Advances in the prediction of stock market volatility

Dissertation
der Wirtschaftswissenschaftlichen Fakultät
der Universität Augsburg
zur Erlangung des akademischen Grades eines Doktors
der Wirtschaftswissenschaften
(Dr. rer. pol.)

vorgelegt von

Herr Dipl.-Math. oec. Alain Hamid

Erstgutachter: Prof. Dr. Yarema Okhrin
Zweitgutachter: Prof. Dr. Dr. h.c. Günter Bamberg
Vorsitzender der mündlichen Prüfung: Prof. Dr. Marco Wilkens

Tag der mündlichen Prüfung: 07.05.2015

Contents

Chapter 1

Introduction

Beginning with the appearance of tradable financial assets, such as stocks, options, futures, bonds etc., there has been a rich interest in the future development of the corresponding asset returns. It is a tedious and ongoing debate among academics whether stock returns are predictable or not. While many empirical studies claim (at least to some extent) the predictability of returns [see Barberis (2000); Lettau and Ludvigson (2001); Campbell and Thompson (2008); Cochrane (2008, 2011); Ferreira and Santa-Clara (2011) and the references therein] others remain skeptical [see Ang and Bekaert (2007); Welch and Goyal (2008); Cenesizoglu and Timmermann (2012) and the references therein]. The *Efficient Market Hypothesis* (Fama, 1965) and the Random Walk Hypothesis, by Malkiel (1973, 2011), examine return predictability from a theoretical point of view. Thereby the Efficient Market Hypothesis states that all information currently available to investors is already incorporated within the asset price, and the Random Walk Hypothesis states that returns follow a random walk and are consequently unpredictable. Since the Random Walk Hypothesis is consistent with the Efficient Market Hypothesis, return predictability has often been viewed as a violation of market efficiency. However more recent research states that return predictability can occur due to a time-varying aggregate risk exposure and in this case, still be consistent with efficient markets [see Rapach and Zhou (2013) for a discussion of this topic].

Though the issue of stock return predictability has been neither empirically nor theoretically solved, it is widely accepted in academia that a large unpredictable component in stock returns makes it in general difficult to forecast returns with statistical significance[1] [see Merton (1980); Rapach and Zhou (2013)].

On the other hand, the variation of asset returns, the so called *volatility*, which composes the centrepiece of this dissertation, is thoroughly predictable (Andersen and Bollerslev, 1998). This quantity is also of great interest, as it forms an important risk measure for financial assets and has consequently a vast application area in portfolio optimization or risk management in general. In particular the recent subprime and Eurozone crisis pointed out the necessity for proper risk management in the financial sector. In this respect it is not surprising, that modelling and predicting return volatility has become an active field of research in the last few decades. In this context, empirical investigations have uncovered basic properties in the dynamics of volatility time series - the so called *stylized facts*. Examples for these stylized facts are long memory (the autocorrelation is high and decays slowly), volatility clustering (phases of high volatility are followed by tranquil phases and vice versa), mean reversion (the time series tends towards its mean after large deviations) or the asymmetric volatility phenomenon also known as Leverage effect (the impact of negative returns on volatility is larger than that of positive returns). After volatility had initially been treated as constant over time [as in the Black Sholes model for option pricing, Black and Scholes (1973)], academics began to develop models that succeeded in mimicking at least some of the stylized facts of volatility series.

The aim of this dissertation is to shed light on different aspects of the task of forecasting volatility in financial markets. Therefore chapter 2 investigates in how far the concept of empirical similarity, developed in the context of case based decision theory, can be exploited to combine different volatility components or forecasts to augment forecasting accuracy. Chapter 3 delivers a substantial empirical study, examining the forecasting performance of different high-frequency measures of daily volatility, applied with different models and loss functions. Within chapter 4 the role of investor attention on volatility is considered. Chapter 5 concludes the dissertation and gives a perspective on

[1] In this context it should be mentioned, that a growing number of articles argues that return forecasts can still be of economic value, although they fail to beat naive benchmarks with statistical significance [see Han (2010); Cenesizoglu and Timmermann (2012)].

related future research topics. The remainder of this introduction describes the basic ideas and findings of the articles embedded in this dissertation.

1.1 The application of the empirical similarity concept in time series forecasting

Since accurate forecasts are in general of crucial interest to decision makers it is not surprising that much effort has gone into the development of forecasting models and model selection tools. Thereby it is a common finding that predictions (for a certain variable of interest) obtained by the combination of different forecasting models frequently outperform individual forecasts [see Aiolfi and Timmermann (2006) and the references therein for an introduction into the topic of model combination]. Though this fact has been well known for many decades (Bates and Granger, 1969), the issue of how forecasts should be combined, more precisely which weights should be assigned to each individual forecast, is still a subject of active research and has not been solved conclusively to this day. Usually the approaches for model combination involve the model success probabilities (Elliott and Timmermann, 2004), which can be seen as the probability of occurrence for the next period's state. Such a probabilistic procedure, which is in line with the expected utility theory of von Neumann-Morgenstern, assumes that decision makers are able to evaluate probabilities. However this is not always guaranteed for real world applications. The first article of this dissertation contributes to this research direction and proposes a new approach for determining the weights for model combination, in situations where decision makers are unable (or unwilling) to evaluate probabilities. The core idea is to involve the concept of *empirical similarity* [see Gilboa et al. (2006, 2011)] to calculate the model weights. Originally the empirical similarity (ES) concept was intended to forecast a variable of interest Y_t by means of a database of cases (problems, situations), where each case consists of past values of m relevant variables (states of nature) $X_i^1, X_i^2, \ldots, X_i^m$ and the respective outcome Y_i for $i = 1, \ldots, t-1$. Thereby the similarity between the current and past cases is involved in a nonlinear way in the forecasting process. The ES model is related to numerous statistic and econometric standard techniques like nonlinear regression, kernel estimation, Bayesian updating, interpolation and autoregressive modelling. As pointed out in Gilboa et al. (2011) the similarity between the ES approach and nonlinear regression by means of the widely used Nadaraya-Watson estimator is remarkably strong, according to the applied formulas. However they do also note the conceptual difference between both models, which lies in the role played by the true underlying data generating process for the estimation procedure. Furthermore Gilboa et al. (2011) illustrate the connection to other statistical techniques like spatial modelling, probability estimation and double kernel density estimation. In our application of the ES approach we measure the empirical similarity distance between the current observation and the last forecasts from different models. A model which recently provided more accurate point forecasts (with higher similarity to the true value) obtains a larger weight compared to inferior models. In general this approach for model combination can be applied to any real-valued time series, but in this article it is explicitly used to combine volatility forecasts or past realized volatilities at different frequencies. The ES volatility forecasts for the S&P 500 index, NASDAQ and Nikkei are compared to the HAR model of Corsi (2009), the equally weighted model (1/3) and the RiskMetrics model (RM). While the HAR model is nothing other than a linear regression of contemporary daily volatility on past average daily (short-run), weekly (medium-run) and monthly (long-run) volatility, the 1/3 model uses the same components with equal weights. In particular two ES models are suggested: First the ES1 approach which combines the HAR components, second the ES2 approach consisting of past daily realized volatility (short-run), RiskMetrics volatility forecasts (medium-run), and the HAR volatility forecasts (long-run). The sample period ranges from 3 January 2000 to 25 February 2013, containing tranquil and turbulent volatility phases. While the ES2 model produces decent daily volatility forecasts, the ES1 approach constantly provides better daily and weekly volatility forecasts than the HAR model in-sample and out-of-sample. The results are significant during and after volatility shocks according to the model confidence set of Hansen et al. (2011). Furthermore the empirical analysis shows that the equal weighted 1/3 model seems to be accurate in calm market phases, where volatility components over different horizons contribute evenly to the total volatility, whereas the ES weights strongly deviate from the 1/3-benchmark in impermanent market phases, explaining the superior performance in this period.

1.2 On the determinants influencing the accuracy of volatility predictions

While the article described in the previous section illustrates one possibility to improve volatility forecasts, there is still a considerable number of factors that influence the accuracy of volatility forecasts. Among these factors are the chosen data, the volatility estimates applied to measure the unobservable true volatility, the models used to compute forecasts and the loss functions to evaluate these forecasts. The second article of this dissertation deals with this issue by means of an extensive empirical study.

Due to the availability of intraday (high frequency) financial data a vast number of measures for daily volatility have been proposed. Typically the application of recently developed measures is motivated by improved theoretical properties as opposed to the standard realized variance (RV) measure. For example more sophisticated measures (theoretically) show less variance, are less biased or robust to jumps or market microstructure noise[2] [see for example Zhang et al. (2005); Zhang (2006) or Barndorff-Nielsen and Shephard (2004); Barndorff-Nielsen et al. (2008)]. Though several measures are suggested from a theoretical point of view, empirical studies fail to identify a volatility measure with general superior performance. There is rather a large number of empirical articles, with contradicting recommendations about the preferable measure [see Ghysels et al. (2006), Patton and Sheppard (2009) and Liu et al. (2012) among others]. We investigate the prediction power of 7 different high-frequency based volatility measures, applied to 3 time series models and evaluated with 4 loss distances of the family of robust loss functions from Patton (2011). We apply the (subsampled) realized variance, bipower variation and realized kernels with different sampling frequencies as estimates for the daily volatility. The usage of these estimates is widespread in the literature and they provide different appealing theoretical properties, like robustness to noise or jumps in the price process. The HAR- (Corsi, 2009), MIDAS- (Ghysels et al., 2006) and ES-model

[2] Microstructure noise is an umbrella term for market frictions like discreteness of prices, bid-ask bounces or unevenly spaced observations that lead to noisy volatility estimates.

(Golosnoy et al., 2014) are used to predict next-periods volatility since they all proved to be successful in volatility forecasting. Given that these models are of (nonlinear) regression type, they fit within our framework, in which one specific volatility proxy (e.g. the RV) is predicted by means of another volatility measure (e.g. bipower variation). Tests on forecasting superiority are conducted via the model confidence set (MCS) as introduced in Hansen et al. (2011). The MCS is suitable for arbitrary loss functions and hence appropriate for our framework. The applied loss functions are robust in the sense that the ranking of different volatility forecasts is the same whether the ranking is done using the true variance, or some unbiased, but possibly imperfect (noisy), volatility proxy. This property is important, since forecasts cannot be compared to the true volatility, which is an unobservable quantity. The data set consists of 18 worldwide stock indices from 01/03/2000 to 02/25/2013. The out-of-sample forecasts are derived within a rolling window approach and average results over all indices are calculated. The main finding is that different models tend to favour different volatility measures, hence the existence of a generally superior measure can be doubted heavily. Since this result holds for a large data set, is statistically significant and proves to be robust to variations of the forecasting horizon and the volatility proxy[3], the contradictory findings in prior empirical literature can be explained by the usage of different volatility models and are obviously not only data driven. Further results are that the choice of the loss function is of high importance, especially when statistically significant results are desired. In addition the assumption that more sophisticated volatility estimates, for example realized kernels, lead to better forecasts, does not hold in general. To the contrary, the standard RV measure performs remarkably well when combined with an appropriate model.

1.3 The role of investor attention on the volatility process

Until now we have examined the quality of volatility forecasts without incorporating external factors like macroeconomic variables or news. While the

[3]Since the true volatility is unobservable, a proxy has to be used instead. It is common to use RV as a proxy, though there is no theoretical justification for this choice.

link between macroeconomic factors and stock market volatility has already been a subject of active research for decades [see Schwert (1989) or Veronesi (1999) among others], the third article of this dissertation focuses on the role of investor attention on the volatility process. Thereby we use Google Trends data as a proxy for investor attention, since Google accounted for 77.46%[4] of search queries worldwide in 2013, and can consequently be assumed to be representative of general internet search behavior.

The inclusion of "online data" such as internet message postings (Kim and Kim, 2014), Facebook users sentiment data (Siganos et al., 2014) or search frequencies [Vlastakis and Markellos (2012), Andrei and Hasler (2013) or Vozlyublennaia (2014)] to model financial time series is a rather new but growing field. Though Vlastakis and Markellos (2012) and Vozlyublennaia (2014) suggest that Google search volume is a driver of future volatility, they focus on in-sample analysis and do not provide any model which successfully exploits search data to improve volatility forecasts. The last article of this dissertation contributes to this research direction and shows how Google search data can be used to improve out-of-sample volatility forecasts for the DJ index with statistical significance. In this context a special type of empirical similarity model, which can be interpreted as a dynamic AR(1) model with a time varying autoregressive coefficient, is applied. In our framework, this model, suggested by Lieberman (2012) and afterwards called the ESL model, incorporates the similarity between last periods Google data and volatility to determine the dynamic autoregressive coefficient. It is noteworthy that the ESL model exploits the information in Google data in a nonlinear way and is able to replicate time series with stationary, unit-root or explosive phases. The capability of the ESL model to emulate time series with varying stationarity behaviour is helpful for volatility modelling, as Hansen and Lunde (2014) find that volatility time series are generally close to unit-root. The basic assumption in our framework is that potential investors first seek to gain information about the market (or stock) before trading, hence high attention to the market (measured via Google data) should lead to a growing number of transactions and hence increasing volatility. The analogical relation should hold the other way round. In fact the empirical analysis for DJ data from 01/16/2004 to 10/18/2013 shows that

[4] See http://www.netmarketshare.com

our approach provides reasonable weekly volatility forecasts and significantly outperforms relevant benchmarks like the HAR model, especially in turmoil market phases. Furthermore we find that simply adding a Google component in a linear regression context does not improve the forecasting quality. This result confirms the assumption that the information content in search engine data does not transmit to the volatility process in a linear way. The practical benefit of improved volatility forecasts is highlighted in a Value-at-Risk (VaR) forecasting exercise, where our approach results in more accurate VaR forecasts while requiring less capital. Though the combination of the ESL model and search engine data results in superior forecasts, this approach cannot be applied to arbitrary markets or forecasting horizons. This restriction is due to the fact that Google does not provide applicable daily search data for periods longer than 3 months and that some stock or indices names have multiple meanings, for example the term S&P500 is also relevant in the rating context, which leads to distortions in the search data.

Bibliography

M. Aiolfi and A. Timmermann. Persistence in forecasting performance and conditional combination strategies. *Journal of Econometrics*, 135(1-2):31–53, 2006.

T. G. Andersen and T. Bollerslev. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4): 885–905, November 1998.

D. Andrei and M. Hasler. Investor attention and stock market volatility. Working paper, 2013.

A. Ang and G. Bekaert. Stock return predictability: Is it there? *Review of Financial Studies*, 20(3):651–707, 2007.

N. Barberis. Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1):225–264, 2000.

O. E. Barndorff-Nielsen and N. Shephard. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1):1–37, 2004.

O. E. Barndorff-Nielsen, P. R. Hansen, A. Lunde, and N. Shephard. Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536, 2008.

J. M. Bates and C. W. J. Granger. The combination of forecasts. *Operation Research Quarterly*, 20(1):451 – 468, 1969.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):pp. 637–654, 1973.

J. Y. Campbell and S. B. Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4): 1509–1531, 2008.

T. Cenesizoglu and A. Timmermann. Do return prediction models add economic value? *Journal of Banking & Finance*, 36(11):2974 – 2987, 2012. International Corporate Finance Governance Conference.

J. H. Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008.

J. H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4): 1047–1108, 2011.

F. Corsi. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196, 2009.

G. Elliott and A. Timmermann. Optimal forecast combinations under general loss functions and forecast error distributions. *Journal of Econometrics*, 122(1):47 – 79, 2004.

E. F. Fama. Random walks in stock market prices. *Financial Analysts Journal*, 21 (5):55–59, 1965.

M. A. Ferreira and P. Santa-Clara. Forecasting stock market returns: The sum of the parts is more than the whole. *Journal of Financial Economics*, 100(3):514 – 537, 2011.

E. Ghysels, P. Santa-Clara, and R. Valkanov. Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131(1-2):59 – 95, 2006.

I. Gilboa, O. Lieberman, and D. Schmeidler. Empirical similarity. *Review of Economics and Statistics*, 88(3):433 – 444, 2006.

I. Gilboa, O. Lieberman, and D. Schmeidler. A similarity-based approach to prediction. *Journal of Econometrics*, 162(1):124 – 131, 2011.

V. Golosnoy, A. Hamid, and Y. Okhrin. The empirical similarity approach for volatility prediction. *Journal of Banking & Finance*, 40(0):321 – 329, 2014.

Y. Han. On the economic value of return predictability. *Annals of Economics and Finance*, 11(1):1–33, 2010.

P. R. Hansen and A. Lunde. Estimating the persistence and the autocorrelation function of a time series that is measured with error. *Econometric Theory*, 30: 60–93, 2 2014.

P. R. Hansen, A. Lunde, and J. M. Nason. The model confidence set. *Econometrica*, 79(2):453–497, 2011.

S.-H. Kim and D. Kim. Investor sentiment from internet message postings and the predictability of stock returns. *Journal of Economic Behavior & Organization*, in press:–, 2014.

M. Lettau and S. Ludvigson. Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, 56(3):815–849, 2001.

O. Lieberman. A similarity-based approach to time-varying coefficient non-stationary autoregression. *Journal of Time Series Analysis*, 33(3):484–502, 2012.

L. Liu, A. J. Patton, and K. Sheppard. Does anything beat 5-minute rv? a comparison of realized measures across multiple asset classes. Working paper, 2012.

B. G. Malkiel. *A Random Walk Down Wall Street(First Edition)*. Norton, 1973.

B. G. Malkiel. *A Random Walk Down Wall Street(Revised Edition)*. Norton, 2011.

R. C. Merton. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8(4):323 – 361, 1980.

A. J. Patton. Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246 – 256, 2011.

A. J. Patton and K. Sheppard. Optimal combinations of realised volatility estimators. *International Journal of Forecasting*, 25(2):218 – 238, 2009.

D. Rapach and G. Zhou. Chapter 6 - forecasting stock returns. In G. Elliott and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 2, Part A of *Handbook of Economic Forecasting*, pages 328 – 383. Elsevier, 2013.

G. W. Schwert. Why does stock market volatility change over time? *The Journal of Finance*, 44(5):1115–1153, 1989.

A. Siganos, E. Vagenas-Nanos, and P. Verwijmeren. Facebook's daily sentiment and international stock markets. *Journal of Economic Behavior & Organization*, in press:–, 2014.

P. Veronesi. Stock market overreactions to bad news in good times: a rational expectations equilibrium model. *Review of Financial Studies*, 12(5):975–1007, 1999.

N. Vlastakis and R. N. Markellos. Information demand and stock market volatility. *Journal of Banking & Finance*, 36(6):1808–1821, 2012.

N. Vozlyublennaia. Investor attention, index performance, and return predictability. *Journal of Banking & Finance*, 41(C):17–35, 2014.

I. Welch and A. Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.

L. Zhang. Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach. *Bernoulli*, 12(6):pp. 1019–1043, 2006.

L. Zhang, P. A. Mykland, and Y. Aït-Sahalia. A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 100(472):pp. 1394–1411, 2005.

Chapter 2

Article 1: The empirical similarity approach for volatility prediction.

V. Golosnoy, A. Hamid and Y. Okhrin. The empirical similarity approach for volatility prediction. *Journal of Banking & Finance*, (2014), 40:321-329.

Chapter 3

Article 2: Prediction power of high-frequency based volatility measures. A model based approach.

A. Hamid. Prediction power of high-frequency based volatility measures. A model based approach. *Review of Managerial Science*, (2015), 9(3):549-576.

Chapter 4

Article 3: Forecasting volatility with empirical similarity and Google Trends.

A. Hamid and M. Heiden. Forecasting volatility with empirical similarity and Google Trends.*Journal of Economic Behavior & Organization*, (2015), 117:62-81.

Chapter 5

Summary and outlook

This dissertation illustrated different advances in the task of volatility prediction in stock markets. Since, in contrast to returns, volatility shows very strong and persistent memory, it is, to some extent, predictable. The demand for accurate volatility forecasts in risk management has led to intensive research in this area among academics and practitioners. The theoretical and empirical findings within this doctoral thesis contribute to this research topic. Thereby chapter 2 developed an approach to apply the empirical similarity concept for model (or component) combination. The forecasts achieved with this new model class outperformed classical (RiskMetrics) and recently developed volatility models (HAR) in terms of in-sample fit and out-of-sample forecasting accuracy for major stock markets (DJ, Nasdaq, Nikkei). Though these results are very convincing, there is still a lot of space left for future research. For example theoretical properties of the suggested ES1 and ES2 models have not been investigated yet. Thereby the exact distribution of the parameter estimates is unknown and tests on parameter significance are so far only asymptotically available when assuming normally distributed error terms. Furthermore the identifiability of the ES models has not been examined from a theoretical point of view. In other words the question if different parameter combinations can describe identical time series is unanswered. Besides this incompleteness of the theoretical background, the success in the forecasting exercise motivates further applications of the ES model. First the general quality of the ES model as a tool for forecast combination could be investigated by comparing it with other recent or standard methods in this context, like pooling (Aiolfi and Timmermann,

2006) or the methods described in Bates and Granger (1969). Moreover an extension to the multivariate case is desirable. Such a multivariate empirical similarity (MES) model could help to get rid of the linear and static parameter estimates obtained when applying the widespread Vector Autoregressive (VAR) model for multivariate time series analysis. As each realisation of a relevant univariate time series at time t can be interpreted as a state of nature in the ES context, a system of contemporary univariate ES models can be a first step towards the MES model. However it is likely that the dependence structure between the univariate time series still has to be modeled separately.

The main contribution of chapter 3 is an empirical analysis which investigates the impact of the interplay between popular volatility measures, models and loss functions on the forecasting performance. The findings help to resolve the puzzling results in prior empirical literature on this topic, where many different volatility measures have been declared to have superior predictive ability. These apparently paradoxical statements are due to the fact that different volatility models usually prefer different volatility measures. The application of different loss functions to evaluate forecasts exacerbates the disagreements. Though the conducted empirical analysis in chapter 3 is already quite extensive there remain some interesting research questions. For instance the relation between data quality and the preferred measure has not been investigated, since average results for a large dataset have been computed. This dataset consists of indices from liquid markets with a high number of daily transactions and other indices with few daily transactions. It is possible that some measures (for instance subsampled ones) cannot be computed reliably when the number of daily transactions is low, since this is equivalent to a small number of intraday return observations. On the other hand the conducted robustness check implies that the ES model produces best forecasts if the volatility proxy to be predicted and the applied measure stem from the same time series. However this result has to be validated by expanding the analysis in section 4.4 of article 2 to all volatility measures used. In addition to the empirical results, a theoretical background explaining why some model-measure pairs are better suited than others could lead to a deeper understanding of the relation between model and measure. Furthermore the link between statistic and economic loss functions, an aspect that is briefly outlined in section 3.1 of article 2 deserves a closer look. Taylor (2013) delivers very interesting preliminary work in this field.

Chapter 4 applies Google search frequencies to the empirical similarity model of Lieberman (2012) (ESL) in order to augment weekly volatility forecasts for the Dow Jones. Thereby Google data is seen as a indicator variable for future volatility. The basic assumption within this approach is that increasing investor interest in the Dow Jones (measured via search volume) translates into higher market participation and consequently increasing volatility. The empirical results support this theory and illustrate that the information content of Google search frequencies can only be used to full capacity when combined with the ESL model, and not as an additional linear regressor in standard time series models. Since the inclusion of search engine data and the application of the ESL model into the task of volatility forecasting is a rather new field, further research questions have emerged. In chapter 4 the predictive ability of Google search data is only investigated for the local market, however it might be possible that search frequencies in economically leading countries can also be exploited for volatility prediction in other countries. This topic is highly related to the growing literature on spill over and contagion effects. Of course it could also be examined if the ESL model can be successfully applied to other indicator variables for volatility, like survey based disagreement measures on the expectation of stock market movements. Li and Li (2014) utilize such survey based data to analyse the link between the beliefs of household investors about macroeconomic conditions and volume in the S&P500. However they do not provide sufficient forecasting tools. Another research question originates from the fact that the ESL model can be interpreted as a dynamic AR(1) model with a time varying autoregressive coefficient and is capable of modelling time series that consist of stationary, unit-root and explosive phases. Hence a closer analysis of the properties of this dynamic autoregressive coefficient might help to derive new tests for structural breaks or bubble occurrences in time series.

To conclude, one can say that there has been a lot of progress in the area of volatility modelling and prediction, however it is likely to remain an active research field in the upcoming years, since there is still a large number of open questions.

Bibliography

M. Aiolfi and A. Timmermann. Persistence in forecasting performance and conditional combination strategies. *Journal of Econometrics*, 135(1-2):31 – 53, 2006.

J. M. Bates and C. W. J. Granger. The combination of forecasts. *Operation Research Quarterly*, 20(1):451 – 468, 1969.

D. Li and G. Li. Are household investors noise traders: Evidence from belief dispersion and stock trading volume. Working paper, 2014.

O. Lieberman. A similarity-based approach to time-varying coefficient non-stationary autoregression. *Journal of Time Series Analysis*, 33(3):484–502, 2012.

N. Taylor. The economic value of volatility forecasts: A conditional approach. *Journal of Financial Econometrics*, 2013.