

The one week STIBOR, Stockholm InterBank Offer Rate, is used as the risk free rate of return and daily data was obtained from Reuters for the period of the study. The rate was at a mean level of 2.81%, going from 3.90% on January 3 2003 to 2.12% on November 25 2004.

3.1 Data Criticism

Data is based on one-minute sampling of tick-data which inevitably involves microstructure biases, as for example mentioned by Fleming and Paye (2005). Problems may arise when there are multiple quotes reported on the same stock during the same one-minute interval. Also, the system might lag, providing imprecise quotes. The data set was provided by Hossein Asgharian in a cleaned and filtered condition, which means that tedious data preparations were spared. The negative aspect is that insight was limited regarding the preparation process.

An obvious weakness of the data set is that it is very limited, both regarding the number of assets and the length of the time series. If more data had been available, statistical significance would increase. Also, the influence of the overall market trend would diminish as one would be able to study different regimes over the business cycle.

4 Methodology and Empirical Results

4.1 Identifying Extreme Events

The methodology described in Section 2.6 is used for identifying extreme events at a significance level of 5%⁴. A 30-day moving time frame is adopted for estimating the empirical distribution, similar to the length in Gustafsson and Hegerin (2005). Longer time frames are restricted by limited data availability and shorter time frames are considered insufficient for finding reasonable estimates.

The empirical distributions are highly dependent on the time frame considered, with daily returns being very different from those on an hourly basis. Due to intraday seasonalities, a distribution is estimated for each trading hour, i.e. one for 09.00-09.59, one for 10.00-10.59 and so on. The day is concluded by the distribution for 17.00-17.30, summing the number of hourly distributions to a total of nine. Each stock is also treated as an individual asset, which means that 13 estimations are concurrently conducted for each distribution.

Svensson (2006) finds evidence that large price changes in the same direction are often clustered within the same day, making it necessary to at least consider longer-term trends than an hour. Stretching the time horizon mitigates the risk of taking a contrarian position that is ruined by another large price change. When considering changes over as long horizons as a day, precision becomes coarse for finding events and one is also likely to detect returns that are part of a strong trend, not necessarily driven by an extreme, sudden change. Thus, a filter evaluates both daily and hourly returns, demanding returns at a specific time to be considered extreme on both scales. As a final step, the ability to detect jumps is utilized, filtering on events that are jump related within the last day. In this way, there are a total of 353 extreme events identified for the stocks during the estimation period⁵.

⁴This is a standard level of significance for statistical tests.

⁵464 events are identified when only filtering on the daily distribution and 369 trades are identified when filtering on both daily and hourly distributions. Thus, the jump detector only filters out 16 events after this.

4.2 Measuring the Reversal Effect

The behavior of the price process is studied for events identified by the approach in Section 4.1 during the estimation period. Just like Fehle and Zdorovtsov (2005), only great price declines are examined. When considering practical implementation issues, it makes sense to limit trading orders to long positions, since the process of shorting stocks is more complicated. Thus, a buy signal is generated after extreme losses and the return from holding the position is evaluated on different time horizons.

Evidence is presented by Fehle and Zdorovtsov (2005) that larger price declines should suggest stronger reversal effects, logically motivated by the market correcting more for a greater initial “mistake”. All events are treated as independent and since properties of trading psychology are assumed to be constant over time and between different stocks, the following cross-sectional regression is conducted

$$R_{i,T} = a_0 + a_1 P_{i,T} + \epsilon, \quad (19)$$

where realized return of trade number i , $R_{i,T}$, is related to the size of an initial price trend, $P_{i,T}$. In line with momentum studies like Jegadeesh and Titman (2001), the holding period, T days, is evaluated symmetrically with the formation period, i.e. T is the same for $R_{i,T}$ and $P_{i,T}$.

The results for different values of T are presented in Table 7 in Appendix. All regressions provide highly significant evidence that negative trends including extreme price declines are followed by a positive reversal effect. It is also indicated that more negative initial returns lead to stronger reversals (the a_1 coefficient is negative and significant), confirming the findings in Fehle and Zdorovtsov (2005).

For finding the most profitable holding period, mean returns for different periods are examined, presented in Figure 2. The ideal graph would present a relationship where the optimal holding period is located on a “bump”. However, the analysis does not seem to favor any specific holding or formation period, with return being a nearly linear function of holding time. Since there is a strong overall positive trend for the market benchmark, it may be hard to separate the reversal effect from the market trend. Instead,

the results from the regressions in Table 7, Appendix B, are examined for guidance, with the coefficient a_1 plotted in Figure 3.




Figure 2: The mean return for holding periods 1–40 days during the estimation period.




Figure 3: Absolute values of the reversion coefficient, a_1 , for different values of holding and formation periods, T .

It is decided to rely on $T = 5$ since the regression coefficients suggest the strongest

relationship⁶ and longer time periods being logically less related to the extreme event (other factors are likely to affect the stock price). The coefficients from this regression are presented in Table 2 and are also used for forecasting returns for out-of-sample data.

Coefficient	Estimate	95%-confidence interval
a_0	0.007	0.001 – 0.012
a_1	-0.355	-0.452 – -0.257

Table 2: Results from regressing initial price trend on realized return.

The 353 5-day-returns are plotted in a histogram, presented in Figure 4. The distribution of returns is obviously skewed in favor of positive returns.




Figure 4: Histogram of returns from holding period $T = 5$ for estimation data.

When examining all 353 extreme events, the mean return over the 5 day holding period is 1.35%. A t-test is applied, testing the one-sided null hypothesis whether mean return is zero with the alternative hypothesis of mean return being positive. The t-test rejects the null hypothesis at the 1%-level, indicating that the events form a reliable trading signal.

⁶A high a_1 coefficient suggests a stronger reversal effect. Also, the constant a_0 is low, which could indicate that the overall upward trend has been filtered out to a larger extent than for other values of T .

4.3 Examples of Extreme Events

Figure 5 gives an example of a successful trade. Clearly, a buy signal is triggered in a declining price environment characterized by large negative jumps. The return for the position over a 5 day holding period is over 13% in this case.




Figure 5: Extreme event for AZN, where the diamond indicates where a trading signal is issued.

Figure 6 gives an example of a failed trade. The trade is entered after a jump, but with jumps often being clustered over time, timing turns out to be wrong. This type of failed trades will occur occasionally, but are hard to avoid. One may consider strategies for exiting such failed trades before the loss grows too large.




Figure 6: Extreme event for ALIV, where the diamond indicates where a trading signal is issued.

4.4 Implementing a Portfolio Strategy

When forming a portfolio strategy, the Black-Litterman model is applied with views attained from the signals triggered by extreme price changes. The use of the Black-Litterman model assures that the money inserted in the strategy is fully invested at all times with a reasonable level of diversification.

4.4.1 The Benchmark and the Simple Form of the View Portfolio

For defining the benchmark, the first approach was to use a continuously updated value weighted index where the portfolio weights of the 13 stocks were in relation to the relative shares of the assets in the total market capitalization. Examining the index portfolio, it was observed that the weight of Ericsson B, the asset with the highest capitalization during the period of study, was extremely high, at times accounting for more than 25 % of the total portfolio. The large impact of Ericsson B could also be seen in its β for the benchmark, calculated over the entire period, reaching 2.5, whereas the β of all other stocks were less than one. In addition, market capitalization data was not available for the entire time span studied, which would reduce the already small data set by 77 days.

This led to the decision to use an equally weighted index as the benchmark. The benchmark will now be a strategy which invests equal amounts in each stock at the start of the period of study and stays that way through continuous readjustments. This also simplifies the process of updating the Black-Litterman model since the negative weights in the view portfolio, i.e. the weights of all assets except the one which is favored by the signal, will always be $(-)/12$. If a value weighted index would be used, the negative view portfolio weights would have to be balanced in accordance with their relative weights in the market portfolio. Another simplifying effect is that the neutral view portfolio will be homogenous with weights of $1/13$ for all assets.

Evaluating the equally weighted index, the cumulative return is plotted in Figure 7. The graph clearly shows a positive overall trend for the considered time period.




Figure 7: The cumulative return for the equally weighted index of the 13 stocks over the time period January 3, 2003, through November 25, 2004.

4.4.2 Estimating the Input Parameters

All input parameters can be estimated from historical values with a suitable timeframe. The choice of timeframe depends on the investment horizon or, in the case of a quantitative trading strategy, the holding period of the strategy. Since the strategy of this study is a short term strategy with a holding period of just 5 days, a rolling timeframe of 30 trading

days is used to estimate the parameters of the model. The rationale behind using such a short estimation period is that this will better reflect temporary trends and will not rely heavily on more consistent historical trends.

Even though the individual parameters are more or less well suited for short estimation periods, all parameters are estimated on the same time frame for consistency. Also, the 30 day time window is consistent with the estimation period used for deriving the empirical distributions in Section 4.1. The drawbacks of using such a short estimation period are that the parameters become very sensitive to noise and that long term structural relationships between certain stocks, e.g. stocks in the same business, might get lost in short term trends and noise. The parameters estimated over the 30 day time period are the covariance matrix, the β 's with respect to the benchmark and the expected benchmark returns.

4.4.3 Calculating View Premium from the Expected Returns

When estimating the model, it is crucial to filter the data in a way that only leaves trades that give a positive expected exceptional return, α . The extreme event related trading signal provides an expected return⁷ consisting of several components. α is then filtered out by subtracting the risk free rate and the stock's β times the expected benchmark return, as

$$\alpha_A = E(R_A) - R_f - \beta_A * \mu_B \quad (20)$$

β and the expected benchmark return are both continuously estimated over the previous 30 days. The risk free rate of return is always known with certainty as the one week STIBOR. After filtering out events with negative values of α , the remaining α 's will serve as the view premiums for assets on which there is a signal. For the estimation period, there are 239 trades entered for the 353 extreme events identified.

Note that (20) is different from the formula originally presented in (15), which differs by also containing a term with the exceptional benchmark return, $-\beta f_B$, on the right hand

⁷Forecasting is made using the regression model from Section 4.2.

side. Such a modification for short term exceptional trends is especially needed when the expected return of the benchmark is estimated as a very long term average, with an estimation period of several years. This study estimates expected benchmark return over such a short period of time that short term trends are assumed to already be captured by the μ_B measure, thus making it unnecessary to introduce a separate estimation procedure for f_B .

4.4.4 Estimating the View Confidence

It is necessary to assess view confidence before entering a position in order to update portfolio weights correctly. The view confidence is equal to the volatility of the view premium during the holding period. In the general examples of the Black-Litterman model seen in the literature, this is set by the individual investor based on the degree of certainty the investor feels about the view, which is as subjective as the view itself. When using the model with a quantitative strategy, it is desirable to find some way to determine view confidence as a function of the signal.

Since higher risk generally is compensated by higher return, it is reasonable to believe that one can relate the size of the volatility, i.e. the view confidence, to the size of the expected excess return. A simple linear regression is made of the 5 day volatility on the expected excess return from each trade⁸ during the first half of data as

$$\Omega_i = b_0 + b_1\alpha_i + \epsilon_i, \tag{21}$$

for event i .

The results from the regression can be seen in Table 3 and indicate that there is indeed a positive relationship between volatility and expected return that is statistically significant at the 5 % level. The same regression coefficients are further used when forecasting the view confidence during the validation period.

⁸Only the extreme events where expected excess return is positive, i.e. $\alpha > 0$, are considered trades.

Coefficient	Estimate	95%-confidence interval
b_0	0.016	0.015 – 0.018
b_1	0.211	0.146 – 0.276

Table 3: Results from regressing expected excess return (α) on view confidence (volatility).

4.4.5 Setting the Risk Aversion Coefficient

The risk aversion coefficient is used to tune the portfolio weights in accordance with the individual investor’s willingness to take on risk. In practice, this is a quite complex input which is often arbitrarily set. In this study, the risk aversion coefficient is varied at three different levels, low (15), medium (25) and high (35). The coefficient is set at levels that do not generate wild portfolio weights that involve shorting or using leverage for financing the portfolio.

The Black-Litterman model was originally developed for assets such as global indexes of equities, bonds and currencies, all of which have higher covariances than individual stocks. The stocks included in the study all have low covariance compared to the empirical covariances of indexed assets. This calls for the use of higher risk aversion coefficients than what is seen in e.g. Litterman and He (1999), who use the value of 2.5 as the approximated world average risk aversion.

4.5 Results from Portfolio Strategies

Table 4 presents the returns for the portfolios formed using Black-Litterman methodology with view inputs formed from trading signals triggered by the extreme events. The active strategy clearly outperforms the index over the estimation period (11.9% higher with medium risk aversion), as well as in out-of-sample (5.1% higher with medium risk aversion). One may draw the same conclusion from risk-adjusted returns, represented by Sharpe ratios. The returns vary by the three levels of risk aversion examined, where higher risk aversion takes less active positions and thus shrinks the active return towards the benchmark. No strategy has taken trading costs into account, which should reduce the

return of both active and indexed returns (since the equally weighted index is rebalanced over time).

A	r_A/SR in sample	r_I/SR in sample	r_A/SR out of sample	r_I/SR out of sample
15	50.8% /2.86	30.3% /1.77	19.1% /1.57	10.4% /0.74
25	42.2% /2.61	30.3% /1.77	15.5% /1.30	10.4% /0.74
35	38.5% /2.45	30.3% /1.77	13.9% /1.17	10.4% /0.74

Table 4: The total returns of the active trading strategy (r_A) for different risk aversion coefficients (A) compared to the index total returns (r_I). Sharpe ratios, SR , are presented for each return.

The positions for all stocks are plotted over time in the same graph, where the case with medium risk aversion, $A = 25$, is presented in Figure 8 for in-sample and Figure 9 for out-of-sample. For other risk aversion coefficients, graphs are presented in Appendix. Clearly, the portfolios are frequently rebalanced with certain periods of especially high activity. Also, the resulting positions are at some points weighing heavily in single assets, for example up to 35% in one stock during the estimation period (see Figure 8).




Figure 8: Positions for all stocks using $A = 25$ for estimation data.




Figure 9: Positions for all stocks using $A = 25$ for out-of-sample data.

4.6 Implementing Trading with Fixed Bets

According to Grinold and Kahn (2000), p. 377, a careful implementation procedure is necessary for not spoiling investment research. Therefore, an alternative implementation approach is evaluated, where fixed trading bets are made on each signal triggered by the extreme events. The strategy is simpler than the Black-Litterman methodology, while also being less sensitive to assumptions about input parameters.

A trading account with 1 SEK is set up and 0.05 SEK⁹ is invested in each trade. At times, the account may be more than fully invested, i.e. a minor degree of leverage may be used, but most of the time, the account is held in cash (assumed to generate no return). Trades are entered only if the expected excess return (α) is positive, where forecasts of expected returns are formed using the results from Section 4.1 and extracted as in 4.4.3. For testing whether higher expected return improves accuracy, trades are also filtered on 1% and 2% respectively.

Table 5 present the results from the estimation period and indicate that hit ratio is high for finding profitable trades; around 70% of trades generate positive return. Sharpe ratios

⁹This level is arbitrarily set, where 5% of the capital per bet provides decent diversification and risk control.

α	No of trades	Hit ratio (%)	Mean return (%)	Sharpe ratio
> 0%	239	69.5	0.12	3.88
> 1%	141	69.5	0.14	3.23
> 2%	76	73.7	0.23	3.24

Table 5: Trading with fixed bets at different expected excess returns, α , for the estimation period.

on an annual basis are also high, taking values over three. This can be compared to other investment strategies like buying and holding the S&P 500 stock index, which generated a Sharpe ratio of 0.47 over the time period 1995–2005 and 0.16 during 2000–2005. The equally weighted benchmark should be the most relevant benchmark, where active trading clearly beats the index Sharpe ratio of 1.77.

Table 5 seems to support the hypothesis that higher forecasted returns generate more accurate trading - the hit ratio increases slightly while the return per trade clearly improves. Mean returns of trades are highly significant, with t-tests rejecting zero mean returns at the 1%-level in all cases.

Figure 10 presents an overview of the trading for trades filtered at $\alpha > 2\%$. Panel 1 and 2 compares trading returns over time, while Panel 3 presents how the total degree of investment varies over time. From Panel 2, trading is consistently generating profits, with no time periods of highly negative returns. Panel 3 indicates that the strategy is never fully invested, implying opportunities for increasing total return by applying a complementary investment strategy.

Out-of-sample trading is also examined, producing almost as impressive returns as for the estimation period. Table 6 presents the results, where stronger proof is provided that higher forecasted returns may improve trading accuracy. Both hit ratios and mean returns become significantly higher when filtering. All Sharpe ratios are higher than the equally weighted benchmark (taking the value of 0.74 over the time period).

Figure 11 presents an overview of out-of-sample trading for trades filtered at $\alpha > 2\%$. From Panel 2, trading returns experience boosts in volatility when active trading intensi-




Figure 10: Panel 1 presents return for the equally weighted index over the estimation period. Panel 2 presents return for the active strategy, filtering trades on $\alpha > 2\%$. Panel 3 presents the total degree of investment.

fies. Just as for estimation data, Panel 3 indicates that the strategy is never fully invested. It should be noted that the scale for the “Index return”-axis in Panel 1,11, is different from Panel 1, 10.

α	No of trades	Hit ratio (%)	Mean return (%)	Sharpe ratio
$> 0\%$	250	60.0	0.024	0.75
$> 1\%$	183	63.9	0.040	1.33
$> 2\%$	99	67.7	0.061	1.38

Table 6: Trading with fixed bets at different expected excess returns, α , for the validation period.




Figure 11: Panel 1 presents return for the equally weighted index for out-of-sample. Panel 2 presents return for the active strategy, filtering trades on $\alpha > 2\%$. Panel 3 presents the total degree of investment.

The distribution of returns is examined for out-of-sample data and presented in Figure 12. The histogram suggests that there is a fairly large amount of negative returns, which may

cause the strategy to perform worse than during the estimation period. One could refine trading by introducing some filter for maximum loss tolerated, a so-called “stop-loss” (see for example Gustafsson and Hegerin (2005) for an implementation of such trading rules).




Figure 12: Histogram of returns from holding period $T = 5$ for validation data, filtered on $\alpha > 0\%$.

4.7 Details on Trades for Individual Stocks

The trades triggered, i.e. extreme events with $\alpha > 0\%$, are examined for each individual stock. Table 8 and 9 in Appendix provide summary statistics for the estimation and validation period, respectively. It is interesting to note that trading is not heavily focused on just a few stocks, but rather evenly distributed over the full sample of available assets. Also, the overall tendency of mean returns, being positive in almost all cases, further strengthens faith in the accuracy of the trading signal.

Maximum returns for some stocks are occasionally extreme (e.g. 40% for ERICB), but mean returns for the full sample are in the same vicinity, (around 0.5%–5% for in-sample and around 1% for out-of-sample). The active strategy manages to generate positive return for some stocks with a long-term losing trend. TEL2B and ELUXB for instance, both experience a negative trend over the validation period (see Table 1 for details), while it is still possible to make profits from trading on extreme events. Also, there seems to

be additional evidence that the strategy performs systematically poorer for the validation period, i.e. the reverting effect becomes weaker. It does not seem to be the case that there are some extreme failed trades that ruin the total return (there are no trades losing more than 10%).

5 Conclusion

For the data set studied, there is a significant reversal effect after extreme price changes. The reverting trend is stronger for larger initial price declines, supporting logical intuition and confirming the empirical findings by other studies. This study also relates the effect to longer term trends, stretching the time period to days instead of hours.

Trading is implemented on both a portfolio and fixed bet basis. The portfolio approach looks nice in theory, but contains a number of abstract input parameters that are hard to estimate in practice. Fixed bets is a more straightforward approach for assessing the overall opportunities for making profit using the trading signal. However, both evaluation methods indicate similar results.

Trading during the estimation period suggests highly significant returns from taking active positions. Out-of-sample also generates very promising results, but slightly less extraordinary than during in-sample. A plausible explanation is that there is a strong upward market trend for the period used for estimation, making a significant impact on some of the estimated parameters. Reversals may be stronger during such positive market regimes. However, during both time periods, adding active bets results in a strategy that beats the benchmark in terms of risk-adjusted returns.

One promising finding is the fact that the regression model for forecasting excess expected return (α) for each trade performs well in both in- and out-of-sample. Filtering on higher levels of α turns trading more accurate in terms of both hit ratio for finding profitable trades and mean returns.

5.1 Suggestions for Future Studies

The study may be further developed by introducing more advanced tools for forecasting and estimation. For example, volatility measures are forecasted using a simple linear regression, with constant coefficients for all events. It may be reasonable to introduce some kind of autoregressive model for improving accuracy of forecasts (see for example Madsen et al. (2004), Chapter 5, for ARCH, GARCH and other ways of modeling volatility).

Also, when making regressions, one may assume that the properties of the trading signal actually varies over time. Considering the results from the study, it is likely to believe that opportunities for profitable trading on reverting trends changes with market conditions. This could be incorporated in the model by continuously evaluating a regression from a recent window of time. For more sophisticated methods, one may turn to programming neural networks, using artificial intelligence to adapt to new states of the market (see for example Campbell et al. (1997), Chapter 12, for more on the subject).

Since failed positions may significantly affect the overall performance, trading may be implemented by introducing trading rules for exiting such trades. For example, a maximum level of loss per trade may be set, introducing a “stop-loss” restriction. Also, one may determine a positive level for closing positions after high profits, hypothesizing that the reverting trend is over. The result would be a more dynamic trading strategy in terms of holding period, while also adding additional parameters that one runs the risk of over-optimizing.

For estimating covariance of assets, β 's and other parameters, a short term estimation period is used. Just as for the volatility forecasts, more advanced methods may be applied. For example, one may include the effect of longer term averages that reflect structural relationships such as stocks operating in the same industry (a rather rigorous approach favored by practitioners such as Grinold and Kahn (2000)).

The constraint of no shorting may be relaxed to some extent when implementing portfolios. Since the default portfolio holds all assets, one may utilize negative views by selling from the initial positions. This means that one could examine extreme positive price changes with the hypothesis that there will be a negative reversion, in analogy with the approach of this study. It is however important to note that implementation may lead to zero weights, generating the problematic corner solutions as discussed in Section 2.7.5.

For more realistic implementation, trading costs should be taken into account. The strategies trade very frequently and it may be necessary to define some thresholds for determining whether it is worth updating the portfolio or not. This would introduce inertia in the trading system, where the degree of inertia is set by the trade-off between expected return for a trade and the related trading cost.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001). The distribution of realized exchange rate volatility, *Journal of the American Statistical Association*, **96**.
- Barndorff-Nielsen, O. and Shephard, N. (2003). Econometrics of testing for jumps in financial economics using bipower variation, *Nuffield College, Oxford University Working Paper*.
- Barndorff-Nielsen, O. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics*, **2**, 1–48.
- Björk, T. (2004). *Arbitrage Theory in Continuous Time*. Oxford University Press, Oxford.
- Black, F. and Litterman, R. (1990). Asset allocation: Combining investors views with market equilibrium, *Fixed Income Research, Goldman Sachs*, **September**.
- Bollerslev, T., Kretschmer, U., Pigorsch, C., and Tauchen, G. (2005). A discrete-time model for daily returns and realized variations: Jumps and leverage effects, *Duke University Working Paper*.
- Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press, Princeton NJ.
- Chiarella, C. and He, T. (2002). An adaptive model on asset pricing and wealth dynamics with heterogeneous trading strategies, *Computing in Economics and Finance*, **135**.
- Cumby, R. E. and Glen, J. D. (1990). Evaluating the performance of international mutual funds, *The Journal of Finance*, **45**, 497–521.
- Elton, E. J., Gruber, M. J., and Blake, C. R. (1996). The persistence of risk-adjusted mutual fund performance, *The Journal of Business*, **69**, 133–157.
- Fama, E. (1970). Efficient capital markets: a review of theory and empirical work, *Journal of Finance*, **2**, 383–417.
- Fehle, F. R. and Zdorovtsov, V. M. (2005). Large price declines, news, liquidity, and trading strategies: An intraday analysis, *University of South Carolina Working Paper*.

- Fleming, J. and Paye, B. (2005). High-frequency returns, jumps and the mixture of normals hypothesis, *Rice University Working Paper*.
- Grinold, R. C. and Kahn, R. N. (2000). *Active Portfolio Management: a Quantitative Approach for Providing Superior Returns and Controlling Risk*. McGraw-Hill, New York, NY.
- Gustafsson, J. and Hegerin, C.-J. (2005). *Systematic FX trading*. Master's Thesis, Centre for Mathematical Sciences, Lund.
- Huang, X. and Tauchen, G. (2005). The relative contribution of jumps to total price variance, *Journal of Financial Econometrics*, **3**, 456–499.
- Idzorek, T. M. (2004). A step-by-step guide to the black-litterman model, incorporating user-specified confidence levels, *Zephyr Associates Working Paper*.
- Jegadeesh, N. and Titman, S. (2001). Momentum, *University of Illinois Working Paper*.
- Litterman, R. and He, G. (December 1999). The intuition behind black-litterman model portfolios, *Investment Management Research, Goldman Sachs Quantitative Resources Group*.
- Madsen, H., Nygaard, J., Lindstrom, E., Baadsgaard, M., and Holst, J. (2004). *Statistics in Finance*. Centre for Mathematical Sciences, Lund.
- Malkiel, B. G. (1995). Returns from investing in equity mutual funds 1971 to 1991, *The Journal of Finance*, **50**, 549–572.
- Markowitz, H. M. (1952). Portfolio selection, *Journal of Finance*, **March**, 77–91.
- Mossin, J. (1966). Equilibrium in a capital asset market, *Econometrica*, **34**, 768–783.
- Qi, M. and Wu, Y. (2005). Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange markets, *Journal of Money, Credit and Banking*, *forthcoming*.
- Sharpe, W. F. (1966). Mutual fund performance, *Journal of Business*, **January**.
- Svensson, L. (2006). *Contrarian Trading Strategies on Extreme Price Changes in the Foreign Exchange Market*. Master's Thesis, Centre for Mathematical Sciences, Lund.

Tauchen, G. and Zhou, H. (2005). Identifying realized jumps on financial markets, *Duke University Working Paper*.

Zawadowski, A. G., Andor, G., and Kertesz, J. (2004). Large price changes on small scales, *Physica A*, **344**, 221–226.

A Empirical Results Appendix

A.1 Regressions of Expected Return

Coefficient/ T	Estimate	95% -confidence interval
$a_0/T = 1$	-0.002	-0.008 – 0.004
$a_1/T = 1$	-0.209	-0.365 – -0.053
$a_0/T = 2$	0.005	0.000 – 0.010
$a_1/T = 2$	-0.128	-0.243 – -0.013
$a_0/T = 3$	0.006	0.001 – 0.010
$a_1/T = 3$	-0.196	-0.300 – -0.092
$a_0/T = 4$	0.007	0.002 – 0.012
$a_1/T = 4$	-0.244	-0.347 – -0.141
$a_0/T = 5$	0.007	0.001 – 0.012
$a_1/T = 5$	-0.355	-0.452 – -0.257
$a_0/T = 6$	0.008	0.003 – 0.014
$a_1/T = 6$	-0.301	-0.398 – -0.204
$a_0/T = 7$	0.013	0.007 – 0.019
$a_1/T = 7$	-0.224	-0.316 – -0.131
$a_0/T = 10$	0.020	0.014 – 0.026
$a_1/T = 10$	-0.125	-0.216 – -0.035
$a_0/T = 15$	0.030	0.023 – 0.038
$a_1/T = 15$	-0.179	-0.271 – -0.086
$a_0/T = 20$	0.044	0.035 – 0.052
$a_1/T = 20$	-0.131	-0.226 – -0.035
$a_0/T = 25$	0.053	0.043 – 0.063
$a_1/T = 25$	-0.112	-0.210 – -0.013
$a_0/T = 30$	0.061	0.051 – 0.072
$a_1/T = 30$	-0.148	-0.250 – -0.046

Table 7: Results from regressing initial price trend on expected return for different formation and holding periods, T .

A.2 Evaluating Trades for Individual Stock

Stock ticker	No of trades	Mean return (%)	Volatility of returns (%)	Min return (%)	Max return (%)
ABB	16	5.3	7.3	-3.4	20.7
ERICB	14	5.5	11.6	-5.0	40.7
STER	19	1.5	4.9	-6.1	14.9
SHBA	21	2.8	2.1	-1.2	8.0
SWMA	29	0.68	2.9	-4.1	6.0
AZN	15	2.8	5.9	-6.1	14.2
TEL2B	15	4.0	2.3	0.0	8.4
ELUXB	15	1.4	5.1	-6.0	12.4
TLSN	14	0.6	3.0	-3.2	6.7
VOLVB	20	3.5	4.6	-5.3	13.4
NDASEK	16	2.5	3.8	-5.2	12.9
ALIV	21	1.3	3.4	-7.0	7.5
HMB	24	0.26	2.7	-5.8	4.4

Table 8: Evaluating trading returns for individual stocks during the estimation period, filtering trades on $\alpha > 0\%$

Stock ticker	No of trades	Mean return (%)	Volatility of returns (%)	Min return (%)	Max return (%)
ABB	18	-0.1	4.5	-9.2	6.2
ERICB	24	1.2	4.9	-8.0	11.3
STER	12	1.1	1.9	-2.1	4.3
SHBA	11	0.5	1.6	-2.5	2.8
SWMA	19	0.6	2.8	-5.6	5.8
AZN	22	-1.0	3.1	-8.8	4.5
TEL2B	26	0.2	3.0	-5.8	5.6
ELUXB	22	0.4	2.4	-6.2	3.8
TLSN	25	0.3	3.1	-7.6	5.2
VOLVB	18	1.1	3.5	-5.1	8.7
NDASEK	17	1.7	2.2	-1.9	6.6
ALIV	17	1.0	2.7	-3.7	5.8
HMB	19	-0.1	2.1	-4.1	3.9

Table 9: Evaluating trading returns for individual stocks during the validation period, filtering trades on $\alpha > 0\%$

A.3 Positions from Portfolio Implementation




Figure 13: Positions for all stocks using $A = 15$ for estimation data.




Figure 14: Positions for all stocks using $A = 35$ for estimation data.




Figure 15: Positions for all stocks using $A = 15$ for out-of-sample data.




Figure 16: Positions for all stocks using $A = 35$ for out-of-sample data.

B Technical Appendix

The methodology of Tauchen and Zhou (2005) for detecting jumps is summarized by Svensson (2006). First, a ratio statistic is defined as

$$RJ_t \equiv \frac{RV_t - BV_t}{RV_t}. \quad (22)$$

$100 * RJ_t$ may be intuitively interpreted as the percentage of total price variation that is related to jumps. Huang and Tauchen (2005) conduct extensive Monte Carlo simulation studies and find nice properties of RJ_t scaled by its asymptotic variance

$$ZJ_t \equiv \frac{RJ_t}{\sqrt{[(\frac{\pi}{2})^2 + \pi - 5] \frac{1}{m} \max(1, \frac{TP_t}{BV_t^2})}}, \quad (23)$$

where ZJ_t is asymptotically distributed as the standard normal. Barndorff-Nielsen and Shephard (2004) introduce TP_t , the tri-power quarticity robust to jumps, and show that this converges as

$$TP_t \equiv m\mu_{4/3}^{-3} \frac{m}{m-2} \sum_{j=3}^m |r_{t,j-2}|^{4/3} |r_{t,j-1}|^{4/3} |r_t|^{4/3} \rightarrow \int_{t-1}^t \sigma^4(s) ds, \quad (24)$$

where $\mu_{4/3} \equiv 2^{2/3} \Gamma(7/6) \Gamma(1/2)^{-1}$.