Executive Summary

Problem and Objective

The aim of this thesis is to extend the existing literature on technical analysis by incorporating additional indicators and markets. Technical analysis is discussed controversially in the investment world and hence, it is interesting to link the ideas of charting to a more accepted concept: trading the trend. Consequently, the indicators included are of trend-following nature. Moving average cross-over rules based on different averaging methods are used for benchmarking, and more complex, forward-looking or adaptive indicators are introduced. Since the early literature on technical analysis focuses on US blue chips, additional indices are included.[1]

The focus is set on learning and voting algorithms according to Hsu and Kuan (2005), which autonomously generate trading signals based on a universe of underlying rules.[2] The corresponding signal quality is statistically assessed. Furthermore, the generally assumed switching strategy is to be overcome and the concept of trading signals is applied to the more complex mean-variance investment approach.

Methods

For the assessment of the signal quality a simple switching strategy is assumed, where the position is either long, neutral or short (i.e. 1, 0 or -1) depending on the trading signal. A strategy is significantly outperforming if its Sharpe ratio is significantly higher than the one of the market. The corresponding test is based on Ledoit and Wolf (2008), and uses circular bootstrapping of the studentized difference in Sharpe ratios. However, as numerous strategies are applied to a unique time series of observations, the dangers of data snooping are immense (Brock et al., 1992). Hence, the hypothesis tests are adjusted for data snooping using the methodology of Romano et al. (2008).

[1] The indices are furthermore split into two groups. Group A consists of the S&P 500, FTSE 100, and Nikkei 225 index. Group B includes the NASDAQ Composite, MSCI Emerging Markets, and the Russell 2000 index.

[2] In a periodic re-evaluation process, the learning algorithm identifies the rule with the historically highest average log-return after transaction costs to follow going forward. The voting algorithm is a generalization thereof where not only one but at least two rules are considered simultaneously.

In a simple mean-variance framework with one risky asset, expectations are based on past observations conditional on the trading signal. The prediction power is evaluated using the difference in cumulative squared forecast errors, where the benchmark model has unconditional historical moments as expectations. The value of the signals is represented with the realized utility for a mean-variance investor, making the performances comparable.

Results

Following the significance test of Ledoit and Wolf (2008) and assuming no transaction costs, significantly outperforming switching strategies are found in every market. As expected, the significance weakens once transaction costs are introduced. However, as soon as data snooping is accounted for, only the MSCI Emerging Markets and the Russell 2000 index allow for significant outperformance.

When only the index can be traded, outperforming with a switching strategy is hard since it is only possible when the index is decreasing. For all indices in Group A, the difference of the cumulative performance to the market is mainly achieved in bearish times. In bullish periods, the best performing strategies fail to outperform and thus to predict the daily market movements. For the MSCI Emerging Markets and the Russell 2000 index, the best performing strategies are better than the market even in bullish periods. Hence, the market timing implied by the strategies seems to work. An exceeding performance can directly be linked to significant autocorrelation coefficients. The higher the autocorrelation, the better the strategies. This is irrespective of the size of the lag considered or the corresponding sign, since the algorithms are automatically choosing an adequate strategy to exploit this information.

In the mean-variance framework, the conditional average return given a buy signal in general is higher than the conditional average return given a sell signal. Using this conditional setup to form expectations, a higher prediction power is achieved. As long as no transaction costs are included, this directly leads to a higher realized utility. With transaction costs, exceeding the benchmark utility is more likely if short positions are allowed.