Efficient Daily Volatility Forecasts and Forecast Revisions Using Intra-day Data

Wesley B. Ballering - 326828

September 24, 2012

Master Thesis

Master Quantitative Finance

Erasmus University Rotterdam

Supervisor

Prof. Dr. D.J.C. van Dijk

Co-reader

Prof. Dr. M.J. McAleer

Abstract

The objective of this thesis is to examine the predictive performance of begin-of-the-day volatility in forecasting end-of-the-day volatility using intra-day data; to search for an efficient timing window to make our forecast and to assess the process of forecast revisions during the day. For this we develop three methods that make use of diurnal patterns: the average seasonal over the last K days, an exponentially weighted average seasonal and a Flexible Fourier Form alternatively we build a method on Mincer-Zarnowitz bias adjustments. As the actual volatility is an unobserved variable several high-frequency volatility estimators: Realized Variance, Bipower variation, Realized Range, Two Time Scales and Kernel estimation are proposed to base ex-ante forecasts and measure ex-post forecasting performance. We conclude that begin-of-the-day volatility is a highly predictive measure for end-of-the-day volatility and can be improved by scaling and bias adjustments. The assessment of efficient timing samples leads us to believe that roughly the first 15-20 min are most informative, with the remark that omitting the first 0-10 minutes of return data could be beneficial as noise prevails during this interval. Furthermore, correlations between lagged and current volatility forecast revisions is found to be small yet statistically significant, implying our models inherited some build in smoothing property.

Keywords: Integrated Volatility, Diurnal Patterns, Volatility Forecasting , Flexible Fourier Form, Intra-day Data, Fixed Event Forecast revisions

Contents

1 Introduction

Volatility figures prominently in the arena of risk management, is central to portfolio formation and constitutes a pivotal role as input parameter for a variety of derivative pricing models. Creating more accurate forecast and developing a greater understanding of the nature of volatility and its determinants is therefore of utmost importance to investors, banks, insurance companies, pension funds, institutions, corporations, governments and all others that are exposed to volatility risk.

Since the introduction of Autoregressive Conditional Heteroskedasticity (ARCH) type models by Engle (1982) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) type models as proposed Bollerslev (1986) there has been a proliferation for models of this type[1]. A different class of volatility models comprises of Stochastic Volatility models, first proposed by Taylor (1982) based on the ideas of Black (1976) and further developed by Nelson (1991) amongst others. Although these models create a nice mindset practically they do, however, have some drawbacks bound to their potentially restrictive parametric nature and empirically dubious underlying assumptions. Therefore, new non-parametric integrated volatility (IV) type estimators, build upon the theory quadratic variation, have been gaining ground. Being non-parametric, rapidly adjusting, intuitive and observable measures of volatility they do indeed have their advantages.

Most well known estimator of ex-post integrated volatility is the Realized Variance (RV), i.e. the sum of frequently sampled squared returns, as discussed by Barndorff-Nielsen and Shephard (2002). In an ideal world this would be an unbiased, consistent and highly efficient estimator of the true integrated volatility. With the use of high frequency intra-day data containing microstructure noises like bid-ask bounce, discretization, liquidity effects, minimum tick sizes, misrecordings etc. this however no longer holds. Therefore other methods of ex-post volatility estimation have been proposed like Realized Bipower Variation (BPV) see Barndorff-Nielsen and Shephard (2003), Realized Range (RR) see Martens and Dijk (2007), Two Time Scales (TTS) see Zhang et al. (2005) and Kernel estimation see Barndorff-Nielsen et al. (2008) among other methods.

In this research the focus is on predicting daily variance[2] using high frequency intra-day data. Specifically return data close to the start of open outcry trade is used to forecast daily variance as obtained through above mentioned IV measures. The informational content of begin-of-the-day volatility in itself is tested but forms a highly biased predictor for daily volatility. In order to rescale begin-of-the-day volatility we build upon the empirical regularity that intra-day volatility exhibits

[1] For a thorough list of available ARCH derivatives see Bollerslev (2008)

[2] During this research daily variance is defined as the integrated variance over open outcry trade only.

strong cyclical behavior in the sense that start and end of the day volatility is on average high, with a gradual decline to the middle. Creating a distinct U, J, or inverse J-shape throughout the trading day, see Wood et al. (1985), Muller et al. (1990) and Baillie and Bollerslev (1991). Subsequently these shapes can thus be exploit to rescale begin-of-the-day volatility to daily proportions. Three approaches are conducted to obtain this assumed deterministic intra-day seasonal pattern. A simple Moving Average method is used to obtain the periodic volatility component as the average volatility over a preceding interval up to the forecast; a Exponentially Weighted Moving Average (EWMA) is used as faster adjusting alternative to the former and a more advanced Fourier Flexible Form (FFF) is used to model the dynamics of intra-day volatility patterns. Additionally a bias adjustment method to begin-of-the-day variance is built upon Mincer-Zarnowitz coefficients using rolling window MZ regressions.

Through this framework we extend the research of Frijns and Margaritis (2008) and Harju and Hussain (2008) in four ways: (1) We use multiple measures of ex-post volatility to gain robustness. (2) We do not restrict ourselves to start-of-the-day volatility samples but rather loop over different start and ending times as to obtain the most efficient estimation sample. Hereby basically investigating whether early return data, just after start open outcry trade, actually increases forecasting ability or weakens it due to the large noise to signal ratio inherent to this data. (3) Evaluate performance using additional measures and (4) enhance forecasting performance by a not earlier conducted bias-adjustment methods and through techniques of Kalman Filtering on the unobserved Mincer-Zarnowitz α and β states. Furthermore, related to the search for an efficient forecasting period, revisions on end-of-the-day volatility forecasts through the day will be evaluated. Under weak-form market efficiency lagged revisions and current revisions should be independent and thus have zero correlation. If this is not the case one could, in the spirit of Franses et al. (2011), investigate whether it is overreaction to news, some kind of smoothing behavior or something else causing correlation in forecast revisions.

To our knowledge Frijns and Margaritis (2008) are the first and only to have investigated forecasting power of start-of-the-day volatility to end-of-the-day volatility. They find that first-hour volatility in itself is highly predictive for end-of-the-day volatility and leads to coefficients of determination, R^2, as high as 68% where less than 30% of daily volatility is observed by that time. Rescaling with an average seasonal leads to improvements and forecast combinations using FFF seasonals with GARCH(1,1) forecasts are superior to any single prediction. Our finding on S&P500 index futures returns and US 30 year treasury bond futures returns support their earlier work. We find a squared robust correlation (R^2_{MAD}) between forecasts and daily volatility of

0.68 in S&P500 after just 30 minutes of start-of-the-day volatility in itself, equaling only 13% of total daily variance. Rescaling by an (Exponentially Weighted) Moving Average Seasonal or FFF seasonal further improves upon these figures, where the latter improves mainly in terms of smaller HMSPE. Subsequently adjusting forecasts using Kalman Filtered Mincer-Zarnowitz states reduces Heteroskedasticity-consistent Mean Squared Prediction Error (HMSPE) by a factor 2 as compared to FFF forecasts. US30 volatility proves somewhat harder to predict yet has greater benefit from seasonal modeling. After 30 minutes, equaling on average 14% of daily variance, squared correlations are on average 0.51. However adjusting these forecasts using FFF seasonal scaling and again Mincer-Zarnowitz α and β states, improves predictability appreciably. Through Kalman filtering procedures R^2_{MAD} on average grew up to 0.57 whereas HMSPE diminishes by a factor 4 against FFF. However simpler alternative found in scaling FFF forecasts by Mincer-Zarnowitz regression coefficients obtained from rolling window regressions yields similar, and during early start of the day even better, results. Leading to R^2_{MAD} of up to 0.75 after 30 minutes and diminishing HMSPE up to a factor 2 compared to FFF forecasts.

Choosing an efficient data sample to base forecasts could improve R^2_{MAD} figures even further. Dependent on the security at hand, the first 0-10 minutes of return data should be omitted in order to increase correlation. Evidently the noise during this early interval corrupts forecasts more than it favors them. Furthermore in terms of efficiency one could argue that, as additional R^2_{MAD} is outweighed by the additional daily variance needed, a saturation point is found. For S&P500 and US30 such can be found at respectively 15 and 20 minutes after start open outcry trade. Meaning roughly interval [5-15] and [5-20] minutes after start of open outcry trade are most efficient to base forecasts for resp. S&P500 and US30 volatility.

For the remainder of this thesis the plan is as follows. Section 2 is used to discuss the data, it's high frequency features and some stylized facts. Section 3 concentrates on different methods for obtaining ex-post volatility measures, different kinds of ex-ante volatility models and statistical evaluation procedures and bias adjustments. Section 4 is used to evaluate forecasting performance and other statistics. Concluding remarks are given in section 5.

2 Data

This section will describe the data used for this thesis with addition of theoretical as well as practical implications. Some stylized facts are given; potential hazards are highlighted and some adjustments are made.

2.1 Data description

Data provided for this research comprehends 10 second price levels[3] for the Standard & Poor's 500 Index futures (henceforth S&P500) starting at 05:00:10 GMT 2003/07/01 up until 21:15:00 GMT 2009/12/31, totaling 6 and a half years or 1642 active trading days. This index is calculated by Standards & Poor's Corporation as the capitalization weighted average of the greatest 500 US based companies which meet their admission requirements. Globally it is seen as the most reliable indicator of stock market developments and therefore much used as subject of research. For reasons of comparability we propose to do the same.

Second data series contains 10 second price levels for 30 year US Treasury Bonds futures (henceforth US30) over the period 05:11:50 GMT 2003/07/01 up until 22:00:00 GMT 2009/12/31 totaling 6 and a half years or 1622 active trading days. Throughout the whole sample electronic trading was available having some substantial implications as: reduced cost of transactions; greater liquidity; increased transparency and tighter spreads, thereby diminishing the microstructure noise effects.

As a basis for later end-of-the-day or daily volatility calculations, i.e. the variable of interest in forecasting, we take the open outcry trading hours in the regarding securities. For the S&P500, whose constituents are solely traded on the NYSE, AMEX or NASDAQ, this means we start at 9:30 EST and end on 16:15 EST. For the US30, traded mostly in New York and Chicago, this means we start at 8:20 EST and end on 17:00 EST. Subsequently all data is used to create forecasts yet with the difference that after-trade and overnight returns are not in itself part of the daily volatility as defined here[4].

[3] For reasons that will become clear in the methodology section, calendar-time-sampling (CTS) is used with 10 second price levels obtained as the last available price from tick-by-tick data. This does impose some stickiness of prices resulting in negative autocorrelation in the return series as further explained in section 2.2.1. Alternatively one could use business-time-sampling (BTS) or tick-time-sampling (TTS).

[4] Ignoring overnight return observations in the creation of daily volatility using high frequency intra-day data is in accordance with Andersen et al. (2001); Thomakos and Wang (2003); Corsi et al. (2008); Wu (2010). Overall there is no consensus on the in- or exclusion of overnight return observations in daily volatility constructed from volatility measures using high frequency intra-day data. For a further exposition on different choices defining daily volatility, see Ahoniemi and Lanne (2010).

The price levels are converted to continuously compounded return series by taking the logarithms and subtracting the previous value. For ease of notation we normalize the daily interval to unity. The return over the i-th interval of length Δ on day t, for $i = 1, 2, ..., n$ and n the number of daily increments, can then be calculated as:

$$r_{t,i} = ln(P_{t,i\Delta}) - ln(P_{t,(i-1)\Delta}) \tag{1}$$

Where $P_{t,i\Delta}$ is the price level at day t after i time intervals of length Δ.

2.2 Stylized facts of high frequency intra-day data

The use of high frequency intra-day data is both revealing and problematic. It has opened up all new possibilities to observe volatility in far greater detail than ever before. Features like diurnal patterns, that would otherwise be hidden in lower data aggregations, can now be revealed. It, however, also poses new difficulties concerning mostly the impact of microstructure noise. Irregular spacing resulting in stickiness of prices; price discreteness; minimum tick-size see Munnix et al. (2010) and bid/ask spread see Gatheral and Oomen (2010), are just some of the noise constituents corrupting high frequency data. As so, Andersen and Bollerslev (1998b) find that these errors have the tendency to mask the strong persistence in underlying latent volatility dynamics. It would therefore be beneficial to (1) select our data carefully, (2) cautiously clean our data and (3) adjust our models as to cope with these disruptive sources. For the interested reader an intuitive and more thorough discussion on microstructure effects is included in Appendix A: Microstructure effects and can be read before continuing.

2.2.1 Data cleaning

Careful cleaning of high-frequency data is an important task. Obvious errors like misrecordings, wrong placement of decimals or outliers can have tremendous impact on volatility estimation. Yet so can cleaning of data, altering it's statistical properties, see Falkenberry (2001) amongst others. Hansen and Lunde (2006) on the other hand show that volatility estimation can improve by tossing out vast amounts of data. In our case we keep removal to a minimum and propose the following adjustments:

1. We delete days with no open outcry trade where open outcry is defined as 9:30 EST until 16:15 EST for S&P500 index futures returns and 8:20 EST until 17:00 EST for the US30 year treasury bonds futures returns. These mostly constitute holidays and some recording

errors. For the S&P500, 32 such days were deleted from our sample, for the US30 these constitute to 24 days being deleted.

2. We check for outliers caused by decimal misplacement or other recording errors. None such errors were found throughout both series.

3. Due to the financial crisis of last couple of years, rather extreme return observations found their way into the S&P500 and US30 data sets. Creating large jumps in the price process, these observations are both interesting and potentially hazardous to later conclusions. In order to overcome such events two options may be considered: 1) thoroughly clean all data series as to suffice underlying model assumptions in an ordinary setting. Yet such rigorous cleaning could alter statistical properties and damage data integrity. Second, one could rely on robust performance measures. The latter is chosen throughout this thesis and indeed has major influence. The introduction to section 4 lays bare a graphical illustration to the potency of ordinary statistics to distort conclusions.

Remaining spikes in the daily pattern are due to regular recurring events like business figures or macro-economic news announcements and should therefore not be altered. They form an integral part of the daily trade pattern and should be accounted for when estimating a seasonal pattern. Most evident spikes are summed in table 1 and synchronize well with general timing of macroeconomic news announcements and trading routines. As can be seen from figure 1 and 2, the S&P500 data is quite sensitive to the opening of exchange markets as well as macroeconomic news announcements whereas for the US30 macroeconomic announcements have by far the greatest impact.

S&P 500 index		US 30y Treasury	
Spikes at:	*Cause:*	*Spikes at:*	*Cause:*
08:30.00-08:30.10	Macroeconomic news announcements at 8:30	08:20.00-08:20.10	Opening markets New York and Chicago
09:30.00-09:30.10	Opening markets New York and Chicago	08:30.00-08:30.10	Macroeconomic news announcements at 08:30
10:00.10-10:00.20	Macroeconomic news announcements at 10:00	10:00.00-10:00.10	Macroeconomic news announcements at 10:00
16:00.00-16:00.10	Brokers clean their sheets	13:01.30-13:01.40	Unknown
16:14.30-16:14.40	Electronic trading about to close	14:59.00-14:59.10	Just ahead of open outcry close
16:30.00-16:30.10	Reopening electronic trading desks	16:59.50-17:00.00	Just ahead of close electronic trading
18:00.00-18:00.10	Reopening electronic trading desks	18:30.00-18:30.10	Reopening electronic trading desks
Note 1: All New York based exchanges open at 9:30 and close at 16:00		Note 1: Bond trade opens at 08:20 and closes at 15:00	
Note 2: S&P500 data assembly is temporarily closed from: 09:15-09:30, 16:15-16:30 and 17:30-18:00		Note 2: US30 data assembly is temporarily closed from 17:00-18:30	

Table 1: Volatility spikes in the 10 seconds absolute return series for the S&P500 index and US 30 year treasury bonds together with their most likely cause. Evidently spikes synchronize with timing of macroeconomic news announcements and recurring practices/routines. (time of day is given in EST)

2.2.2 Diurnal patterns

Another stylized fact of high-frequency data, especially for equity returns, is that it exhibits greater volatility during the start and end of the day with a gradual decline to the middle. Among the first to notice this typical U-, J- or inverted J-shaped pattern in equity data were Wood et al. (1985) and Harris (1986). Muller et al. (1990) and Baillie and Bollerslev (1991) found equal characteristics for foreign exchange markets and Harju and Hussain (2006, 2008) further assess the implications of these patterns. It is this empirical regularity that will be utilized in this paper to produce superior volatility forecasts. As can be seen in figure 1, S&P500 future returns exhibit a comparable U-shaped pattern, which is far more pronounced for high first-hour volatility days than is the case for low first-hour volatility days. They do, however, not support earlier findings that high volatility days tend to follow an inverse J-shape where low volatility days experience more of a U-pattern. Yet this could be due to round the clock electronic trading. As securities are restricted to mere traditional trading hours, overnight information can only demonstrate itself during daytime trading. Resulting in higher start-of-the-day volatility and therefore a more inverse J-shaped pattern. Figure 2 shows the average absolute daily return for the US 30 future returns. These exhibit a less stylized shape, although high volatility days now do on average follow a inverse J like pattern.

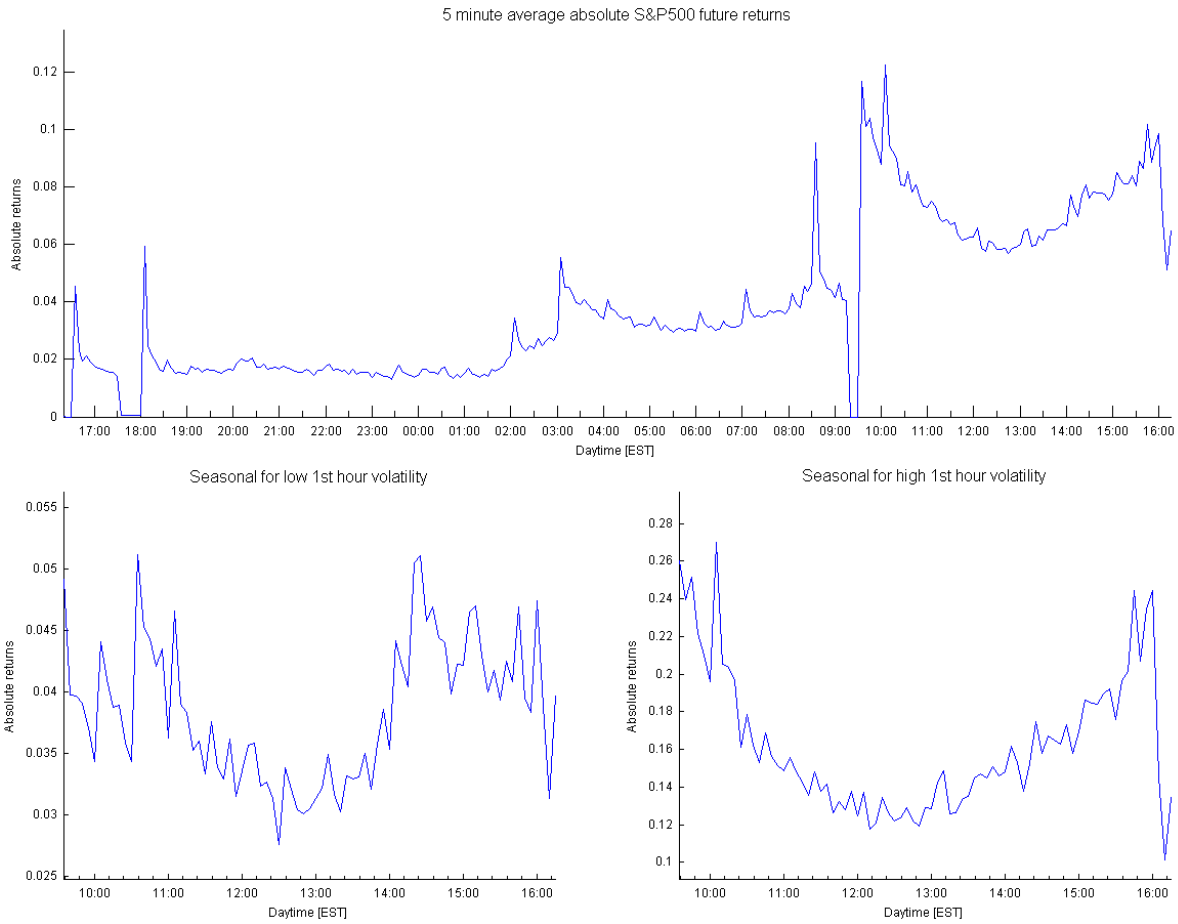

Figure 1: Mean absolute S&P 500 index futures return series at 5 minute sampling frequency over period 2003/07/01 - 2009/12/31 and seasonal shapes for 20% lowest and highest first hour volatility days. The seasonal shape during open outcry trading hours, i.e. 09:30 - 16:15, is far more pronounced for high volatility days than is the case for low volatility days. Note that electronic trading, or at least the assembly of data by Standard & Poor's Corporation, is paused during the following periods: 09:15-09:30, 16:15-16:30 and 17:30-18:00.

In line with Andersen and Bollerslev (1997) the autocorrelation patterns for both S&P500 and US30 absolute return series were plotted to gain a more thorough insight in this seasonal effect. The correlograms of these series together with autocorrelation plots for deseasonalized returns are presented in Appendix A: figure A.3[5]. It is evident from these plots that the S&P500 and US30 data both experience strong cyclical behavior with peaks at the 1 day lag due to the seasonal. When deseasonalized however, much of the cyclical effect is taken away, leaving a smooth geometric decay. Only little peaks at the daily lag remain indicating that the seasonal component is not totally filtered out, yet it did a pretty good job.

[5]Deseasonalizing of returns is done by fitting a Fourier Flexible Form to the day and subsequently calculating returns as $r_{t,n,deseason} \equiv r_{t,n}/s_{t,n}$ with $s_{t,n}$ the seasonal component of returns estimated by a Fourier Flexible Form as discussed in section 3.2.5. As is done in this section we take $J = 1$ and $P = 1$ for S&P500 index future returns and $J = 1$ and $P = 2$ for US 30 year treasury bond futures returns. To account for the flat non trading times during the day, an extra dummy variable was introduced in deseasonalizing as to adjust for this.

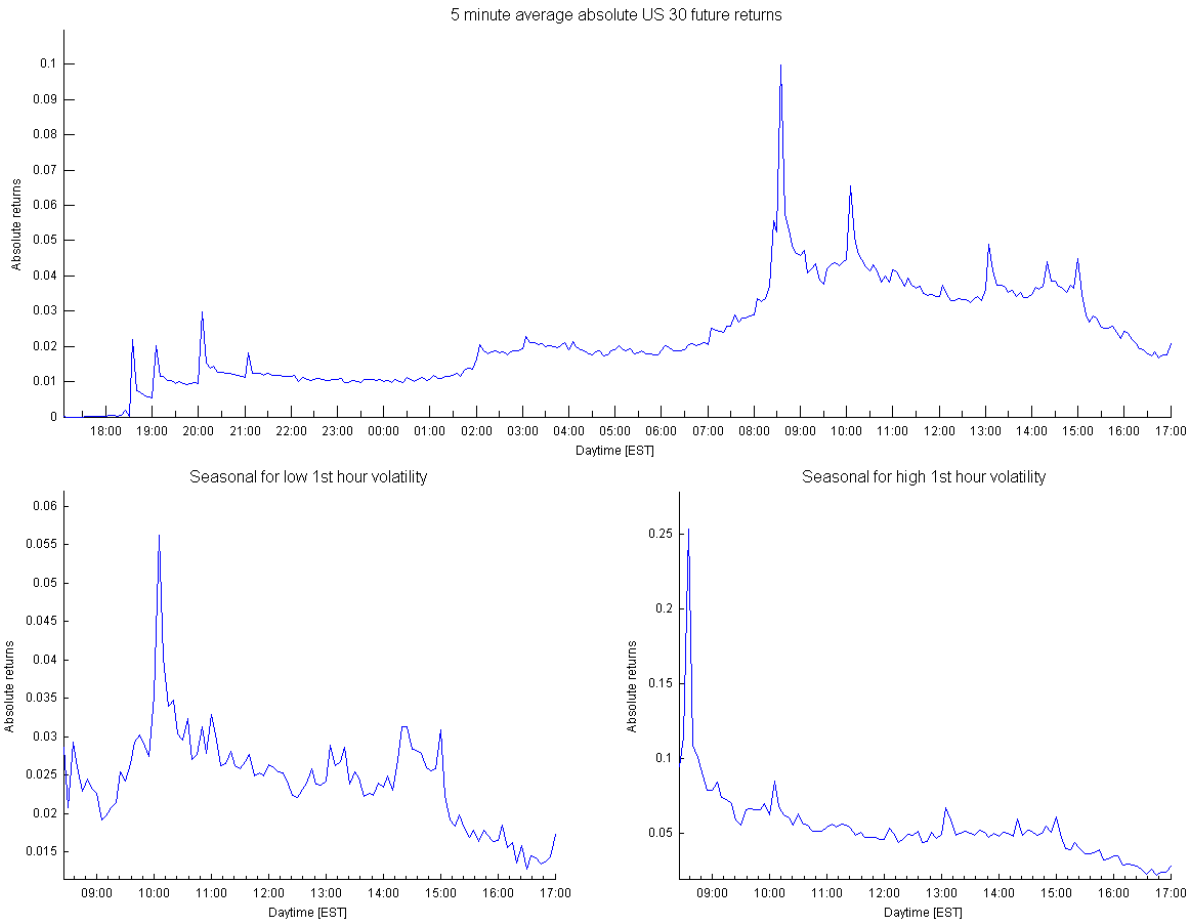

Figure 2: Mean absolute US 30 futures return series at 5 minute sampling frequency over period 2003/07/01 - 2009/12/31 and seasonal shapes for 20% lowest and highest first hour volatility days. The seasonal shape during open outcry trading hours, i.e. 08:20 - 17:00, is less erratic than is the case for low volatility days. The seasonal shape is however clearly less stylized than the one for equity future returns. Note that electronic trading is paused during the period: 17:00-18:30.

Correlations for absolute S&P500 series are up to 0.4 and the mean level dies out only at a very slow rate. They are all positive implying that a volatile start of the day is most often followed by a volatile day overall, as consistent with extensive literature documenting volatility clustering in asset returns, dating back at least to at least Mandelbrot (1963) and Fama (1965). The correlations for absolute US 30 year treasury bond futures feature about the same properties. They also exhibit strong cyclical behavior caused by intra-day seasonal volatility patterns. Biggest difference is to be found in the lower overall level of correlations, around 0.2, demonstrating again the less pronounced diurnal pattern US bond future returns face. Taking 24 hour trading the autocorrelations even turn negative confronting half a day lag, whereas this trend disappears looking at the open outcry returns only[6]. It seems that a day starting rough settles down stronger

[6]Open outcry return series are taken from 9:30-16:15 for S&P500 data and 8:20-17:00 for US30 data, as opposed

in the after trading hours, as if to take a breath.

2.3 Data statistics

After adjusting the sample, e.g. cleaning for outliers and removing sparse trading days the following summary statistics for 5 minute return series and open-outcry-only 5 minute returns series can be obtained.

	S&P 500	S&P 500*	US 30	US 30*
Minimum [%]	-2,8422	-2,7886	-5,1923	-2,1815
Maximum [%]	4,3634	3,7844	3,3424	3,3424
Mean [%]	1,9443E-05	-8,7488E-05	-2,7858E-06	1,0129E-04
Median [%]	0,0000	0,0000	0,0000	0,0000
Standard Deviation	0,0814	0,1254	0,0426	0,0582
Skewness	1,1662	0,2860	-6,6170	0,5015
Kurtosis	110,1776	34,3244	953,0225	105,1092
AC(1) return	-0,0249	-0,0216	-0,0455	-0,0349
AC(1) absolute return	0,4366	0,3813	0,3068	0,2667
Percentage of daily volatility	100,0000	54,8850	100,0000	59,0105
Percentage zero returns	30,1062	7,8728	46,2801	27,8947

Table 2: Descriptive statistics for 5 minute S&P500 index returns and US 30 year treasury bond futures returns from July 1st, 2003 until December 31, 2009. Series denoted with a star (*) are taken over open-outcry trade only.

The S&P500 is most volatile, especially during regular trading hours, with a standard deviation of 0.1254 or an annualized standard deviation of approximately 33.78%[7]. Al returns are highly leptokurtic, sample minimum/maximum returns are between 34/53 and 121/78 standard deviations from their mean for resp. S&P500 and US30. Assuming normality the probability of observing such extreme values are practically zero. First order return autocorrelations are small, negative, yet significant for both markets, this can typically be attributed to market microstructure effects, see Hansen and Lunde (2004). Whereas first order autocorrelations for absolute returns are high and highly significant implying a great deal of volatility persistence.

to round the clock returns for the standard sample. This yields better insight in daily volatility faced by the majority of (institutional) investors.

[7]Approximate annualized Standard Deviations are obtained as $std_{5m} \cdot \sqrt{288} \cdot \sqrt{252}$, for 288 daily 5 minute increments and approximately 252 trading days a year. Note that this is an approximation as the 'square root of time rule' only holds for i.i.d. return data. For further elaboration on this matter see Danielsson and Zigrand (2006).

3 Methodology

This section discusses the methodologies used in creating ex-post volatility measures and ex-ante volatility forecasts. Furthermore methods for interpretation of Mincer-Zarnowitz regressions are presented and variance forecast revision are discussed. Let us start by explaining what integrated volatility (IV) is and why we use it.

3.1 Integrated Volatility

Volatility is the degree to which financial returns (log price differences or relative price changes) fluctuate over time. Widely it is seen as the most important measure of risk. Measuring, modeling and understanding this latent variable has therefore been a point of great attention. On lower data aggregations ARCH type models form a workable solution to modeling volatility, working with high frequency data though, this approach seems ill suited. As noted by Andersen and Bollerslev (1997) "ARCH models imply a geometric decay in the return autocorrelation structure and simply cannot accommodate strong regular cyclical patterns" as often observed in high frequency intra-day data. The parameter-free diffusion models reviewed in this section are better suited to handle this type of data. To put our mindset in mathematical terms, we take the price process $X_{t,i} = ln(P_{t,i})$ to follow an Itô process,

$$dX_t = \mu_t dt + \sigma_t dB_t \tag{2}$$

where B_t is a standard Brownian motion, μ_t is a drift term and σ_t the instantaneous volatility of returns process X_t, assumed to follow a càdlàg process. Integrated variance is now defined as the integral:

$$IV = \int_0^1 \sigma_t^2 dt \tag{3}$$

The integral of the instantaneous variance between successive time stamps 0 and 1. The designation Integrated Volatility and Integrated Variance are in practice used alongside each other to denote a measure of riskiness. Here we will however reserve the name integrated volatility for the square root of integrated variance. If prices $P_{t,i}$ could be observed continuously and without error this would be the whole story. Taking Realized Variance (RV) as measure would, as can be seen later on, make IV visible and proxy it in a unbiased, consistent and highly efficient way. Unfortunately, in reality we do not observe such true log prices $X_{t,i}$ but rather a proxy $Y_{t,i}$,

$$Y_{t,i} = X_{t,i} + \varepsilon_{t,i} \tag{4}$$

where noise constituents $\varepsilon_{t,i}$ are mutually independent and and jointly independent of X. Consisting of earlier mentioned market microstructure effects, this noise has a major effect on volatility estimation when using truly high frequency data like one second prices or even higher. Making the increments infinitely small the volatility of the true price process diminishes to zero. The volatility of the error term however stays more or less the same yielding an exploding noise-variance to signal-variance ratio dependent on the sample rate. Following the theory of quadratic variation this can be made clear following the understated formulas

$$r_{t,i} = \Delta Y_{t,i} = \Delta X_{t,i\Delta} + \Delta \varepsilon_{t,i\triangle} \qquad \text{with } \Delta Y_{t,i} = Y_{t,i\Delta} - Y_{t,(i-1)\Delta}$$

$$Var(r_t) = \sum_{i=1}^{n} (\Delta Y_{t,i\Delta})^2 = \sum_{i=1}^{n} (\Delta X_{t,i\Delta})^2 + \sum_{i=1}^{n} (\Delta \varepsilon_{t,i\triangle})^2 + 2\sum_{i=1}^{n} (\Delta X_{t,i\Delta}, \Delta \varepsilon_{t,i\triangle}) \tag{5}$$

Giving the conditional mean of the variance of r_t:

$$E\left[Var(r_t) \middle| X_{process}\right] = \sum_{i=1}^{n} (\Delta X_{t,i\Delta})^2 + E\left[\sum_{i=1}^{n} (\Delta \varepsilon_{t,i\triangle})^2\right]$$

$$\text{with} \quad \sum_{i=1}^{n} (\Delta \varepsilon_{t,i\triangle})^2 = \sum_{i=1}^{n} \varepsilon_{t,i\triangle}^2 + \sum_{i=1}^{n} \varepsilon_{t,(i-1)\Delta}^2 - 2\sum_{i=1}^{n} \varepsilon_{t,i\Delta}\varepsilon_{t,(i-1)\Delta}$$

$$\text{so that} \quad E\left[\sum_{i=1}^{n} (\Delta \varepsilon_{t,i\triangle})^2\right] = 2n\sigma_{\varepsilon,t}^2$$

$$E\left[Var(r_t) \middle| X_{process}\right] = \sum_{i=1}^{n} (\Delta X_{t,i\Delta})^2 + 2n\sigma_{t,\varepsilon}^2 \tag{6}$$

where $Var(r_t)$ is the RV estimate, $\sum_{i=1}^{n} (\Delta X_{t,i\Delta})^2$ is the true IV when $\triangle \to 0$ and therefore $n \to \infty$, σ_{ε}^2 is the variance of the error term and n is the number of daily increments. As sampling rate grows, so does the error variance with a factor $2n$. When n goes to infinity the error variance explodes taking over all of the estimated variance. We are now actually consistently measuring noise variance instead of IV.

One question though naturally comes to mind: why do we actually want to obtain the variance

of X_t rather than the variance of Y_t. As in principle Y_t conceals the variance of market prices so this is the variance we face in practice. Mean reason is that the error variance is tied to transaction costs and is artificially created by the mechanics of the trading system whereas the variance of X_t tells us more about the volatility of the underlying process. Moreover as mentioned above, the variance of Y_t would depend on data frequency which is not a desirable feature. Hence $Var(X_t)$ is our target for this research. In the remainder of this subsection we briefly go through a selection of well-known IV estimation measures. Again for the interested reader a thorough expositions of the variance estimators, all underlying assumption, sensitivities and details on implementation are to be found in Appendix B.

3.1.1 Realized Variance

Realized variance (RV) is the first and most basic measure of integrated volatility. According to the theory RV can recover the volatility defined by the quadratic variation of a semi martingale price process[8]. Using high frequency returns it can be estimated by $\sum_1^n (Y_{t,i\Delta} - Y_{t,(i-1)\Delta})^2$, resulting in the formula for realized variance.

$$RV_t = \sum_{i=1}^{n} r_{t,i}^2 \tag{7}$$

where $r_{t,i}$ is the return $\left(Y_{t,i\Delta} - Y_{t,(i-1)\Delta}\right)$ over period $[i-1, i]$. In an ideal world with no market frictions this would be an unbiased, consistent and highly efficient estimator for integrated variance. One would merely have to drive up the sampling frequency to obtain a more accurate estimate and as $n \to \infty$, realized variance converges to IV. Unfortunately reality is not this structured. Practical implementation has to confront the fact that prices are not recorded continuously and markets are not frictionless. Leading RV to be biased and inconsistent resulting in great errors if one was to use RV on high sampling frequencies without correction.

Having its flaws, the ease of computation makes RV a much used IV estimation technique nonetheless. To deal with the bias in practice, another solution is favorite. Sparse sampling is used as to mitigate microstructure effects. Though this tactic merely limits the influence of noise rather than corrects for it, Hansen and Lunde (2006) advocate it works quite well.

[8] See Jacod and Shiryaev (1987) and Barndorff-Nielsen and Shephard (2003) for further exposition on quadratic variation (QV) assuming semimartingale properties for log price process.

3.1.2 Realized Bipower Variation

Due to its quadratic form realized volatility is quite sensitive to outliers and jumps in the return process. A more robust estimator can be found in Realized Bipower Variation (BPV). By taking the product of subsequent return observations it mitigates the influence of outliers and is far less affected by jumps in the (log) price process.

For this research Realized Bipower Variation is defined as:

$$BPV_t = \sum_{i=1}^{n-1} |r_{t,i}||r_{t,i+1}|, \qquad r = s = 1 \tag{8}$$

Details on working, implementation and derivation are again given in Appendix B

3.1.3 Realized Range

Another intuitive measure of volatility estimation is the Realized Range (RR) which makes use of the difference in maximum and minimum observed prices during a certain period of time. Properly scaled, daily high-low range is up to 5 times more efficient than realized variance using daily squared returns. Or correspondingly performs similar to realized variance sampled at 4 to 8 times higher frequency. If this result holds for every sampling rate, in theory we'd have an ever more efficient estimator than RV. In this mindset Martens and Dijk (2007) proposed to use high-low range with intra-day data and dubbed the resulting estimator Realized Range.

To formalize the setting we again take $P_{t,i\Delta}$ to be the last observed price in the i-th interval of length $\triangle$, as we have already assumed in equation (1). $H_{t,i} = sup_{(i-1)\Delta<j<i\Delta} P_{t,j}$ is the highest observed price or supreme and $L_{t,i} = inf_{(i-1)\Delta<j<i\Delta} P_{t,j}$ is the lowest observed price or infimum. The scaled high-low range estimator is than defined as

$$\frac{(lnH_{t,i} - lnL_{t,i})^2}{4ln2} \tag{9}$$

Aggregating over the n daily intervals gives the realized range:

$$RR_t = \frac{1}{4ln2} \sum_{i=1}^{n} (lnH_{t,i} - lnL_{t,i})^2 \tag{10}$$

To further improve upon this method Martens and Dijk propose the use of a bias adjusted RR estimator. It is this bias adjusted RR version that will be used throughout this research. Further details and intuition can be found in Appendix B.

3.1.4 Two Time Scales Estimation

First two volatility estimators have one thing in common. Taking the empirically much used 5 minute intervals, they throw away vast amounts of data just to reduce the influence of microstructure effects. In our case this sparse sampling encompasses throwing away approximately 97% of the data. Pure silliness from a statistical point of view. Zhang et al. (2005) therefore developed a new high frequency estimator which uses all data: Two Time Scale estimation

In mathematical terms they pursue the following:

$$TTS_t = \frac{N}{N-\bar{n}} \left(\frac{1}{K} \sum_{k=0}^{K-1} RV_t^{(k)} - \frac{\bar{n}}{N} RV_t^{all} \right) \tag{11}$$

with,

$$RV_t^{(k)} = \sum r_{t,i}^2 \qquad for\ i = 1 + \frac{k}{K}, 2 + \frac{k}{K}, ..., n - 1 + \frac{k}{K}$$

$$RV_t^{all} = \sum_{j=1}^{N} r_{t,j}^2 \qquad for\ j = 1, 2, ..., N$$

$$\bar{n} = \frac{N - K + 1}{K}$$

With N the number of observed prices per day; n the number of intervals (e.g. 288 5min intervals) and $K = N/n$. That is, the two time scales estimator takes the average RV for day t over multiple subsamples and corrects the bias with $\frac{\bar{n}}{N}$ times the realized variance for day t calculated using all available data. For further details on this measure it is advised to read the appropriate section in Appendix B.

3.1.5 Realized Kernel Estimation

The last discussed and widely accepted estimator of integrated variance we test here is the Realized Kernel (RK). These type estimators, as introduced by Barndorff-Nielsen et al. (2008), are based on the ideas of Hansen and Lunde (2004). They noticed that microstructure noise causes the high frequency intra-day returns to be autocorrelated resulting in biased RV estimation. They figured that the empirical autocorrelation function up to a certain lag H can thus be used to correct the bias in RV through a correction that works in the same way as in which robust covariance estimators of Newey and West (1987) achieve their consistency. The realized kernels

are build upon the same principle. The kernel function, $K_t(Y_{t,\Delta})$, consists of RV,$\gamma_0(Y_{t,\Delta})$, which gets corrected by the empirical autocorrelations, $\sum_{h=1}^{H} k\left(\frac{h-1}{H}\right)\{\gamma_h(Y_{t,\Delta}) + \gamma_{-h}(Y_{t,\Delta})\}$, to adjust for market frictions. Formally, the kernel function of Barndorff-Nielsen et al. can thus be written as:

$$K_t(Y_{t,\Delta}) = \gamma_0(Y_{t,\Delta}) + \sum_{h=1}^{H_t} k\left(\frac{h-1}{H_t}\right)\{\gamma_h(Y_{t,\Delta}) + \gamma_{-h}(Y_{t,\Delta})\} \tag{12}$$

$$\gamma_h(Y_{t,\Delta}) = \sum_{j=1}^{K} (Y_{t,\Delta,j} - Y_{t,\Delta,j-1})(Y_{t,\Delta,j-h} - Y_{t,\Delta,j-h-1}) = \sum_{j=1}^{K} (R_{t,\Delta,j} R_{t,\Delta,j-h}) \tag{13}$$

Where H_t is the bandwidth parameter which is to be estimated and can be seen as the estimated number of autocorrelations needed to bias adjust the Kernel estimate; $k(x)$ is the chosen kernel weighing function; $\gamma_0(Y_{t,\Delta})$ the realized variance; $\gamma_h(Y_{t,\Delta})$ the h-th order autocovariance of the observed log return series; K equals the size of sparse sampling intervals: N/n; and $Y_{t,\Delta,j}$ is the observed log price level at time j in increment Δ during day t.

As weighing function the Parzen kernel is used. Given by

$$k(x) = \begin{cases} 1 - 6x^2 + 6x^3 & 0 \leq x \leq 0.5 \\ 2(1-x)^3 & 0.5 \leq x \leq 1 \\ 0 & x > 1 \end{cases} \tag{14}$$

And the preferred bandwidth equals

$$H_t^* = c^* \cdot \xi_t^{4/5} \cdot K^{3/5}, \qquad \text{with } c^* = \left\{ \frac{k"(0)^2}{\int_0^1 k(x)^2 dx} \right\}^{1/5} \text{ and } \xi_t^2 = \frac{\omega_t^2}{\sqrt{n \int_0^n \sigma_t^4 dt}} \tag{15}$$

Where an accent denotes the derivative of the function and double accents denote the second derivative, not a transpose.

A graph containing all above estimated IV measures is included in Appendix B.

3.2 Forecasting

Now that ex-post measures of volatility have been developed we can focus on actual forecasting. Here fore we explore 3 models, a seasonal moving average, a exponentially weighted moving average seasonal and a Fourier Flexible Form. Forecasting performance is subsequently benchmarked against a Random Walk and a GARCH(1,1) model as estimated on daily return data and compared to begin-of-the-day volatility as reference point.

3.2.1 Begin-of-the-day volatility (reference)

Objective of this paper is to asses the forecasting performance to daily volatility through efficient timing using begin-of-the-day volatility. A natural starting and convenient reference point is thus to test the forecasting value of begin-of-the-day volatility in itself and subsequently assess added value by more advanced models later on. If we take $RV_t = \sum_{i=1}^{n} r_{t,i}^2$ to be the end-of-the-day realized variance for day t with $i = 1, ..., n$ ranging over open outcry trade, begin-of-day variance can be written as

$$RV_t^{n^*} = \sum_{i=1}^{n^*} r_{t,i}^2 \tag{16}$$

with $1 \leq n^* \leq n$ and shifting for different subsets as $n^* \in [1, ..., n]$. Note that RV in 16 could be substituted by every other measure: BPV, RR, TTS or Kernel estimates. Furthermore, in line with research by Bedendo and Hodges (2004) on S&P500 futures contracts, we make use of the widely popular 5 minute sampling intervals for all measures.

To investigate the informativeness of $RV_t^{n^*}$ we run MZ regressions as discussed in section 3.3.1 with $RV_t^{n^*}$ as dependent variable c.q. 'forecast'. Obviously such 'forecasts' are biased and should be scaled first, yet the correlation between Realized daily volatility and $RV_t^{n^*}$ is invariant to scaling and can as such be used to review forecasting potential of the regressor. Resulting regression R^2_{MAD} coefficient thus indicates the degree in which forecasts are correlated to daily variance. Formally this correlation should in turn be adjusted acknowledging we do not actually know true daily integrated variance obliging us to use RV_t in stead of IV_t in the regression. However research to the proper adjustment within our setting is left to further research. Furthermore to make the squared correlation coefficients better comparable among different subsets of information we have to take in to account the variance already observed. Solutions to these problems will be discussed in section 3.3.

3.2.2 Direct MZ Scaling

Begin-of-the-day volatility should be scaled to daily proportions in order to form an unbiased estimate of daily volatility. To do so one could utilize the seasonal pattern as attempted in following sections. Alternatively one could track the bias as apparent from Mincer-Zarnowitz regressions and subsequently utilize these to bias adjust forecasts without the need of explicitly estimating a seasonal pattern. Using an (arbitrarily chosen) 20 day rolling window, Mincer-Zarnowitz regressions can be performed obtaining α and β estimates over the whole sample. These estimates are thus used to adjust Begin-of-the-day variance measurements to form unbiased daily variance forecasts, i.e.

$$RV^{n^*}_{t,scaled} = \alpha_{[t-20,t-1]} + \beta_{[t-20,t-1]} RV^{n^*}_t \tag{17}$$

Where α and β are estimated via OLS MZ regressions and RV can be replaced by any other volatility measure. According statistics can thereafter again be obtained via MZ regressions.

3.2.3 Seasonal Moving Average

To further improve upon our reference 'model' we follow the methods set out by Taylor and Xu (1997). They were among the first to use an average seasonal, over every interval $i = 1, ..., n$ using the previous H days, as the proposed seasonal for a coming day. The intuition behind this is that securities, especially equity, typically exhibit quite a persistent U, J or inverse J shaped volatility pattern throughout the day. The idea is that if the pattern remains persistent, begin-of-the-day volatility can be scaled to end-of-day volatility using this diurnal shape. Furthermore the shape tends to shift between U and J for respectively low and high volatility periods, see Andersen and Bollerslev (1994). By taking the average seasonal over the previous H days one can accommodate for this shifting behavior, making the method a bit more flexible. To assess the sensitivity of the seasonal to arbitrarily chosen H we take the average over 1 month i.e. 22 days and 1 year i.e. 252 days (in previous work by Frijns and Margaritis (2008) a 200 day interval was chosen). For the interval length Δ we take 5 minutes, thus ranging i from 1 up to $n = 288$. Our seasonal forecast can now be calculated as

$$s_{t,i,temp} = \frac{1}{H} \cdot \sum_{h=1}^{h=H} r^2_{t-h,i} \tag{18}$$

The seasonal is thus obtained as the average realized variance on interval i over the previous H days. Note again that multiple competing measures may be used to replace RV. However as the

seasonal is now build of only 5 minute intervals this would create a rather arbitrary and erratic diurnal signatures having low predictive value. In addition to the methods used in Frijns and Margaritis (2008) we therefore smooth the obtained open outcry seasonal using a moving average filter over respectively 30 (2.5 hours) and 10 (1 hour) timestamps for S&P500 and US30 futures data[9]. Note that the edges of the seasonal c.q. near start and end of open outcry use less than the 30 resp 10 timestamps and that the aftertrade/overnight seasonal was left as was. The smoothed seasonal can thus be written as:

$$s_{t,i,smooth} = s_{t,i} = \frac{1}{a} \sum_{i-a/2}^{i+a/2} s_{t,i,temp} \tag{19}$$

with a the filter span being an integer value. Hereby creating a flexible seasonal curve in that sense that it adjusts for shifting shapes and reduces the influence of outliers by a Moving Average Filter. To see why this makes sense take a random seasonal using (18) from S&P500 data. Figure 3 depicts the original and filtered seasonals for $a = \{10, 20, 30, 40\}$. The unfiltered series can hardly be thought persistent and are therefore less usable for forecasting purposes.

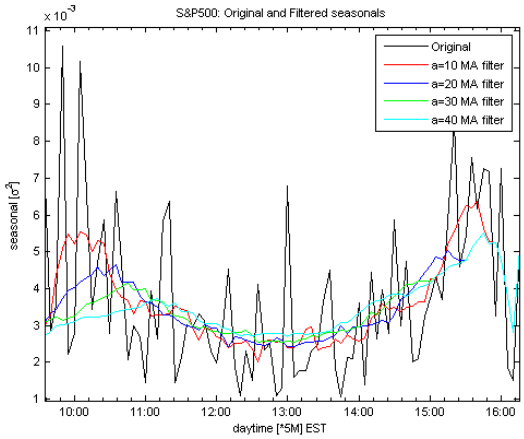

Figure 3: Original and filtered seasonals for S&P500 futures returns using a MA filter with span: 10, 20, 30 and 40 times the standard 5 minute interval length.

Division of the summation of $s_{t,i}$ over the daily intervals, $\sum_{i=1}^{n} s_{t,i}$, by the seasonal summation up to n^*, and multiplication by the readily observed start-of-the-day variance now represents the estimator of integrated volatility.

[9] As the US30 futures returns are far less erratic than is the case for S&P500 futures, these need less of a filter to smooth out undesirable spikes. 150 and 60 minute span values, 'a', are obtained as empirical optimum to a one year estimation sample. Alternatively one might try estimating parameters H and a, though such an approach might be quite hard to impossible.

$$\frac{RV_t^{n^*}}{I\tilde{V}_t} = \frac{\sum_{i=1}^{n^*} s_{t,i}}{\sum_{i=1}^{n} s_{t,i}} \quad \rightarrow \quad I\tilde{V}_t = RV_t^{n^*} \cdot \frac{\sum_{i=1}^{n} s_{t,i}}{\sum_{i=1}^{n*} s_{t,i}} \tag{20}$$

$\sum_{i=1}^{n} s_{t,i} / \sum_{i=1}^{n*} s_{t,i}$ can thus be seen as a scaling factor, scaling the begin-of-the-day variance up to daily proportions. Thus if the seasonal over the past H trading days contains valuable information, e.g. having a (slightly) different shape as opposed to the overall average shape, R^2 statistics could benefit and increase. However differences are not expected to be big as (20) is merely a rescaled version of (16).

3.2.4 Exponentially Weighted Moving Average Seasonal

In addition to simply taking the (equally weighted) moving average over the past H days, one can also define the seasonal as the exponentially weighted moving average (EWMA) over previous days. That is we weight past volatility observations by $(1-\lambda)\,\lambda^{lag}$, where λ is the decay factor, so that we can write the seasonal as

$$s_{t,i,temp} = \sum_{h=1}^{t-1} (1-\lambda)\,\lambda^{h-1} \cdot RV_{t-h,i} \tag{21}$$

By defining the seasonal this way more weight is given to recent observations, making the seasonal faster adjustable depending on the parameter setting of the decay facter, lambda. Assessment of return volatility as measured on interval i for days 1 to T over both S&P500 and US30 data, returns a parameter value lambda of approximately 0.94[10]. Furthermore to ensure the sum of weights converged to approximately one, an expanding estimation sample is taken of at least 252 days.

Final seasonal is then again obtained as $s_{t,i,smooth} = s_{t,i} = \frac{1}{a}\sum_{i-a/2}^{i+a/2} s_{t,i,temp}$. Further details for this model equal those of the seasonal moving average discussed before.

3.2.5 Fourier Flexible Form

Another method for explicit modeling of the diurnal pattern concerns fitting a functional form. This method, roughly developed by Gallant (1984), was in current form first applied to high frequency intra-day data by Andersen and Bollerslev (1997, 1998b) and later by Harju and Hussain

[10] For the whole sample of S&P500 and US30 return data, EWMA volatility forecasts were created for every interval over consecutive days. Looping over different values of lambda and maximizing the Mincer-Zarnowitz regression R^2_{MAD} lead to the optimal value of lambda: around 0.88 for S&P500 and 0.97 for US30 data. The mean 0.925 is rather close to 0.94 as suggested in RiskMetrics as developed by J.P. Morgan in the early '90s, be it for daily volatility.

(2006, 2008); Frijns and Margaritis (2008); Martens (2001) amongst others for modeling as well as filtering of high frequency data[11]. The Fourier Flexible Form (FFF) basically fits a highly flexible function through the estimated volatility data, accurately following the seasonal over the previous H days. And as daily volatility is not bound by economic or financial law or equilibrium we may find ourselves free to fit any flexible form[12]. There are a few advantages to this method: to start all data is used to form an average seasonal throughout the day, as opposed to data only within the same daily interval; second, FFF will produce a smooth function reducing the potential hazardous influence of outliers; third, the shape of the seasonal can be made dependent on other variables. The last property is especially interesting as it enables us to make the FFF dependent on start-of-the-day volatility and reviewing figure 1 and 2 first hour volatility can have quite an influence on the seasonal pattern for the rest of the day.

In accordance to Andersen and Bollerslev intra-day returns are decomposed as

$$r_{t,i} = E\left[r_{t,i}\right] + \sigma_{t,i} s_{t,i} Z_{t,i} \tag{22}$$

Additionally assuming intraday return volatility to be homoskedastic though a single day after the seasonal component has been stripped off leads to

$$r_{t,i} = E\left[r_{t,i}\right] + \frac{\sigma_t s_{t,i} Z_{t,i}}{N^{1/2}} \tag{23}$$

Where $Z_{t,i}$ is a zero mean, unit variance i.i.d. random variable assumed to be independent of $\sigma_{t,i}$; $s_{t,i}$ is a deterministic intra-day periodic component; $r_{t,i}$ is the return over interval i; $E\left[r_{t,i}\right]$ is the unconditional expectation of the return and N is the number of daily return intervals. By taking squares on both sides and subsequently applying logarithmic transformations we can rewrite (23) as

$$ln(r_{t,i}^2) = ln(E\left[r_{t,i}\right]^2) + ln\left(\frac{\sigma_t^2 s_{t,i}^2 Z_{t,i}^2}{N}\right) \qquad \rightarrow$$

[11] Apart from modeling daily volatility using Fourier Flexible Forms one can also use them to filter the diurnal pattern from intra-day data. As so, one can further investigate the dynamics properties driving daily volatility. Such a filtering approach works quite well as can be seen from A.3.

[12] Research brought forth a whole line of different flexible forms. Problems like choosing a suitable FFF have again a whole line of literature amongst themselves, see Thompson (1988); Muller et al. (1990); Dacorogna et al. (1993). Alternatively to FFF one can for instance use Cubic- or quadratic-Spline functions and quite some other Taylor, Laurent or Fourier inspired flexible forms to obtain a seasonal pattern.

$$x_{t,i} \equiv 2 \cdot ln\left(\left|r_{t,i} - E\left[r_{t,i}\right]\right|\right) - ln\left(\sigma_t^2\right) + ln(N) - E[ln(Z_{t,i}^2)] = ln\left(s_{t,i}^2\right) + ln\left(Z_{t,i}^2\right) - E[ln(Z_{t,i}^2)] \quad (24)$$

Modeling of the functional form is then based on ordinary least squares (OLS) regression of $x_{t,i} = f(\theta_t; \sigma_t, i) + \varepsilon_{t,i}$, where $\varepsilon_{t,i} \equiv ln(Z_{t,i}^2) - E[ln(Z_{t,i}^2)]$, and $f(\theta_t; \sigma_t, i)$ takes the parametrized form[13]:

$$ln\left(s_{t,i}^2\right) = f(\theta_t; \sigma_t, i) = \sum_{j=0}^{J} \sigma_t^j \cdot \left[\mu_{0,j} + \mu_{1,j}\frac{i}{N_1} + \mu_{2,j}\frac{i^2}{N_2} + \sum_{d=1}^{D} \lambda_{d,j} I_{i=D_d} + \sum_{p=1}^{P}\left(\gamma_{p,j} \cdot cos\left(\frac{pi2\pi}{N}\right) + \delta_{p,j} \cdot sin\left(\frac{pi2\pi}{N}\right)\right)\right] \quad (25)$$

The first part of this functional form is used to catch linear and quadratic terms that may partly shape the daily seasonal pattern. The terms N_1 and N_2 herein are normalizing constants defined as $N_1 = (N+1)/2$ and $N_2 = (N+1) \cdot (N+2)/6$. Furthermore some dummy variables, I_i ,are taken into account to accommodate sudden jumps in the seasonal shape, mostly located near the beginning and the end of the day. The second part of the form fits a certain number of sinusoids to further imitate the diurnal pattern. Last, if $J > 0$ the whole form is multiplied by the standard deviation, σ_t^j. Making the overall shape a function of start-of-the-day volatility. This last property could be of critical value since Andersen and Bollerslev (1994) found different diurnal shapes specific to high and low volatility periods and figure 1 and 2 lead to the same conclusion.

Practical estimation is best performed by first generating the series $x_{t,i}$ as defined in equation (24). $E\left[r_{t,i}\right]$ can therefore be replace by the mean return, $\bar{r}_t$, and σ_t^2 by its estimated counterpart $\hat{\sigma}_t^2$ as obtained from daily realized variances or competing measures. Subsequently treating $\hat{x}_{t,i}$ as the dependent variable enables us to estimate the FFF coefficients via ordinary least squares regression[14]. With regard the coefficients J and P, optimal values are searched for by means of a grid leading to the lowest Schwartz Information Criterion (SIC) or Akaike information Criterion

[13]Note that the breakdown of $x_{t,i}$ contains the term $-E[ln(Z_{t,i}^2)]$ in addition to the regular form of Andersen and Bollerslev. This additional term is needed to make the OLS residual term mean zero without adding this term to the constant from seasonal parametrization. Omitting this term would thus lead to biased estimates of the seasonal pattern. A simulation of the seasonal can be found in Appendix B: Seasonal simulation experiment and confirms this.

[14]Note that the Fourier Flexible Form is originally designed for fitting average volatility patterns across n intra-day intervals. OLS estimation on multiple days can however be conducted. Utilizing the additional information from previous H days enhances the efficiency of the estimation.

(AIC)[15]. These measures lay different weights on the inclusion of additional parameters and as such they give quite different optimal values. During this research we test for the informativeness of start-of-the-day volatility, therefore it seems sensible to take at least $J = 1$. For the P parameter SIC is followed as AIC seems to have the tendency to take in too many parameters in comparison to the number of observations, n. Concluding, we set $J = 1$ and $P = 1$ for S&P500 and $J = 1$ and $P = 2$ for US30 data. During the discussion slightly different parameter settings are considered to gain some insight in the sensitivity to this choice[16]. Last parameter, D, relates to the number of dummy variables which can accommodate abrupt changes, e.g. near the opening and closing of the market. D is set to 3 for S&P500 and 4 for US30, one dummy at the start of the day, one for the end, one to accommodate the consistent presence of 10:00 AM EST Macroeconomic news announcements and one additional fourth dummy to accommodate 08:30 AM EST news announcements for US30 series only.

Once $\hat{f}(\hat{\theta}_t; \sigma_t, i)$ is obtained, the intra-day seasonal volatility can be retrieved as: $\hat{s}_{t,i} = exp\left(\hat{f}(\hat{\theta}_t; \sigma_t, i)/2\right)$. The seasonal forecast than equals $\hat{s}_{t+1,i} = exp\left(\hat{f}(\hat{\theta}_t; \sigma^{n^*}_{t+1}, i)/2\right)$, where $\sigma^{n^*}_{t+1}$ represents the start-of-the-day volatility up until n^*. Now that the seasonal pattern for day $t+1$ has been predicted, we proceed in the same manner as before. That is, equation (20) is followed to obtain the final forecast and (28) leads to the desired statistics[17]. Note that in order to fit the seasonal pattern to absolute returns one has to multiply by the daily volatility and divide by the square root of the number of intra-day intervals, i.e. $\sigma_t \cdot \hat{s}_{t,i}/\sqrt{N}$. To illustrate this fit on the average absolute return series for S&P 500 index future returns and US 30 year treasury bonds future returns, two graphs are included in figure 4. The fit on the average seasonal patterns are remarkably well. FFF therefore seems an appropriate method to estimate diurnal patterns. However it must be held in mind that the fit shown here is over the *average* diurnal pattern. The daily volatility pattern as developed through above mentioned measures is far less stylized c.q. much more erratic.

[15]The SIC, aka BIC, weights the model fit against the amount of parameters to be estimated using the following formula: $SIC(p) = ln(\hat{\sigma}^2_p) + \frac{p \cdot ln(n)}{n}$, where p is the number of explanatory variables, $\hat{\sigma}^2_p$ is the Maximum Likelihood estimator of the error variance and n is the number of daily observations. AIC takes the form $AIC(p) = ln(\hat{\sigma}^2_p) + \frac{2p}{n}$. Note that for SIC and AIC to be comparable among models they have to be estimated against the same dependent variable, as is the case here.

[16]See Appendix B, figure B.3 and B.4 for some graphical fits and SIC/AIC statistics on mean absolute daily volatility for J ranging from 0,...,2 and P ranging from 1,...,9.

[17]Note that the flexible Fourier Form will only be fitted to the periods over which we have defined end-of-the-day volatility, i.e. 09:30 until 16:15 ET for S&P 500 index and 8:20 up until 17:00 for US 30 year treasury bonds.

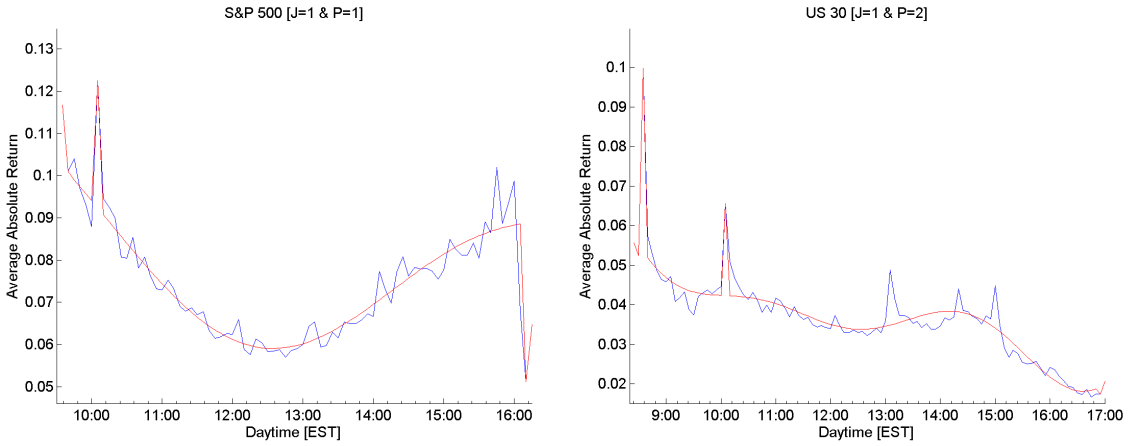

(a) FFF fit to mean absolute S&P500 returns **(b)** FFF fit to mean absolute US30 returns

Figure 4: Fourier flexible form fit on absolute S&P 500 index and US 30 year treasury bond returns data using respectively J=1, P=1 and J=1, P=2. The red line are the fitted values and the blue line the average observed absolute return over 5 minute intervals during open outcry trading times. The S&P 500 data depicts an apparent diurnal pattern whereas this is less pronounced in US 30 data.

3.2.6 Random Walk (Benchmark 1)

Most basic yet sometimes hard to beat 'model' is the random walk (RW). This method assumes non predictable future states whereby the best forecast is simply given by the current value. The random walk is created by taking end-of-the-day variance as obtained by realized variance or competing measures based of 5 minute intra-day returns. Stated in a formula RW is given by:

$$\sigma^2_{RW,t+1} = \sigma^2_{RW,t} + \varepsilon_t = RV_t + \varepsilon_t \text{ so that}$$

$$\hat{\sigma}^2_{RW,t+1} = RV_t \tag{26}$$

3.2.7 GARCH(1,1) model (Benchmark 2)

It might strike that we do not use Generalized Autoregressive Conditional Heteroskedasticity (GARCH) or other ARCH type models estimated on intra-day data to benchmark our forecasting performance. The reason is that standard ARCH models assume exponential decay in the return autocorrelation function making them fundamentally unable to cope with the strong daily cyclical behavior as apparent from figure A.3. Supplemented with a seasonal model to capture the intraday periodicity, however, GARCH could be of value. Such a model has earlier been proposed by

Andersen and Bollerslev (1997). Here we use a GARCH(1,1) model estimated on daily return data to benchmark forecasting performance. Furthermore we assess whether GARCH forecasts make a significant addition to the FFF forecasts described below.

For readers unfamiliar with the GARCH(1,1) model, this model it is expressed as:

$$r_t = \mu + \varepsilon_t \quad \text{with } \varepsilon_t \sim N(0, \sigma_t^2) \text{ and}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 \sigma_{t-1}^2 \tag{27}$$

That is, the estimated variance for day t is equal to a constant plus α_1 times the squared previous return innovation and α_2 times the previous variance. By modeling volatility this way it is made dependent on time so that it can account for conditional heteroskedasticity, the property that return series experience periods of high and low volatility. Subsequently it is made dependent on its own past, the autoregressive property.

3.3 Statistical evaluation

Statistical evaluation of the above stated models over different times of the day will be conducted by Mincer-Zarnowitz type regressions. As the regression coefficients will play a significant role in the sections to come and especially later evaluation, some additional explanation and intuitive visualization is included in Appendix B:Mincer-Zarnowitz regression coefficients, for the non familiar reader.

3.3.1 Mincer-Zarnowitz regressions

To assess the forecasting performance of above stated models, Mincer-Zarnowitz type regressions are used. Noting however that true integrated volatility is unobserved normal MZ regressions are infeasible. Using an unbiased estimator of σ_t, the feasible alternative $RV_t = \alpha + \beta \cdot \hat{\sigma}_t^2 + \varepsilon_t$ normally still yields unbiased and consistent estimates of α and β, protecting the validity of the test. Estimates will however become less accurate as the variance of $(\sigma_t^2 - RV_t)$ grows, losing power to detect deviations from optimality. As the residuals from the MZ regressions will in general be heteroskedastic, Patton and Sheppard (2008) argue that size and finite sample power properties can be improved using Generalized Least Squares (GLS), resulting in smaller standard errors of

the estimates and guarding against heteroskedasticity. That is, we use the regression[18]

$$\frac{RV_t}{\hat{\sigma}_t^2} = \frac{\alpha}{\hat{\sigma}_t^2} + \beta + \varepsilon_t \tag{28}$$

Great advantage to this method lies in the lower weights attributed to extreme observations, lowering the influence of outliers on the regression coefficients α and β. This however comes at the cost of lower R^2. To see this note that OLS minimizes the total sum of squares, $SS_{err} = \sum \left(RV_t - \left(\alpha + \beta \cdot \hat{\sigma}_t^2\right)\right)^2$, and thereby maximizes R^2 as defined by $R^2 = 1 - {}^{\text{SS}_{\text{err}}}/_{\text{SS}_{\text{tot}}}$. In fact, in the GLS estimation procedure, R^2 is not well defined, see Buse (1973); Battese and Griffiths (1980); Blomquist (1980). That is, under the use of GLS coefficients the SS_{err} can become greater that SS_{tot}, leading to R^2 values in the range of $[-inf, 1]$. Clearly losing its appealing interpretation as the percentage of variance explained. Despite the common use of GLS standard textbooks do not ellaborate on the goodness of fit in such settings, nor has research delivered an unambiguous procedure.

To nonetheless obtain a goodness of fit measure in the interval $[0, 1]$ we note that under the OLS setting R^2 also equals $corr\left(RV_t, \hat{\sigma}_t^2\right)^2$. Which forms a more direct and intuitive goodness of fit measure, invariant to OLS, WLS or GLS setting. Using a squared robust correlation statistics as summed by Lee et al. (2006); Shevlyakov and Smirnov (2011) could therefore yield a solution. Through simulation studies they find that under a contaminated normal distribution r_{MAD} approximates the underlying correlation best, yielding least bias. r_{MAD}^2 is therefore used as to compare forecasting statistics. To formalize:

$$r_{MAD} = \frac{MAD^2(u) - MAD^2(v)}{MAD^2(u) + MAD^2(v)}$$

$$u = \frac{x - med(x)}{\sqrt{2}MAD(x)} + \frac{y - med(y)}{\sqrt{2}MAD(y)} \text{ and } v = \frac{x - med(x)}{\sqrt{2}MAD(x)} - \frac{y - med(y)}{\sqrt{2}MAD(y)}$$

$$MAD(x) = med\{|x_1 - med(x)|, ..., |x_n - med(x)|\}$$

Whether such a squared robust correlation coefficient still caries the same interpretation in terms of explained variance as its normal counterpart is not further examined here and might be an

[18]The use of squared returns in MZ regressions has caused concern among researchers as statistical inference relies on fourth powers of the returns. Greatly magnifying the impact of large return observations. Often proposed alternative lies in the use of transformed series, e.g. $|r_t| = \alpha + \beta\sqrt{\hat{\sigma}_t^2} + \varepsilon_t$ or $ln\left(r_t^2\right) = \alpha + \beta ln\left(\hat{\sigma}_t^2\right) + \varepsilon_t$. Patton and Sheppard (2008) however ague that such regressions can result is size distortions.

interesting topic for further research.

As mentioned above, the normal MZ regression is infeasible due to volatility being an unobserved variable. Acknowledging this, the R^2 of the feasible MZ regressions should be adjusted accordingly. Using correlation statistics instead to infer R^2, adjustments are unsure. The normal adjustments (details on this can be found in Appendix B: R^2) are therefore omitted and the quest for an appropriate adjustment in this conjuncture is left to future research. As the MZ model is wrongly specified this leads to artificially small standard errors, in turn leading to over-rejection of $H_0 : \alpha = 0, \beta = 1$. Such adjustments could therefore benefit this research, leading to keener conclusions.

3.3.2 Variance Ratio, HMSPE and Marginal R^2

To further asses the performance, we additionally introduce the Variance Ratio (VR), HMSPE (Hetroskedasticity-consistent Mean Squared Prediction Error) and what we shall call the marginal R^2: R^2_{marg}. These additional features make the squared correlation somewhat better comparable along estimation time i. The HMSPE is a hetroskedasticity consistent equivalent of the MSPE and is preferred here as especially S&P500 data contains heteroskedastic periods

$$HMSPE = \frac{1}{T}\sum_{t=0}^{T-1}\left(1 - \frac{RV_{t+1}}{\tilde{IV}_{t+1}}\right) \tag{29}$$

The variance ratio is defined as the average proportion of daily variance already observed up until the point of estimation,

$$VR^{n^*} = \frac{1}{T}\sum_{t=1}^{T}\frac{RV_t^*}{RV_t}, \tag{30}$$

This is an important additional feature since the regression outcome will on average only become better as more of the volatility reveals itself. Ultimately reaching 1 as all realized variance is observed. Offsetting these results against the variance ratio gives better insight in true performance. Take for example a R^2 as high as 0.8 after just 1 hour of trading, yet if 90% of daily variance is already observed by that time such results are not at all impressive.

A second measure is found in the combination of the squared correlation (henceforth abbreviated as: R^2_{MAD}) and variance ratio. If an estimation sample begins at the start of the day and continues up until the end it will most certainly have a R^2_{MAD} of approximately 1. This is clearly not the most efficient forecast en should therefore be penalized. Although mathematically disputable the Variance Ratio could act as such a penalty. That is by subtracting the percentage

of observed volatility from the R^2_{MAD} a single measures, R^2_{marg}, is again obtained giving some extra insight in the marginal effects of extra observations on R^2_{MAD}. In fact measuring efficiency as the increase in information should at least lead to an equal percentage increase of R^2_{MAD}[19].

3.3.3 Mincer-Zarnowitz alpha, beta and errors terms

In addition to R^2_{MAD}, R^2_{marg} and VR statistics it is of interest how Mincer-Zarnowitz parameters $\hat{\alpha}$ and $\hat{\beta}$ evolve as a function of both t and n^*. That is, how they evolve over time as high or low volatility periods often associated with recessions and expansions respectively pass by. Giving inside in the performance during these different regimes. And, how the parameters behave through the day yielding insights potentially enabling the creation of better, i.e. unbiased, forecast as their evolution might be predictable. One certainty is beforehand obvious, as $n^* \to n$ parameter values of $\hat{\alpha}$ and $\hat{\beta}$ inevitably go to their boundaries of respectively 0 and 1. The properties of ε_t in (31), especially the variance thereof, is also closely related to the development of n^*. Diminishing as the forecasts include more of the days volatility, eventually reaching zero for $n^* = n$. First glance at these features can be obtained by OLS estimation of α and β for all n^*.

$$RV_t = \alpha_{n^*} + \beta_{n^*} \cdot \tilde{IV}_t + \varepsilon_t \tag{31}$$

Furthermore relaxing the condition of fixed parameters over time, i.e. replacing α_{n^*} and β_{n^*} by $\alpha_{n^*,t}$ and $\beta_{n^*,t}$ respectively, these parameters can be approximated by OLS regression using a rolling window around the timestamps. Though taking to long a window could blunt interesting developments due to for instance crisis periods, it could also result in low accuracy due to the influence of outliers. Too short a window on the other hand could also induce high estimation insecurity on the accurateness of coefficient estimates. An optimum can likely be found but since this is not a road we embark to follow, instead we treat $\alpha_{n^*,t}$ and $\beta_{n^*,t}$ as unobserved time series and reveal them using Kalman Filtering, a somewhat arbitrary 20 day rolling window is chosen for illustrative purposes.

Some illustrative results are depicted in figure 5 and 6. Figure 5 shows the evolution of α_{n^*} and β_{n^*} over the time of day accompanied by their 95% confidence intervals. It stands out that the MZ regression coefficients behave in quite a stylized way as $n^* \to n$, that is they seem to follow

[19]The observation of additional volatility information on the current day must especially be held in mind comparing forecast performance to a GARCH(1,1) benchmark, as this model has no such additional information. To adjust for this the following (robust) correlation could be used: $corr\left(RV_t - RV_t^{n^*}, \tilde{IV}_t\right)$. However as the main target here is to see whether this additional information improves over conventional estimation methods, adjustments need not be made.

some sort of exponential from the beginning at start of open outcry trade towards their boundaries of 0 and 1 at the end of the day. Figure 6 shows the evolution of $\alpha_{n^*=30min,t}$ and $\beta_{n^*=30min,t}$ over time using 20 day rolling window OLS regressions, taking RV as volatility measure, fixing the daytime at 30 minutes after start of open outcry trade and using the Seasonal Moving Average with H=22 days to forecast the intraday variance pattern. It can be seen from this figure that the evolution of α is highly correlated to the stock price volatility (0.6785 and 0.4235 for resp. S&P500 and US30). Meaning forecasting performance decreases during more volatile periods, as can be expected. The effects on beta (correlations of resp. 0.2051 and -0.3384) are somewhat less pronounced.

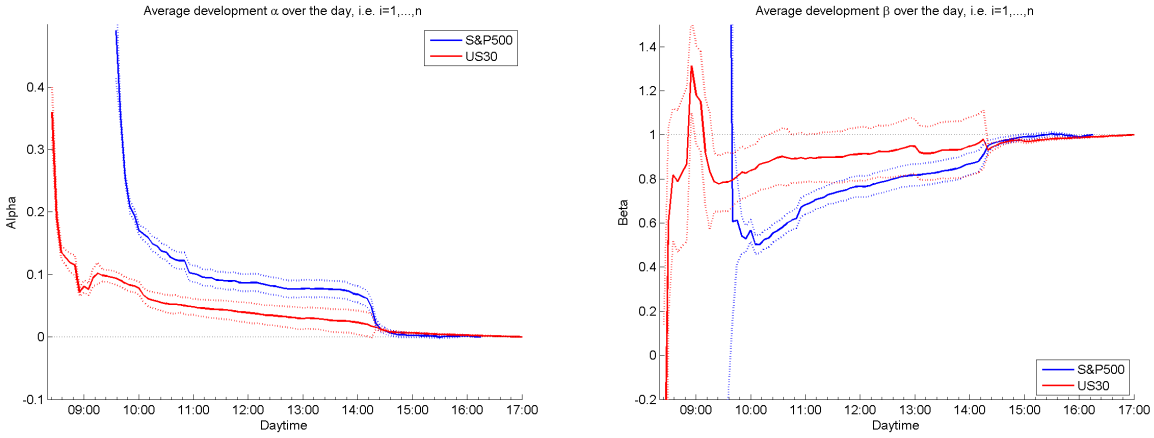

(a) MZ $\alpha\,(n^*)$ **(b)** MZ $b\,(n^*)$

Figure 5: Figures depict the evolution of separately GLS-estimated Mincer-Zarnowitz regression coefficients combined with their 95% confidence intervals. Separately, mend in the sense that GLS regressions are performed for every 5 minute interval i in n over time $t = 1, ..., T$.

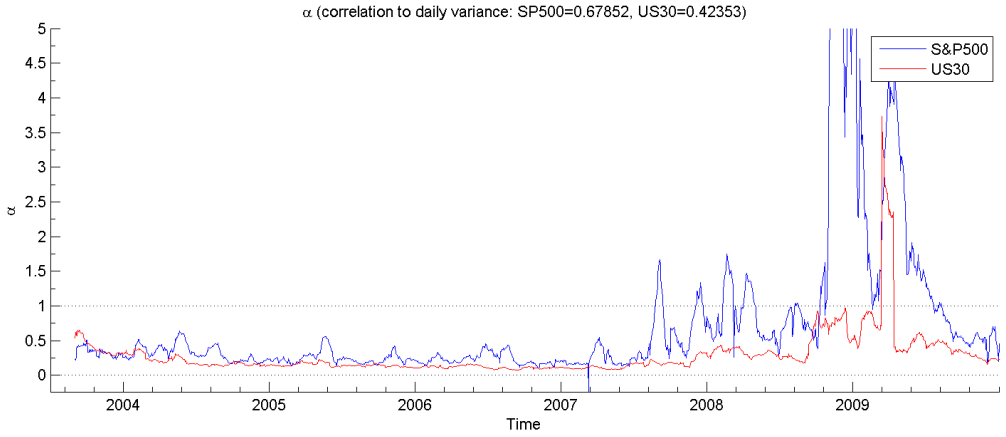

(a) MZ $\alpha\,(n^*)$

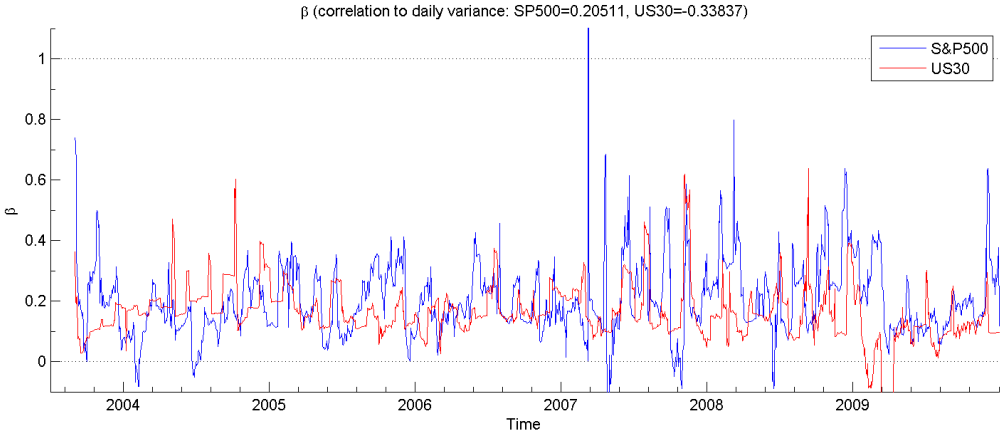

(b) MZ $b\,(n^*)$

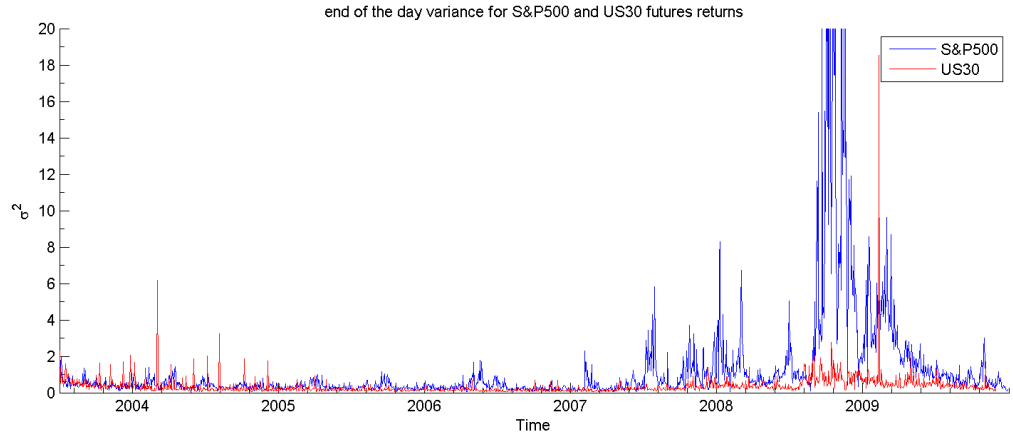

(c) Daily variance, σ^2

Figure 6: Figures depict the evolution of MZ α and β over time for $i = 6$ (30 minutes after opening) with H=22 and their correlation to the daily volatility. MZ α en β estimates are obtained from moving window OLS regressions over 20 subsequent days.

Moreover these features: the stylized form and evolution, seem robust over time and through different volatility measures, forecasting methods, estimation sample (H) and security type making

them susceptible to modeling. This is the road we will follow next. A first attempt to model the unobserved states of alpha and beta over time, $t = 1, ..., T$, is made. This is done by modeling the states as multivariate AR(1) models for $n^* = 1, ..., n$, yielding the so called seemingly unrelated time series equation (SUTSE) model:

$$\alpha_{t,n^*} = c_{\alpha,n^*} + \varphi_{1,n^*}\alpha_{t-1,n^*} + \zeta_{t,n^*} \text{ for } n^* = 1, ..., n \text{with } \zeta_t \sim N(0, Q_\zeta) \tag{32}$$

$$\beta_{t,n^*} = c_{\beta,n^*} + \varphi_{2,n^*}\beta_{t-1,n^*} + \eta_{t,n^*} \text{ for } n^* = 1, ..., n \text{with } \eta_t \sim N(0, Q_\eta) \tag{33}$$

Which can be written more conveniently in state space notation as

$$x_t = C + Fx_{t-1} + \omega_t \text{ with } \omega_t \sim N \begin{pmatrix} 0, & Q_\zeta & 0 \\ 0, & 0 & Q_\eta \end{pmatrix} \tag{34}$$

With x_t a $[2n \times 1]$ vector of unobserved states $\{\alpha_{t,n^*}, \beta_{t,n^*}\}'$; C a $[2n \times 1]$ vector of constants; F a $[2n \times 2n]$ diagonal matrix of slope coefficients and ω a $[2n \times 1]$ vector of mean zero, normally distributed i.i.d. error terms with blockdiagonal covariance matrix. Note that the assumption of blockdiagonal Q is not required within the state space c.q. Kalman filter setting. It is merely taken as to forestall the curse-of-dimensionality by lowering the number of parameters to be estimated. In fact one could expect high volatility periods to create larger, more erratic alphas as indeed observed in figure 6 sub a). Countered by lower betas as the correlation between forecast en predicted value drops during highly volatile periods. Consequently resulting in negative correlation between unobserved states α and β.

For the retrieval of unobserved states α and β Kalman filtering techniques are conducted. The measurement equation for this purpose is defined by:

$$RV_t = \alpha_{t,n^*} + \beta_{t,n^*} \cdot \tilde{IV}_{t,n^*} + v_{t,n^*} \text{ for } n^* = 1, ..., n \tag{35}$$

$$\text{with } \tilde{IV}_{t,n^*} = RV_t^{n^*} \frac{\sum_{i=1}^{n} s_{t,i}}{\sum_{i=1}^{n^*} s_{t,i}} \text{ and } v_t \sim N(0, R)$$

Which again put in state space form becomes

$$z_t = H_t x_t + v_t \text{ with } v_t \sim N(0, R) \tag{36}$$

With z_t a $[n \times 1]$ vector of end-of-the-day variance measurements; H_t a $[n \times 2n]$ transition matrix containing a diagonal matrix of ones to accommodate the constant term and a second additive diagonal matrix containing $I\tilde{V}_t$, i.e. $\begin{bmatrix} I_n & I\tilde{V}_t \end{bmatrix}$; x_t a $[2n \times 1]$ state vector $[\alpha_{t,1}, ..., \alpha_{t,n}, \beta_{t,1}, ..., \beta_{t,n}]'$; v_t a $[n \times 1]$ vector of mean zero, normally distributed error terms with $[n \times n]$covariance matrix R. For the interested reader initial values, intuitive examples, explanations to and details on Kalman Filtering and smoothing can be found in Appendix B: The Kalman Filter.

Expectation Maximization (EM)

The standard Kalman filter as discussed above assumes transition matrix F, constants vector C, initial states distribution, $x_0 \sim N(\mu_0, \Sigma_0)$ and the state and observations covariance matrix, respectively Q and R, to be known. Originating from physics this is not an odd thing as these matrices can, it that case, often be derived with quite some precision by differential equations or repeated experiments, e.g. on the sensors accuracy. In economics none such recurring real time experiments can be set up (apart from simulation) and no such a priori knowledge can thus be obtained. Therefore one would like to estimate the quantities $\theta = \{\mu_0, \Sigma_0, C, F, Q, R\}$. There are a few possibilities to do such a thing, yet first we need a data likelihood function before an estimation algorithm can take over.

Given the state space model, the data likelihood can be written as: $L(Z) = \prod_{t=1}^{T} p(z_t|Z_{t-1})$ with $Z_t = \{z_1, z_2, ..., z_t\}$ or by log transformation as the sum: $lnL(Z) = \sum_{t=1}^{T} p(z_t|Z_{t-1})$. As v_t is assumed $N(0, R)$, the measurement equation has multivariate normal distribution with conditional mean $E[z_t|Z_{t-1}] = H_t x_t$ and conditional variance $var(z_t|Z_{t-1}) = H_t P_{t|t-1} H_t' + R = S_t$. Now noting that the multivariate normal probability density function is given by

$$2\pi^{-n/2} |S_t|^{-1/2} exp\left(-\frac{1}{2}\left(z_t - H_t x_{t|t-1}\right)' S_t^{-1}\left(z_t - H_t x_{t|t-1}\right)\right) \tag{37}$$

The data likelihood can be written as:

$$lnL(Z_t;\theta) = \sum_{t=1}^{T} ln\left(2\pi^{-n/2}\right) + \sum_{t=1}^{T} ln\left(|S_t|^{-1/2}\right) + \sum_{t=1}^{T} -\frac{1}{2}\left(z_t - H_t x_{t|t-1}\right)' S_t^{-1}\left(z_t - H_t x_{t|t-1}\right)$$

$$lnL(Z_t;\theta) = -\frac{n}{2}\sum_{t=1}^{T} ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T} ln(|S_t|) - \frac{1}{2}\sum_{t=1}^{T} e_t' S_t^{-1} e_t \text{ with } e_t = z_t - H_t x_{t|t-1}$$

Ignoring the constant term as this has no effect on the maximization and multiplying by -2 we get

$$-2lnL\left(Z_t;\theta\right)=\sum_{t=1}^{T}ln\left(|S_t|\right)+\sum_{t=1}^{T}e_t'S_t^{-1}e_t \tag{38}$$

Yielding the prediction error decomposition as first written down by Schweppe (1965). Equation (38) can now be minimized to obtain estimates for θ. Here we discuss two possible options to tackle such a problem.

First, as the likelihood is often a highly complicated possibly nonlinear function of the unknown parameters, θ, estimates are often obtained by iterative (numerical) procedures based on Newton-Raphson type optimization algorithms to minimize the likelihood in (38). See for example the Newton-Raphson, Gauss-Newton (scoring), Davidson-Fletcher-Powell methods or the numerical quasi-Newton procedure of Broyden-Fletcher-Goldfarb-Shanno (BFGS) and the Berndt-Hall-Hall-Hausman (BHHH) method. Generally the procedure involves the following steps:

1. Select initial values say $\theta_{(0)}$.
2. Run the Kalman Filter to get (new) states using the initial parameter values, $\theta_{(0)}$. Obtain a set of innovation e_t and error covariances S_t
3. Run the Newton-Raphson procedure to obtain a new set of parameters $\theta_{(1)}$
4. Repeat step 2 using the new parameters, $\theta_{(j)}$. Repeat step 3 to get $\theta_{(j+1)}$ and keep iterating until the likelihood has converged or the difference decreased to a predetermined small amount, that is $-2lnL\left(Z_t;\theta_{(j)}\right)\approx-2lnL\left(Z_t;\theta_{(j+1)}\right)$

This procedure, advocated by for instance Jones (1980), does however have some drawbacks. First it needs the gradient matrix (Jacobian) of first order partial derivatives to determine the direction of a step and the Hessian to modify the step size. In practice it is often infeasible or impossible to obtain such expression analytically and one is therefore designated to rely on feasible numerical evaluation methods. Second drawback lies in the fact that sequential iteration not necessarily improve the likelihood. Advantage on the other hand is that the algorithm is generally fast and standard errors of the estimates can be obtained as a byproduct of the estimation procedure.

Second, Shumway and Stoffer (1982) showed how the conceptually simpler expectation maximization algorithm of Dempster et al. (1977) could be used for parameter estimation. Basically the algorithm consists of Two alternating steps, an Expectation or E-step where one calculates the

conditional expectation of the complete data likelihood measured by the so called Q-function and the Maximization or M-step where new parameters θ are collectively estimated by maximizing this conditional likelihood.

However before we jump into the specifics, let again start from the data likelihood. Their basic idea was that if we can observe states, $X_t = x_1, x_2, ..., x_t$, and measurements, $Z_t = z_1, z_2, ..., z_t$, than we could consider these the complete data $\{X_t, Z_t\}$ with the joint density[20]:

$$f_\theta(X_t, Z_t) = f_{\mu_0,\Sigma_0}(x_0) \cdot \prod_{t=1}^{T} f_{C,F,Q}(x_t|x_{t-1}) \prod_{t=1}^{T} f_R(z_t|x_t)$$

Under the assumption of Gaussian error terms for the initial states and independent and multivariate normally distributed error terms v_t and ω_t, that is $\underset{[2nx2n]}{x_0} \sim N(\mu_0, \Sigma_0)$, $\underset{[2nx2n]}{x_t} \sim N(C + Fx_{t-1}, Q)$, $\underset{[nxn]}{z_t} \sim N(Hx_t, R)$, the complete data likelihood, $lnL(X, Z; \theta)$, may be written a

$$\begin{aligned} lnL(X, Z; \theta) = ln\left(2\pi^{-2n/2}\right) + ln\left(|\Sigma_0|^{-1/2}\right) - \frac{1}{2}(x_0 - \mu_0)' \Sigma_0^{-1}(x_0 - \mu_0) + \\ \sum_{t=1}^{T} ln\left(2\pi^{-2n/2}\right) + \sum_{t=1}^{T} ln\left(|Q|^{-1/2}\right) + \sum_{t=1}^{T}\left(-\frac{1}{2}(x_t - C - Fx_{t-1})' Q^{-1}(x_t - C - Fx_{t-1})\right) + \\ \sum_{t=1}^{T} ln\left(2\pi^{-n/2}\right) + \sum_{t=1}^{T} ln\left(|R|^{-1/2}\right) + \sum_{t=1}^{T}\left(-\frac{1}{2}(z_t - H_t x_t)' R^{-1}(z_t - H_t x_t)\right) \end{aligned} \quad (39)$$

Note that due to the determinant in 39 covariance matrices Σ_0 and Q had to be split in two diagonal blocks, corresponding to the α and β part, in order to correctly calculate the likelihood. For the latter matrix multiplication no such adjustment has to be made. Further, again ignoring the constant term and multiplying by a factor -2 we get:

$$\begin{aligned} -2lnL(X, Z; \theta) = ln(|\Sigma_0|) + (x_0 - \mu_0)' \Sigma_0^{-1}(x_0 - \mu_0) + \\ T \cdot ln(|Q|) + \sum_{t=1}^{T}(x_t - C - Fx_{t-1})' Q^{-1}(x_t - C - Fx_{t-1}) + \\ T \cdot ln(|R|) + \sum_{t=1}^{T}(z_t - H_t x_t)' R^{-1}(z_t - H_t x_t) \end{aligned} \quad (40)$$

Where we emphasize that the quantities in this equation are dependent on current parameters

[20] x_0, x_t and z_t are assumed mutually independent of each other and for all t. Their joint density may therefore be written as the product of their marginal densities.

θ. Next step is to minimize (40) conditional on all measurements and parameter estimates form last iteration. This can be done by the Kalman Smoother and One-Lag Covariance smoother equations, formulas here fore can be found in Appendix B: table B.3. This smoothing basically consists of a backward recursion through the Kalman filter, estimating the optimal states, in the sense of mean squared error (MSE), having already observed all the data, i.e. $E[x_t|Z_T]$. Consequently we get:

$$argmin: Q_{function}\left(\theta^j|Z_T,\theta^{j-1}\right) = E\left[-2lnL\left(X,Z;\theta\right)|Z_T,\theta^{j-1}\right]$$

$$Q_{function}\left(\theta^j|Z_T,\theta^{j-1}\right) = ln\left(|\Sigma_0|\right) + tr\left(\Sigma_0^{-1}\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)' + P_{0|T}\right]\right) +$$
$$T\cdot ln\left(|Q|\right) + tr\left(Q^{-1}\left[S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right]\right) +$$
$$T\cdot ln\left(|R|\right) + tr\left(R^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t - H_tx_{t|T}\right)\left(z_t - H_tx_{t|T}\right)' + H_tP_{t|T}H_t'\right)\right]\right) \quad (41)$$

with

$$S_0 = \sum_{t=1}^{T}\left(x_{t-1|T}\right)$$

$$S_1 = \sum_{t=1}^{T}\left(x_{t|T}\right)$$

$$S_{00} = \sum_{t=1}^{T}\left(x_{t-1|T}x_{t-1|T}' + P_{t-1}^T\right)$$

$$S_{10} = \sum_{t=1}^{T}\left(x_{t|T}x_{t-1|T}' + P_{t,t-1|T}\right)$$

$$S_{11} = \sum_{t=1}^{T}\left(x_{t|T}x_{t|T}' + P_{t|T}\right)$$

As we are now conditioning on Z_T, already knowing all measurements, we thus take the smoothed state estimates from the Kalman Smoother. $x_{t|T}$ denotes these smoothed state estimate, i.e. $E[x_t|z_1,...,z_T]$; $P_{t|T}$ equals the smoothed state covariance, $P_{t|T} = cov\left(x_t|z_1,...,z_T\right)$ and $P_{t,t-1|T} = cov\left(x_t,x_{t-1}|z_1,...,z_T\right)$ is the lag-one smoothed covariance. For the transformation we made use

of the matrix trace properties 'tr': $x'Ax = tr\left(x'Ax\right) = tr\left(Axx'\right)$ in which the first equality holds as long as x is a vector or scalar and $\sum tr\left(Axx'\right) = tr\sum\left(Axx'\right)$. Furthermore note that through the definition of the covariance, $E\left[xx'\right] = cov(x, x') + E\left[x\right]E\left[x\right]'$, holds. Translated to the situation at hand this clarifies where the extra covariance term comes from: $E\left[\sum_{t=1}^{T}(x_t x_t')|Z_T;\theta\right] = \sum_{t=1}^{T}E\left[x_t x_t'|Z_T;\theta\right] = \sum_{t=1}^{T}\left(P_{t|T} + x_{t|T}x_{t|T}'\right)$. More thorough derivations to (41) can be found in Appendix B: Unconstrained EM M-step derivations.

Calculation of (41) comprehends the Expectation or E-step from the EM algorithm. The Maximization step or M-step is used to obtain new optimal estimates by minimizing (41) with respect to the parameters at iteration j. If all data, i.e. the measurements and states $\{z_t, x_t\}$, were known and an error distribution was assumed, standard Maximum Likelihood Estimates (MLE) could be obtained through multivariate normal theory. Yet as we do not have all data EM gives us a iterative procedure to obtain MLE of θ based on the incomplete data, z_t, by successively maximizing the conditional expectation of the complete data likelihood.

To obtain the M-step estimates, additive parts of the Q-function can be optimized separately, yielding the updated parameters $\theta^j = \{\mu_0, \Sigma_0, C, F, Q, R\}$. This can be done by taking the partial derivative of the Q-function and solving when set equal to zero. In that case we get the following updating formulas for the unrestricted parameters:

$$\mu_{0,(j)} = x_{0|T} \tag{42}$$

$$\Sigma_{0,(j)} = \left(x_{0|T} - \mu_0\right)\left(x_{0|T} - \mu_0\right)' + P_{0|T}(\text{fixed}) \tag{43}$$

$$C_{(j)} = T^{-1}\left(S_1 - FS_0\right) \quad \rightarrow \quad C_{(j)} = T^{-1}\left(S_1 - \left(S_{10} - T^{-1}S_1S_0'S_{00}^{-1}\right)\left(I - T^{-1}S_0S_0'S_{00}^{-1}\right)^{-1}S_0\right) \tag{44}$$

$$F_{(j)} = \left(S_{10} - CS_0'\right)S_{00}^{-1} \quad \rightarrow \quad F_{(j)} = \left(S_{10} - T^{-1}S_1S_0'S_{00}^{-1}\right)\left(I - T^{-1}S_0S_0'S_{00}^{-1}\right)^{-1} \tag{45}$$

$$Q_{(j)} = T^{-1}\left(S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right) \tag{46}$$

$$R_{(j)} = T^{-1} \sum_{t=1}^{T} \left(\left(z_t - H_t x_{t|T} \right) \left(z_t - H_t x_{t|T} \right)' + H_t P_{t|T} H_t' \right) \tag{47}$$

Completing the M-step (Thorough derivations of equation (42) - (47) can be found in Appendix B). Note that the initial mean and covariance cannot be estimated simultaneously so it is conventional to fix one or both.

To summarize the general procedure:

1. Initialize the procedure by selecting initial values for the unknown parameters $\theta^{(0)}$ and fix either μ_0 or Σ_0 or both. For iterations $j = 1, 2, ...$

2. Perform the E-step: Compute the complete-data likelihood $-2lnL\left(X, Z; \theta\right)$, through equation (41)

3. Use the Kalman Filter, the Kalman Smoother and the One-Lag Covariance Smoother to obtain smoothed values for $x_{t,n|T}$, $P_{t,n|T}$ and $P_{t,t-1,n|T}$, for $t = 1, ..., T$ using parameters $\theta^{(j-1)}$. Use these to calculate S_0, S_1, S_{00}, S_{10} and S_{11}

4. Perform the M-step: Update the estimates $\mu_0,\ \Sigma_0,\ C,\ F,\ Q,\ R$ to get $\theta^{(j)}$

5. Repeat steps 2-4 until the maximum number of iterations has been reached or convergence has succeeded to desired precision, i.e. the likelihood and the estimates remain stable.

Or put Graphically:

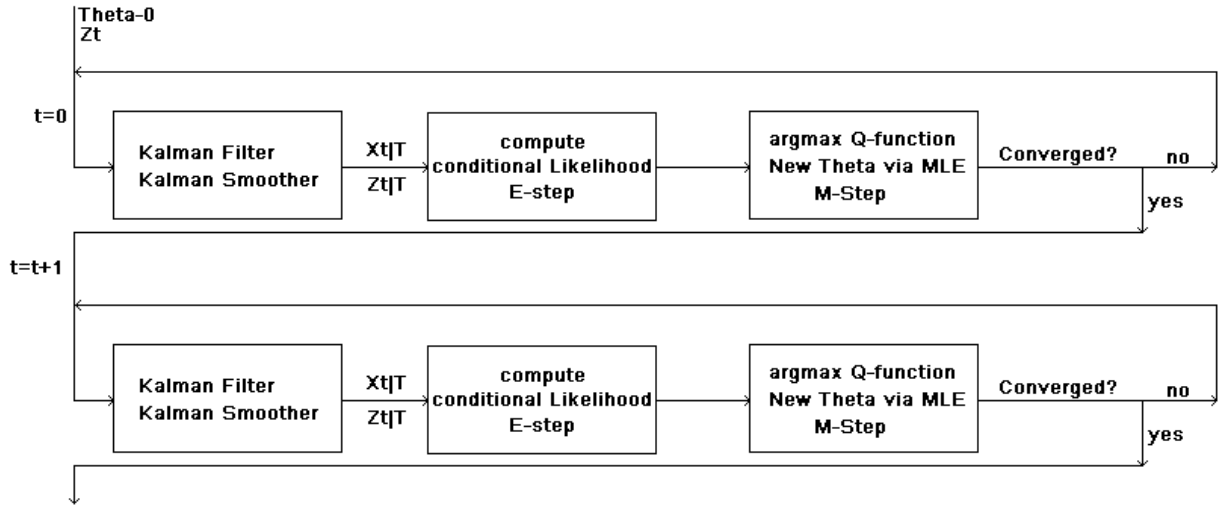

Figure 7: Schematic representation of the EM-algorithm

Although EM is generally considered one of the most powerful algorithms for maximum likelihood estimation in incomplete data problems, being computationally relatively simple and numerically stable, it does also have it's drawbacks. Advantage that the EM procedure does not need the second order partial derivatives of the likelihood function with respect to the parameters is countered by the disadvantage that standard errors of the parameter estimates cannot be obtained directly. This matrix can however be approximated by perturbing the likelihood function in the neighborhood of the maximum or alternatively by use of nonparametric bootstrap procedures as shown by Stoffer and Wall (1991). Second little inconvenience is the additional need for the Kalman Smoother and One-Lag Covariance smoother next to the the Kalman Filter.

Normally this procedure would suffice. In our case however one would like to impose some restrictions and/or parametrizations on the parameter space θ as to diminish the number of estimates. Unfortunately in that case elegant parameter updating formulas as (42) - (47) are hard to find or may not exist at all enforcing the need for yet another algorithm within the EM algorithm for numerical optimization of the likelihood function, (41). For this purpose a simple yet effective custom made Newton Raphson like algorithm was developed. Details can be found in Appendix B: Newton Raphson Algorithm

Restrictions and Parametrizations

In order to obtain the model proposed in the former Kalman Filter section, at least two parameter restrictions are to be imposed. First of which is on the AR coefficients matrix F. As we propose an AR(1) model over every instance i through the day and not a first order Vector Autoregressive model or 'VAR(1)', F should be a diagonal matrix having AR coefficients on the diagonal and zeros elsewhere. Furthermore we assume the eigenvalues of F to be less than 1 in absolute value to secure stability of the filter[21]. This assumption can be weakened, see Harvey (1991) section 4.3, but it is retained for simplicity. Second restriction lies in the blockdiagonal matrix Q with Q_ζ and Q_η on its diagonal and zeros elsewhere.

Further restrictions are not per see needed in small dimensional problems. But parameter estimation could seriously degrade in higher dimensional cases as the amount of parameters to be estimated increases exponentially. Some restrictions or parametrizations could in this case help confine the estimation errors, be it at cost of some additional estimation bias. Which basically makes it a bias-variance tradeoff problem as to obtain lower MSPE $\left(MSPE\left(R\right) = bias\left(R\right)^2 + variance\left(R\right)\right)$. Matrices suitable to such an approach are in this case R and Q.

To start with R: as $n^* \to n$ in (B.31), α and β will with certainty converge to respectively 0 and 1. Both reaching these limits as all of the days volatility is incorporated in the prediction. Moreover, as can be seen from figure 5, α and β follow on average quite a smooth path towards their end values. It will therefore come as no surprise that v_t will also follow a resembling stylized path with the estimation insecurity tightening (decreasing covariances) as n^* comes closer to n. The covariance of v_t is therefore also bounded by zero for $\sigma_{v,n}\sigma_{v,n}$, i.e.

$$R\left(n \times n\right) = \begin{bmatrix} \sigma_{v,1}^2 & \cdots & \sigma_{v,1}\sigma_{v,n} \\ \vdots & \searrow & \vdots \\ \sigma_{v,n}\sigma_{v,1} & \cdots & 0 \end{bmatrix}$$

The way in which this decrease manifests itself could be parsimoniously parametrized or alternatively restricted as to estimate it using less parameters than the original problem. It should however be noted that R is the covariance matrix of the measurement equation and should therefore be at least positive semi-difinite. If not, R wouldn't actually be a covariance matrix! More strictly, as R is used in a Maximum Likelihood Estimation procedure here, it should be positive definite and well-conditioned. That is all eigenvalues of the matrix must be > 0 so that it is non-

[21] In a diagonal matrix the eigenvalues are the same as the diagonal elements themselves, meaning all AR coefficients must be between -1 and 1.

singular (i.e. full rank) and the ratio of maximum and minimum singular value is not too large. Furthermore it must be noted that the implied correlation matrix should always be within the boundaries of $[-1, 1]$. If these properties are not satisfied one would encounter problems inverting the covariance matrix as needed in MLE, yielding inaccurate estimates and therefore wrong likelihood values. Additionally if the determinant of R becomes smaller than zero, the likelihood even yields imaginary outcomes by log transformation on $det(R)$. A seriously unwanted effect. Luckily positive definiteness of the matrix also excludes such occurrences.

For Q the same approach was taken. As can be seen from equation (B.28) and (B.29), when $n^* \to n$, $\varphi \to 1$. This is again a direct consequence of the way we predict end-of-the-day volatility and consequently what the outcome of α and β will be when $n^* \to n$. The pattern in which Q diverges to zero is however far less stylized and cannot sensible be parametrized.

$$Q\,(2n \times 2n) = \begin{bmatrix} \sigma^2_{\zeta,1} & \cdots & \sigma_{\zeta,1}\sigma_{\zeta,n} & 0 & \cdots & 0 \\ \vdots & \searrow & \vdots & \vdots & \ddots & \vdots \\ \sigma_{\zeta,n}\sigma_{\zeta,1} & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & \sigma^2_{\eta,1} & \cdots & \sigma_{\eta,1}\sigma_{\eta,n} \\ \vdots & \ddots & \vdots & \vdots & \searrow & \vdots \\ 0 & \cdots & 0 & \sigma_{\eta,n}\sigma_1 & \cdots & 0 \end{bmatrix} \tag{48}$$

For this research a number of parametrizations, restrictions and enhancements are put on trial. Details here fore and graphical representations of R, Q_ζ and Q_η together with their parametrized/restricted counterparts based on a first year estimation sample from rolling window M-Z regressions can be found in Appendix B: Restrictions and Parametrizations.

3.4 Forecast Revisions

Preliminary results on the R^2_{MAD} from previous section Minzer-Zarnowitz regressions proof forecast do on average improve during the day, especially after open outcry trading has started. Therefore we now turn to the investigation on how revisions of these forecasts evolve over time and what inherent properties they possess. If consecutive revisions are correlated we could investigate whether it is overreaction to news, some kind of smoothing behavior or other forces that drive forecasts and their revisions. Subsequent these patters may be usable to enhance volatility forecasting performance or at least give a more thorough insight in the forecasting process.

To obtain the desired results the regression,

$$I\hat{V}_{t,i} - I\hat{V}_{t,i-1} = \alpha + \beta \left(I\hat{V}_{t,i-1} - I\hat{V}_{t,i-2} \right) + \varepsilon_{t,i}, \tag{49}$$

is used. Under weak-form market efficiency news reaching the market should be implemented adequately and direct. Hence it should hold that consecutive forecast revisions are uncorrelated, i.e. $\beta = 0$. Resent studies by for example Clements (1997) or Isengildina et al. (2006) find however, that this hypothesis is often rejected. Leading to the case $\beta > 0$ commonly referred to as "forecast smoothing" or $\beta < 0$, referred to as "over reaction", that is overreaction to news which gets partly undone during the next forecasting round. Franses et al. (2011) on the other hand first decompose forecast in a model-based part and an expert intuition part. Their research exposes that a $\beta \neq 0$ could have multiple causes, depending on intuition- and news variance and mutual covariances. However as our forecast consist only of econometric model predictions we adhere to their special case (i) and should thus expect β to equal zero[22].

[22]Note that former research is solely performed on longer horizon forecast revisions for e.g. macroeconomic variables. Hence empirical findings on positive or negative β coefficients do not carry forth to this research. Theoretical findings however do.

4 Results and discussion

Throughout this section above stated methods and models are put to the test. Interesting findings will be reported and graphically supported. Furthermore the impact of and sensitivity to ad hoc parameter choices are discussed.

4.1 Forecasting results and Statistical evaluation

Economic forecasts and volatility forecasts all the same usually get updated as time goes by, incorporating the newly available information to make more accurate forecasts. However, as news reaching the market contains varying degree of signal to noise, forecasts could also worsen. To gain a thorough insight in this process the evolution and properties of the R^2_{MAD},$\hat{\alpha}$ and $\hat{\beta}$ from Mincer-Zarnowitz type regressions, as proposed in (25), are offset against starting and stopping time of intra-day volatility observations.

Yet before we proceed we briefly discuss the sensitivity to extreme return observations. Throughout this research we rely on robust performance measures as to strengthen conclusions in the potential presence of extreme observations, and indeed such more advances procedures prove worthwhile. Take figure 8 for example, sub a) and b) display the squared correlations for respectively OLS and GLS Mincer-Zarnowitz regressions on US30 data. It then becomes remarkably clear that extreme returns observations poses the potency to distort conclusions in an ordinary (OLS) setting.

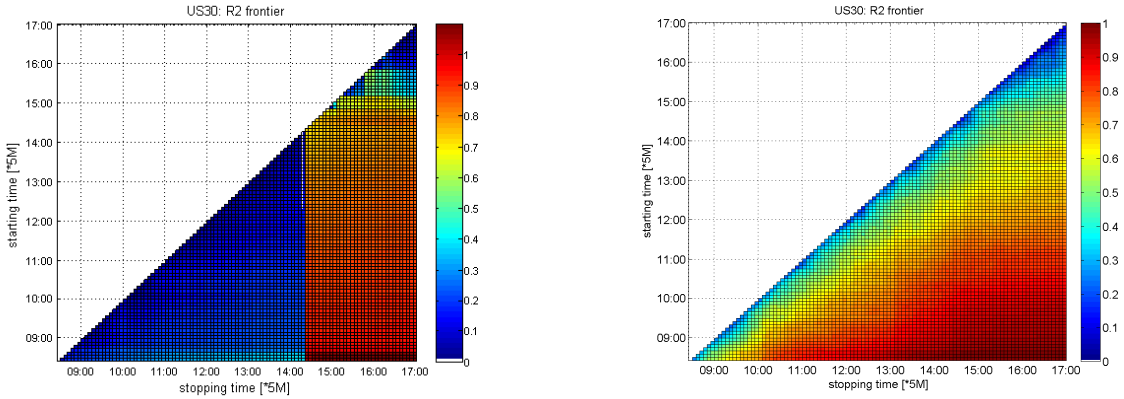

(a) R^2 frontier including outliers

(b) R^2_{MAD} frontier excluding outliers

Figure 8: Displayed figures show R^2 resp. R^2_{MAD} from OLS-MZ and GLS-MZ regressions on US30 data with Moving Average Seasonal forecasts using H=22. Figure a) is unadjusted for heteroskedasticity/outliers in the return series. In figure b) strange distortions are resolved due to the more robust performance measure, consequently following a smoother path. Producing irrefutable evidence for the R^2 sensitivity to extreme return observations and the value of robust statistics.

To elaborate upon the gain through the models employed we start by discussing the performance of 'Begin-of-the-day volatility' as basis and subsequently deal with the four models (MZS, MA, EWMA, FFF) and compare these to conventional benchmarkes as RW and GARCH(1,1). Subsection 2 continues with the properties of MZ $\alpha$and β as estimated through Kalman Filtering and subsection 3 elaborates upon properties of forecast revisions.

4.1.1 Begin-of-the-day variance (reference)

To test whether the idea of scaling begin-of-the-day variance has some forecasting potential, i.e. forms an informative regressor, statistics hereof are included in this discussion. Subsequently it is of interest whether scaling this variance to daily proportions by Mincer-Zarnowitz coefficients or a diurnal pattern like the Seasonal Average, Exponentially Weighted Moving Average or Fourier Flexible Form improves the correlation to end-of-the-day variance. Note again that begin-of-the-day variance in itself is not a complete forecast. It is biased by construction as estimates are unscaled. MZ β can therefore not be expected to equal 1.

For evaluation we do not only look at statistics of first 15, 30, 60 and 120 minute intervals but rather for an efficient timing interval which will be sought for by maximization of R^2_{MAD}. Table 3 displays the results for S&P500 and US30 data.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2561^*$ (0,0049)	$15,8645^*$ (2,5868)	0,4875	0,0730	0,4145	0,8642	$0,1636^*$ (0,0094)	$10,0993^*$ (3,2247)	0,2989	0,0978	0,2011	0,8309
	BPV	$0,1649^*$ (0,0027)	$10,7495^*$ (2,4899)	0,4738	0,0701	0,4037	0,8688	$0,1027^*$ (0,0049)	$7,7055^*$ (2,0507)	0,2808	0,0869	0,1940	0,8462
	RR	$0,0995^*$ (0,0056)	$13,8934^*$ (0,5038)	0,6454	0,0690	0,5765	0,8685	$0,0898^*$ (0,0043)	$10,9593^*$ (0,8843)	0,4614	0,0768	0,3846	0,8587
	TTS	$0,1829^*$ (0,0064)	$12,9811^*$ (0,5822)	0,6163	0,0692	0,5471	0,8691	$0,1269^*$ (0,0079)	$10,9710^*$ (1,5133)	0,4070	0,0945	0,3125	0,8338
	Kernel	$0,1785^*$ (0,0068)	$11,9005^*$ (0,4845)	0,6206	0,0755	0,5451	0,8579	$0,1263^*$ (0,0054)	$10,6478^*$ (1,0351)	0,4507	0,0946	0,3561	0,8328
	Mean	0,1764	13,0778	0,5687	0,0714	0,4974	0,8657	0,1219	10,0766	0,3798	0,0901	0,2897	0,8405

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1723^*$ (0,0042)	$7,3480^*$ (0,3320)	0,6294	0,1265	0,5028	0,7706	$0,1163^*$ (0,0145)	$7,0886^*$ (1,8380)	0,5033	0,1498	0,3535	0,7436
	BPV	$0,1090^*$ (0,0020)	$6,6374^*$ (0,3355)	0,6111	0,1308	0,4803	0,7634	$0,1005^*$ (0,0039)	$2,1281$ (0,9914)	0,4115	0,1385	0,2730	0,7597
	RR	$0,0670^*$ (0,0048)	$7,9053^*$ (0,2147)	0,7406	0,1256	0,6150	0,7683	$0,0724^*$ (0,0063)	$6,5220^*$ (0,5753)	0,5558	0,1305	0,4253	0,7658
	TTS	$0,1301^*$ (0,0058)	$7,1567^*$ (0,2299)	0,6881	0,1286	0,5595	0,7647	$0,0974^*$ (0,0106)	$6,7296^*$ (0,9136)	0,5025	0,1481	0,3544	0,7437
	Kernel	$0,1266^*$ (0,0071)	$7,1470^*$ (0,2300)	0,7239	0,1308	0,5931	0,7607	$0,1005^*$ (0,0096)	$6,5541^*$ (0,7479)	0,5770	0,1448	0,4322	0,7478
	Mean	0,1210	7,2389	0,6786	0,1285	0,5501	0,7655	0,0974	5,8045	0,5100	0,1423	0,3677	0,7521

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1319^*$ (0,0061)	$3,6538^*$ (0,1297)	0,7627	0,2430	0,5196	0,5872	$0,0920^*$ (0,0077)	$3,7911^*$ (0,3810)	0,5753	0,2270	0,3483	0,6212
	BPV	$0,0828^*$ (0,0029)	$3,5792^*$ (0,1178)	0,7469	0,2390	0,5079	0,5929	$0,0525^*$ (0,0019)	$3,8826^*$ (0,2470)	0,5427	0,2159	0,3268	0,6370
	RR	$0,0497^*$ (0,0057)	$4,0415^*$ (0,1126)	0,7977	0,2477	0,5500	0,5748	$0,0556^*$ (0,0059)	$4,1482^*$ (0,2854)	0,6583	0,2090	0,4493	0,6385
	TTS	$0,0994^*$ (0,0075)	$3,8609^*$ (0,1269)	0,8015	0,2438	0,5577	0,5828	$0,0745^*$ (0,0082)	$4,1378^*$ (0,3600)	0,6399	0,2254	0,4145	0,6209
	Kernel	$0,0872^*$ (0,0078)	$3,9218^*$ (0,1156)	0,8055	0,2471	0,5584	0,5771	$0,0774^*$ (0,0093)	$4,1190^*$ (0,3652)	0,6708	0,2196	0,4512	0,6277
	Mean	0,0902	3,8114	0,7829	0,2441	0,5387	0,5830	0,0704	4,0157	0,6174	0,2194	0,3980	0,6291

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0798^*$ (0,0073)	$2,4128^*$ (0,0743)	0,8711	0,3933	0,4777	0,3868	$0,0478^*$ (0,0090)	$2,4255^*$ (0,1897)	0,7862	0,3909	0,3953	0,3963
	BPV	$0,0458^*$ (0,0039)	$2,4852^*$ (0,0705)	0,8589	0,3857	0,4732	0,3965	$0,0272^*$ (0,0026)	$2,4205^*$ (0,1196)	0,7801	0,3841	0,3959	0,4048
	RR	$0,0338^*$ (0,0057)	$2,5005^*$ (0,0643)	0,8889	0,4035	0,4853	0,3699	$0,0327^*$ (0,0079)	$2,5058^*$ (0,1822)	0,8381	0,3748	0,4632	0,4064
	TTS	$0,0705^*$ (0,0078)	$2,4419^*$ (0,0736)	0,8877	0,3937	0,4941	0,3837	$0,0396^*$ (0,0097)	$2,5090^*$ (0,1934)	0,8253	0,3912	0,4341	0,3930
	Kernel	$0,0556^*$ (0,0077)	$2,5094^*$ (0,0654)	0,9011	0,3958	0,5053	0,3800	$0,0459^*$ (0,0126)	$2,5010^*$ (0,2296)	0,8356	0,3806	0,4550	0,4036
	Mean	0,0571	2,4700	0,8815	0,3944	0,4871	0,3834	0,0386	2,4724	0,8131	0,3843	0,4287	0,4008

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:10	$0,1563^*$ (0,0055)	$4,9651^*$ (0,2000)	0,7128	0,1779	0,5349	0,6871	08:20	09:45	$0,0771^*$ (0,0071)	$3,0628^*$ (0,2455)	0,7026	0,2819	0,4207	0,5406
	BPV	09:30	10:15	$0,0885^*$ (0,0029)	$4,5634^*$ (0,1656)	0,7257	0,1923	0,5333	0,6641	08:25	10:10	$0,0336^*$ (0,0024)	$2,5950^*$ (0,1310)	0,7679	0,3566	0,4113	0,4610
	RR	09:30	10:00	$0,0670^*$ (0,0048)	$7,9053^*$ (0,2147)	0,7406	0,1256	0,6150	0,7683	08:40	10:10	$0,0335^*$ (0,0075)	$3,7356^*$ (0,2561)	0,8352	0,3473	0,4879	0,5696
	TTS	09:30	09:50	$0,1682^*$ (0,0064)	$9,9273^*$ (0,4086)	0,6647	0,0892	0,5755	0,8332	08:20	10:00	$0,0480^*$ (0,0081)	$2,9814^*$ (0,2040)	0,7883	0,3233	0,4650	0,4797
	Kernel	09:30	09:50	$0,1516^*$ (0,0074)	$9,6662^*$ (0,3582)	0,6940	0,0955	0,5985	0,8220	08:35	10:05	$0,0422^*$ (0,0118)	$3,9729^*$ (0,3236)	0,8243	0,3394	0,4849	0,5784
	Mean	9:30	10:01	0,1263	7,4054	0,7076	0,1361	0,5714	0,7549	8:28	10:02	0,0469	3,2695	0,7837	0,3297	0,4540	0,5259

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	16:15	08:35	$0,7905^*$ (0,0025)	$-3,9259^*$ (0,3935)	0,6898	0,0000	0,6898	0,7692	18:25	08:35	$0,0345^*$ (0,0175)	$1,8736^*$ (0,1988)	0,6796	0,0978	0,5818	16,1001
	BPV	04:20	08:30	$0,0734^*$ (0,0025)	$6,7204^*$ (0,2414)	0,6845	0,0000	0,6845	0,7405	20:10	08:30	$0,0349^*$ (0,0044)	$2,1167^*$ (0,1411)	0,6291	0,0504	0,5786	0,4063
	RR	16:15	09:00	$0,5406^*$ (0,0007)	$-4,2321^*$ (0,3059)	0,7112	0,0000	0,7112	0,7098	01:00	08:35	$0,0662^*$ (0,0090)	$2,1489^*$ (0,2196)	0,7168	0,0768	0,6399	0,4500
	TTS	16:30	09:30	$0,3728^*$ (0,0118)	$0,3287^*$ (0,3346)	0,7149	0,0000	0,7149	0,6644	18:05	08:45	0,0159 (0,0177)	$1,8844^*$ (0,1841)	0,7408	0,1330	0,6078	11,8562
	Kernel	16:15	09:10	$0,7629^*$ (0,0028)	$-2,7417^*$ (0,4354)	0,7229	0,0000	0,7229	0,7735	18:25	08:35	0,0125 (0,0216)	$1,9407^*$ (0,2030)	0,7132	0,0946	0,6187	8,3767
	Mean	13:55	8:57	0,5081	-0,7701	0,7047	0,0000	0,7047	0,7315	15:13	8:36	0,0328	1,9929	0,6959	0,0905	0,6054	7,4379

(j) Statistics for the maximum based on R^2_{marg} over 24 hour trade.

Table 3: Overall reference statistics for S&P500 and US30 futures. Standard errors of the estimates are given between parenthesis and a star (*) is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

Beholding sub-table a-d) there are a few differences between the two time series. Observing only the first 15 minutes of S&P500 index futures returns we are able to produce an average R^2_{MAD} up to 0.57 observing only 7% of daily variance itself. Doubling this to 30, 60 and 120 minutes leads to a steady increase up to respectively 0.68, 0.78, and even 0.88 with respectively 13%, 24%, 39% of daily variance observed. This is a major increase upon Random Walk and GARCH(1,1) statistics which on average respectively reach R^2_{MAD} of 0.70 and 0.55. Begin-of-the-day variance can therefore be thought a high potential regressor in the sense that scaled to daily proportions it could form an unbiased and highly predictive forecast[23]. US30 volatility seems only somewhat less predictable. Here the first 15, 30, 60, 120 min. on average lead to R^2_{MAD} of 0.38 (VR = 9%), 0.51 (VR = 14%), 0.62 (VR = 22%) and 0.81 (VR = 38%) resp. Meaning it takes up to approximately 30 minutes of intraday data to beat the random walk (R^2_{MAD}= 0.51) and only 20 minutes to overtake GARCH (R^2_{MAD}= 0.45) taken from start of open outcry trading to beat

[23]Note that the comparison between RW/GARCH and begin-of-the-day based forecasts is intended to illustrate the potential gain on conventional methods in using begin of the day volatility. For a fair comparison on predictive performance, readily observed variance should first be subtracted from end-of-day variance.

GARCH predictability. In terms of HMSPE both series immediately defeat RW and GARCH. However this can be thought more due to small daily volatility figures occurring more frequently than due to more accurate 'predictions'. As such large volatility figures are simply dwarfed and potential errors therefore confined.

Sub-tables e/f) present us with some other interesting insights. When obtaining maximum R^2_{MAD} for S&P500, first 5 minutes of data is sometimes left aside, for US30 this even comprises of the first 5-10 minutes. As the penalty term, VR, is calculated as the percentage of variance one could have observed until stopping time, this means that the R^2_{MAD} for an interval starting at start of open outcry trade actually has lower R^2_{MAD}. In other words omitting these first observations leads to higher predictability in terms of correlation. One could therefore argue that data from the first few minutes of the day is corrupted in such a way that it contains more noise than signal.

The high information value especially at the beginning of the day can further be assessed by the R^2_{MAD} and R^2_{marg} frontiers. In the case of realized variance, these plots are shown in figure 9. They can be seen as the different squared correlation coefficients obtained from begin-of-the-day volatility while diversifying starting and stopping time by 5 minute intervals. It then becomes clear that the frontier is upward sloping as more of the days volatility reveals itself, i.e. additional information leads to an *almost* continues growing ability to forecasts. Furthermore it gives clear insight in the marginal effect of new information as the slope decreases over time.

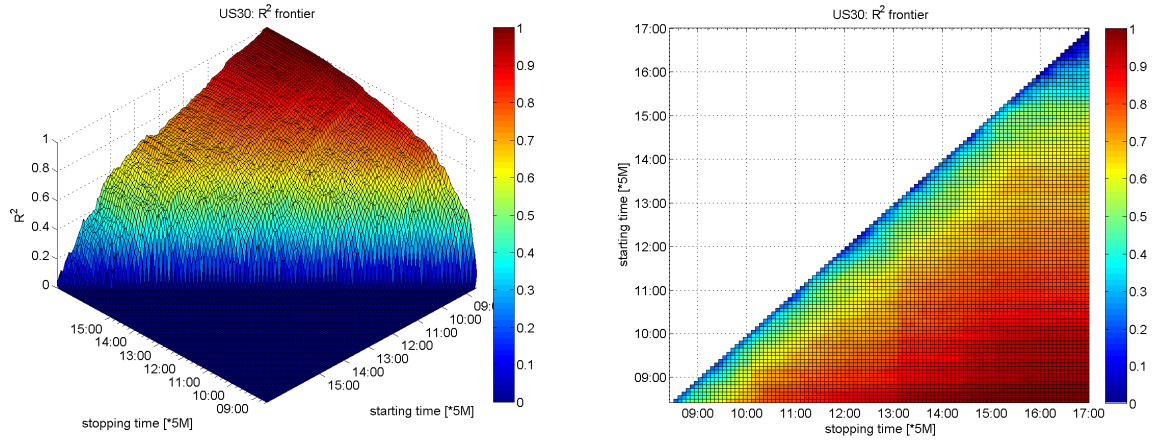

(a) R^2_{MAD} frontiers in 3D and top view

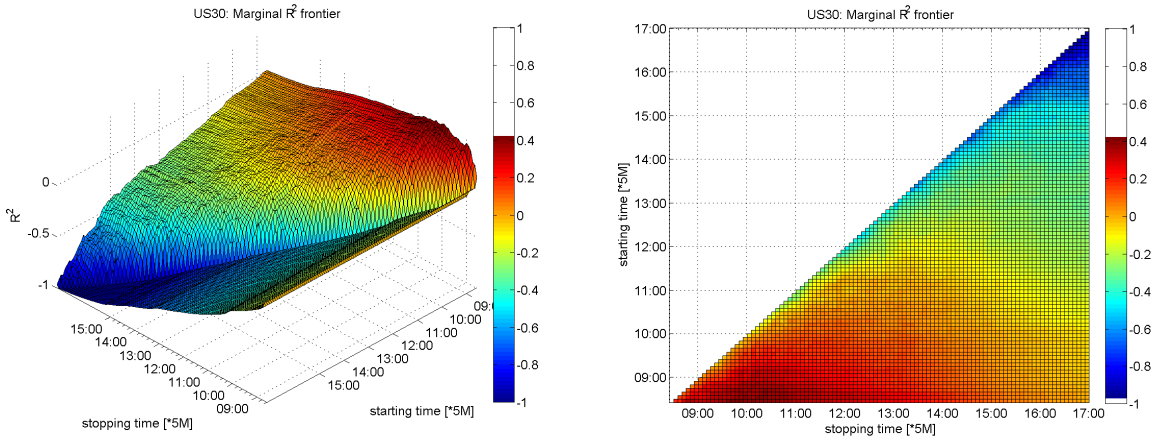

(b) R^2_{marg} frontiers in 3D and top view

Figure 9: Sub-figure a) reveals the R^2_{MAD} for US30 data using Realized Variance as proxy. Sub-figure b) shows the R^2_{marg}. Note that in both cases half of the graph takes the lowest possible point or is left blank as the start of observations cannot overtake the stopping time.

Plots displayed above represent R^2_{MAD} and R^2_{marg} for Realized Variance. Similar plots can be obtained for other measures but were omitted due to the marginal extra insight (they are available upon request). Analyzing these figures, there are again a few notable appearances. First of which is the steeper rise in R^2_{MAD} near the start of open outcry markets, demonstrating higher informational value during these periods, directing the correlation steeply upwards. R^2_{marg} on the other hand rises near the opening to its maximum and steadily declines afterward as the penalty for observing more of the days return volatility becomes more severe. Meaning the extra R^2_{MAD} to be gained from additional observations is in time overruled by the additional return volatility needed obtain such predictions.

Expanding these figures to incorporate overnight information in the same manner as done before (figure 9 can be thought only the tip of an iceberg as similar forecasts can be estimated

using overnight information), figure 10 displays some new interesting findings. For instance we see that the information obtained overnight does not seem to create an additional advantage over RW predictions (black striped line). Near the start of the day an advantage can be observed yet this is not consistent among other volatility measures and cannot be preserved as compared to open outcry data only (red line). This latter property could however well be due to the way end-of-the-day variance is defined (variance over open outcry trade returns only, therefore not including overnight returns) then by the estimates being actually less informative.

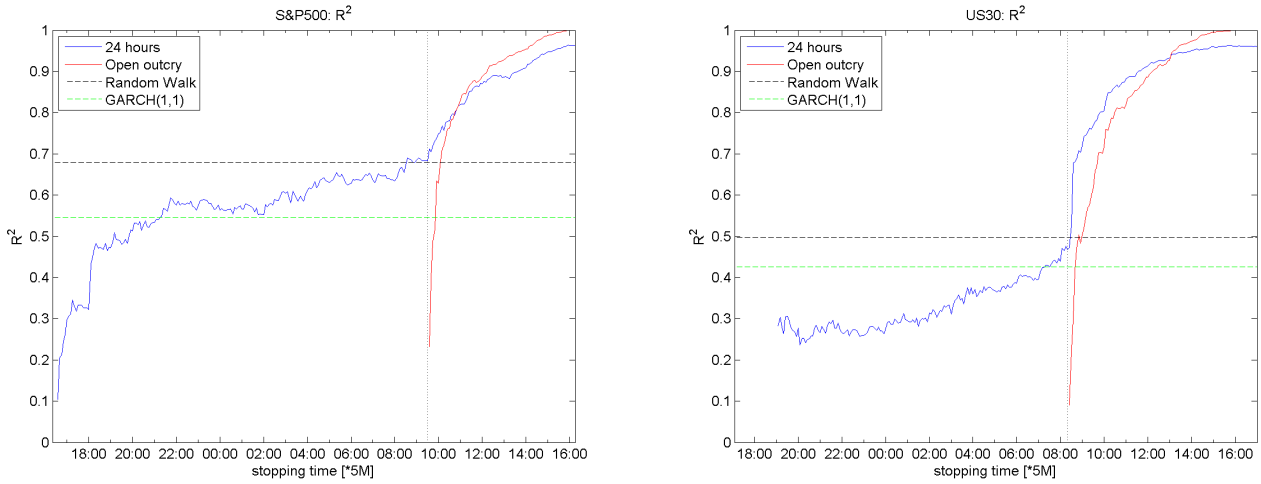

(a) S&P500: R^2_{MAD} **(b)** US30: R^2_{MAD}

Figure 10: R^2_{MAD} coefficient lines: first including overnight information (blue line), second using open outcry information only (red line), third a comparative line is inserted for Random Walk estimation (black striped line), forth is a comparative line for GARCH(1,1) estimation (green striped line). The vertical dotted line indicates start of open outcry markets.

4.1.2 Direct Mincer-Zarnowitz Scaling

As mentioned, Begin-of-the-day variance needs to be scaled to daily proportions in order to create an unbiased forecast. MZ regression coefficients could be used for such adjustments. Associated R^2_{MAD} frontiers using RV are depicted in figure 12. Henceforth only top-viewed figures are presented as these are easier interpretable.

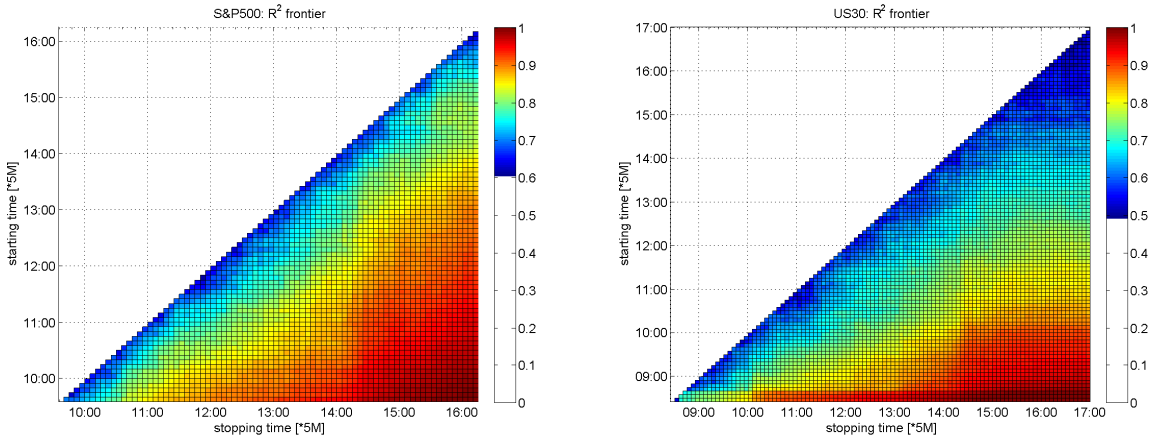

(a) S&P500: R^2_{MAD} **(b)** US30: R^2_{MAD}

Figure 11: Figures display the squared correlations of forecasts and estimated daily variance figures, i.e. R^2_{MAD}, for S&P500 as well as US30 futures return data using Realized Variance as volatility proxy. Note that in both cases half the graph is left blank as the start of observations can never overtake the stopping time.

Evidently such method not only creates unbiased estimates but also increases correlation significantly. Considering table 4, overall R^2_{MAD} is drastically higher than is the case for begin-of-the-day variance. Moreover the advantage is withhold over time (15, 30, 60, 120 min) and is present for both securities. Where greatest profit is pleasingly gained near start of open outcry trade. After just 15 minutes of return observations squared correlations are up to 0.7 growing to 0.75, 0.8 and 0.88 for resp. 30, 60 and 120 minutes.

Obtaining the most efficient sample (table 4e) S&P500 and US30 now only need 15 resp 20 minutes of return observations to reach saturation in the sense that more extra observed variance is needed than can be gained in estimation. Reliable forecasts can therefore be obtained far earlier in the day, where it is again noted that US30 estimation notoriously leaves the first 5 minutes of observations out of estimation. For S&P500 this comprises of the first 0 to 5 minutes.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,0430^*$ (0,0130)	$0,9361$ (0,0349)	0,6934	0,0733	0,6202	1,2734	$0,0075$ (0,0084)	$0,9922$ (0,0455)	0,6925	0,0974	0,5952	0,4921
	BPV	$0,0189^*$ (0,0070)	$0,9558$ (0,0324)	0,6833	0,0703	0,6130	1,9638	$0,0092^*$ (0,0039)	$0,9535$ (0,0396)	0,7317	0,0863	0,6454	0,4354
	RR	$0,0297^*$ (0,0085)	$0,9352^*$ (0,0303)	0,7048	0,0691	0,6357	0,6934	$0,0039$ (0,0066)	$0,9986$ (0,0442)	0,7065	0,0764	0,6300	0,3235
	TTS	$0,0377^*$ (0,0117)	$0,9378^*$ (0,0309)	0,6909	0,0695	0,6214	0,7387	$0,0061$ (0,0089)	$0,9991$ (0,0487)	0,7039	0,0942	0,6097	0,4059
	Kernel	$0,0371^*$ (0,0119)	$0,9398^*$ (0,0299)	0,6907	0,0757	0,6151	0,7824	$0,0040$ (0,0106)	$1,0123$ (0,0527)	0,7079	0,0941	0,6139	0,4003
	Mean	0,0333	0,9409	0,6926	0,0716	0,6211	1,0903	0,0061	0,9911	0,7085	0,0897	0,6188	0,4114

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,0387^*$ (0,0111)	$0,9309^*$ (0,0305)	0,7282	0,1269	0,6013	0,9626	$0,0062$ (0,0085)	$0,9972$ (0,0466)	0,7230	0,1492	0,5737	0,5215
	BPV	$0,0141^*$ (0,0065)	$0,9707$ (0,0308)	0,7121	0,1312	0,5808	1,2526	$0,0678^*$ (0,0048)	$0,4384^*$ (0,0514)	0,7527	0,1378	0,6149	0,3982
	RR	$0,0270^*$ (0,0074)	$0,9294^*$ (0,0268)	0,7340	0,1259	0,6082	0,5986	$0,0017$ (0,0068)	$1,0098$ (0,0455)	0,7525	0,1300	0,6225	0,3402
	TTS	$0,0402^*$ (0,0107)	$0,9184^*$ (0,0293)	0,7705	0,1290	0,6415	0,5852	$0,0027$ (0,0095)	$1,0156$ (0,0523)	0,7553	0,1476	0,6077	0,4281
	Kernel	$0,0465^*$ (0,0112)	$0,9124^*$ (0,0288)	0,7395	0,1311	0,6083	0,5883	$0,0010$ (0,0109)	$1,0214$ (0,0547)	0,7626	0,1443	0,6183	0,5170
	Mean	0,0333	0,9324	0,7369	0,1288	0,6080	0,7975	0,0159	0,8965	0,7492	0,1418	0,6074	0,4410

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,0231^*$ (0,0089)	$0,9575$ (0,0258)	0,7919	0,2432	0,5487	0,4769	$0,0034$ (0,0089)	$1,0077$ (0,0494)	0,7620	0,2266	0,5354	0,3799
	BPV	$0,0144^*$ (0,0052)	$0,9498$ (0,0258)	0,7897	0,2393	0,5504	0,6479	$0,0043$ (0,0034)	$0,9879$ (0,0358)	0,7813	0,2153	0,5660	0,3105
	RR	$0,0536^*$ (0,0092)	$0,8538^*$ (0,0347)	0,8251	0,2479	0,5771	0,3891	$0,0007$ (0,0067)	$1,0146$ (0,0458)	0,7794	0,2085	0,5709	0,2855
	TTS	$0,0346^*$ (0,0097)	$0,9315^*$ (0,0276)	0,8076	0,2441	0,5635	0,3616	$0,0006$ (0,0093)	$1,0217$ (0,0515)	0,7825	0,2251	0,5574	0,3497
	Kernel	$0,0649^*$ (0,0125)	$0,8754^*$ (0,0336)	0,8171	0,2473	0,5698	0,3778	$-0,0022$ (0,0112)	$1,0332$ (0,0565)	0,7874	0,2191	0,5683	0,4729
	Mean	0,0381	0,9136	0,8063	0,2444	0,5619	0,4507	0,0014	1,0131	0,7785	0,2189	0,5596	0,3597

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0267^*$ (0,0075)	$0,9366^*$ (0,0228)	0,8781	0,3939	0,4842	0,3423	$-0,0005$ (0,0083)	$1,0197$ (0,0476)	0,8811	0,3906	0,4905	0,3075
	BPV	$0,0150^*$ (0,0045)	$0,9412^*$ (0,0232)	0,8708	0,3864	0,4843	0,5145	$0,0017$ (0,0030)	$0,9968$ (0,0324)	0,8667	0,3835	0,4833	0,2404
	RR	$0,0356^*$ (0,0067)	$0,8975^*$ (0,0265)	0,8918	0,4041	0,4877	0,2828	$0,0014$ (0,0069)	$1,0070$ (0,0483)	0,8737	0,3743	0,4994	0,3092
	TTS	$0,0586^*$ (0,0102)	$0,8749^*$ (0,0306)	0,8941	0,3943	0,4999	0,2876	$-0,0008$ (0,0084)	$1,0220$ (0,0486)	0,8814	0,3907	0,4907	0,2841
	Kernel	$0,0514^*$ (0,0100)	$0,8930^*$ (0,0278)	0,9002	0,3963	0,5039	0,2766	$-0,0028$ (0,0108)	$1,0311$ (0,0564)	0,8837	0,3803	0,5034	0,4926
	Mean	0,0375	0,9087	0,8870	0,3950	0,4920	0,3408	-0,0002	1,0153	0,8773	0,3839	0,4935	0,3268

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	09:40	$0,4874^*$ (0,0281)	$-0,0031^*$ (0,0785)	0,7004	0,0519	0,6485	1,5975	08:25	08:35	$0,0073$ (0,0088)	$0,9970$ (0,0476)	0,7007	0,0974	0,6033	0,4354
	BPV	09:30	09:40	$0,0191^*$ (0,0082)	$0,9746$ (0,0381)	0,6807	0,0497	0,6310	2,3744	08:25	08:35	$0,0061$ (0,0036)	$0,9755$ (0,0370)	0,7375	0,0863	0,6512	0,4198
	RR	09:35	09:40	$0,0277^*$ (0,0089)	$0,9453$ (0,0310)	0,6975	0,0475	0,6500	1,9010	08:25	08:40	$0,0026$ (0,0066)	$1,0069$ (0,0441)	0,7347	0,0971	0,6376	0,3087
	TTS	09:30	09:40	$0,0515^*$ (0,0134)	$0,9179^*$ (0,0354)	0,6884	0,0468	0,6417	0,9883	08:25	08:45	$0,0025$ (0,0090)	$1,0154$ (0,0494)	0,7525	0,1326	0,6199	0,3624
	Kernel	09:35	09:55	$0,0280^*$ (0,0098)	$0,9499^*$ (0,0251)	0,7461	0,1134	0,6328	0,7295	08:25	08:35	$0,0093$ (0,0109)	$0,9942$ (0,0543)	0,7201	0,0941	0,6260	0,4660
	Mean	9:32	9:43	0,1227	0,7569	0,7026	0,0619	0,6408	1,5181	8:25	8:38	0,0056	0,9978	0,7291	0,1015	0,6276	0,3985

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	21:20	09:15	$0,0520^*$ (0,0125)	$0,9231^*$ (0,0341)	0,7249	0,0000	0,7249	1,4093	05:20	08:35	$0,0212^*$ (0,0080)	$0,9307$ (0,0438)	0,7285	0,0974	0,6311	0,9304
	BPV	01:50	09:35	$0,0264^*$ (0,0068)	$0,9234^*$ (0,0320)	0,7305	0,0276	0,7030	1,8146	06:35	08:35	$0,0057$ (0,0032)	$0,9657$ (0,0327)	0,7531	0,0863	0,6668	0,4608
	RR	22:50	09:05	$0,0783^*$ (0,0152)	$0,8178^*$ (0,0547)	0,7454	0,0000	0,7454	1,2962	03:35	08:40	$0,0053$ (0,0062)	$0,9833$ (0,0419)	0,7589	0,0971	0,6618	0,6126
	TTS	00:00	09:30	$0,0387^*$ (0,0108)	$0,9401^*$ (0,0292)	0,7237	0,0000	0,7237	0,9624	07:10	08:50	$0,0037$ (0,0088)	$1,0040$ (0,0484)	0,7810	0,1476	0,6334	0,5738
	Kernel	21:30	09:15	$0,0540^*$ (0,0119)	$0,9191^*$ (0,0304)	0,7325	0,0000	0,7325	1,1690	06:35	08:35	$0,1089^*$ (0,0080)	$0,5422^*$ (0,0437)	0,7528	0,0941	0,6587	0,6198
	Mean	13:30	9:20	0,0499	0,9047	0,7314	0,0055	0,7259	1,3303	5:51	8:39	0,0290	0,8852	0,7549	0,1045	0,6504	0,6395

(j) Statistics for the maximum based on R^2_{marg} over 24 hour trade.

Table 4: Overall Mincer-Zarnowitz bias adjusted variance forecasting statistics for S&P500 and US30 futures. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

4.1.3 Seasonal Moving Average

Second, a deterministic assumed diurnal pattern is obtained through the Seasonal Moving Average method and subsequently used for scaling begin-of-the-day variance. Associated R^2_{MAD} frontiers using RV are depicted in figure 12.

Figure 12 clearly shows the similarities and dissimilarities between both securities. It is apparent that estimating US 30 year treasury bond future return volatility is somewhat more difficult. However the gain from using a diurnal patterns is also greater. From the tables it becomes clear that the coefficients of determination for US30 data increase over the full range of the day as compared to the reference point - begin-of-the-day variance.

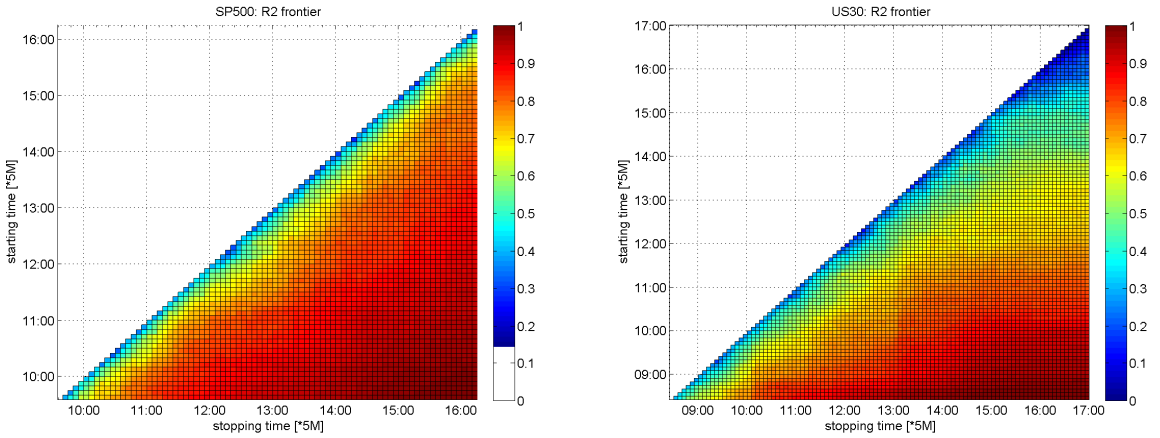

(a) S&P500: R^2_{MAD} **(b)** US30: R^2_{MAD}

Figure 12: Figures display the squared correlations of forecasts and estimated daily variance figures, i.e. R^2_{MAD}, for S&P500 as well as US30 futures return data using Realized Variance as volatility proxy. Note that in both cases half the graph is left blank as the start of observations can never overtake the stopping time.

After half an hour of open outcry trade the difference is up to 0.05 in favor of the Seasonal Moving Average where the sensitivity to parameter choice H (backward estimation window for seasonal) seems negligible. Looking at mutual in stead of average performance statistics, about the same margins can be gained. Nonetheless mutual statistics can be somewhat apart, demonstrating the usefulness of comparing multiple measures instead of just one. Sensitivity to the moving average filter span 'a', is small. Omitting this filter all together would result in the loss of just half a percent and therefore still a gain as compared to begin-of-the-day variance. This in contrast to S&P500 data where the smoothing filter has a greater effect: the differences here mount up to 0.04. This greater impact was to be expected as the unfiltered seasonal shape of S&P500 series is far more erratic than is the case for US30 data. Nonetheless, even using the filter the Seasonal Moving Average method on average performs on par with begin-of-day variance. This might feel somewhat counter intuitive as the S&P500 on average seems to experience a stronger diurnal pattern. Yet this is also the problem: on average the seasonal is quite well behaved, on a daily scale however, the seasonal is far less stylized due to the volatile nature of the S&P500 series. US30 series on the other hand experience a less stylized average seasonal pattern yet one that is more persistent through time and thereby more predictable. In other words US30 data strokes better with the deterministic assumption on the seasonal pattern.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2548^*$ (0,0049)	$0,6122^*$ (0,0996)	0,4755	0,0730	0,4025	4,0545	$0,1359^*$ (0,0055)	$0,8182$ (0,1540)	0,3701	0,0978	0,2723	3,8770
	BPV	$0,1634^*$ (0,0028)	$0,3891^*$ (0,0908)	0,4714	0,0701	0,4013	4,1770	$0,0804^*$ (0,0033)	$0,6695^*$ (0,0983)	0,3675	0,0869	0,2807	2,9452
	RR	$0,1022^*$ (0,0055)	$0,5066^*$ (0,0188)	0,6234	0,0690	0,5545	2,0110	$0,0905^*$ (0,0029)	$0,5460^*$ (0,0384)	0,5254	0,0768	0,4486	2,0350
	TTS	$0,1844^*$ (0,0064)	$0,4991^*$ (0,0228)	0,5909	0,0692	0,5217	2,4565	$0,1233^*$ (0,0049)	$0,6565^*$ (0,0762)	0,4870	0,0945	0,3924	3,2921
	Kernel	$0,1804^*$ (0,0068)	$0,4561^*$ (0,0191)	0,6481	0,0755	0,5726	3,0248	$0,1266^*$ (0,0036)	$0,6186^*$ (0,0558)	0,5110	0,0946	0,4164	3,2562
	Mean	0,1770	0,4926	0,5619	0,0714	0,4905	3,1448	0,1113	0,6618	0,4522	0,0901	0,3621	3,0811

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1707^*$ (0,0042)	$0,5657^*$ (0,0258)	0,6235	0,1265	0,4969	1,7687	$0,1149^*$ (0,0083)	$0,8670$ (0,1780)	0,5273	0,1498	0,3775	1,1071
	BPV	$0,1088^*$ (0,0020)	$0,4755^*$ (0,0247)	0,6224	0,1308	0,4916	2,2580	$0,0726^*$ (0,0026)	$0,6667^*$ (0,1003)	0,4618	0,1385	0,3233	1,0031
	RR	$0,0718^*$ (0,0047)	$0,5819^*$ (0,0163)	0,7108	0,1256	0,5852	1,1372	$0,0803^*$ (0,0041)	$0,6434^*$ (0,0501)	0,6271	0,1305	0,4967	0,7099
	TTS	$0,1321^*$ (0,0058)	$0,5501^*$ (0,0179)	0,6940	0,1286	0,5654	1,3520	$0,1049^*$ (0,0060)	$0,7856^*$ (0,0886)	0,5893	0,1481	0,4412	0,9909
	Kernel	$0,1316^*$ (0,0070)	$0,5406^*$ (0,0179)	0,7303	0,1308	0,5995	1,4037	$0,1152^*$ (0,0057)	$0,7150^*$ (0,0739)	0,5944	0,1448	0,4496	0,9589
	Mean	0,1230	0,5428	0,6762	0,1285	0,5477	1,5839	0,0976	0,7355	0,5600	0,1423	0,4177	0,9540

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1358^*$ (0,0060)	$0,5671^*$ (0,0206)	0,7593	0,2430	0,5163	0,8807	$0,0983^*$ (0,0050)	$0,7859^*$ (0,0684)	0,6163	0,2270	0,3892	0,4503
	BPV	$0,0847^*$ (0,0028)	$0,5279^*$ (0,0177)	0,7643	0,2390	0,5254	0,9818	$0,0549^*$ (0,0015)	$0,7711^*$ (0,0491)	0,5826	0,2159	0,3668	0,4463
	RR	$0,0589^*$ (0,0055)	$0,6154^*$ (0,0176)	0,8139	0,2477	0,5662	0,7018	$0,0658^*$ (0,0042)	$0,7484^*$ (0,0471)	0,7076	0,2090	0,4986	0,3029
	TTS	$0,1054^*$ (0,0073)	$0,6069^*$ (0,0202)	0,8106	0,2438	0,5668	0,7152	$0,0868^*$ (0,0053)	$0,8276^*$ (0,0635)	0,7015	0,2254	0,4761	0,4102
	Kernel	$0,0973^*$ (0,0075)	$0,6059^*$ (0,0183)	0,8066	0,2471	0,5595	0,7300	$0,0946^*$ (0,0059)	$0,7952^*$ (0,0626)	0,7072	0,2196	0,4876	0,3886
	Mean	0,0964	0,5846	0,7909	0,2441	0,5468	0,8019	0,0801	0,7856	0,6630	0,2194	0,4437	0,3997

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0908^*$ (0,0068)	$0,7342^*$ (0,0232)	0,8689	0,3933	0,4756	0,2580	$0,0574^*$ (0,0068)	$0,8782^*$ (0,0619)	0,8020	0,3909	0,4111	0,1783
	BPV	$0,0529^*$ (0,0036)	$0,7274^*$ (0,0211)	0,8505	0,3857	0,4648	0,2907	$0,0306^*$ (0,0022)	$0,8716^*$ (0,0407)	0,7948	0,3841	0,4106	0,1768
	RR	$0,0477^*$ (0,0053)	$0,7649^*$ (0,0203)	0,8951	0,4035	0,4916	0,2120	$0,0425^*$ (0,0063)	$0,8352^*$ (0,0569)	0,8418	0,3748	0,4670	0,1269
	TTS	$0,0824^*$ (0,0073)	$0,7537^*$ (0,0231)	0,8885	0,3937	0,4949	0,2195	$0,0537^*$ (0,0072)	$0,8821$ (0,0619)	0,8225	0,3912	0,4313	0,1613
	Kernel	$0,0731^*$ (0,0072)	$0,7629^*$ (0,0204)	0,9032	0,3958	0,5074	0,2152	$0,0620^*$ (0,0094)	$0,8558^*$ (0,0723)	0,8420	0,3806	0,4614	0,1575
	Mean	0,0694	0,7486	0,8812	0,3944	0,4869	0,2391	0,0492	0,8646	0,8206	0,3843	0,4363	0,1602

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:25	$0,1373^*$ (0,0063)	$0,5558^*$ (0,0213)	0,7551	0,2288	0,5262	1,0248	08:20	09:50	$0,0827^*$ (0,0054)	$0,8297^*$ (0,0613)	0,7289	0,2929	0,4359	0,2782
	BPV	09:30	10:10	$0,0965^*$ (0,0029)	$0,4770^*$ (0,0198)	0,7220	0,1748	0,5472	1,8268	08:20	10:10	$0,0331^*$ (0,0021)	$0,8600^*$ (0,0406)	0,7929	0,3566	0,4363	0,2029
	RR	09:35	09:55	$0,0739^*$ (0,0043)	$0,6026^*$ (0,0164)	0,7173	0,1074	0,6099	1,1831	08:20	08:55	$0,0775^*$ (0,0040)	$0,6826^*$ (0,0491)	0,6529	0,1442	0,5087	0,5470
	TTS	09:30	10:25	$0,1091^*$ (0,0073)	$0,5876^*$ (0,0201)	0,8006	0,2287	0,5719	0,8252	08:20	09:20	$0,0868^*$ (0,0053)	$0,8276^*$ (0,0635)	0,7015	0,2254	0,4761	0,4102
	Kernel	09:30	09:55	$0,1398^*$ (0,0070)	$0,5155^*$ (0,0176)	0,7157	0,1131	0,6026	1,6918	08:20	09:35	$0,0885^*$ (0,0065)	$0,8084^*$ (0,0637)	0,7469	0,2507	0,4962	0,3042
	Mean	9:31	10:10	0,1113	0,5477	0,7421	0,1706	0,5716	1,3103	8:20	9:34	0,0737	0,8016	0,7246	0,2540	0,4706	0,3485

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	16:15	09:10	$0,8072^*$ (0,0032)	$-1,2949^*$ (0,2102)	0,6930	0,0000	0,6930	2,7760	20:40	08:45	$0,0487^*$ (0,0156)	$0,9250$ (0,1048)	0,7260	0,1351	0,5909	0,7084
	BPV	16:15	09:20	$0,0638^*$ (0,0031)	$1,0269$ (0,0298)	0,7030	0,0000	0,7030	1,9555	18:35	08:30	$0,0392^*$ (0,0040)	$0,8176^*$ (0,0563)	0,6507	0,0504	0,6003	0,6214
	RR	16:45	07:45	$0,0864^*$ (0,0042)	$1,0091$ (0,0280)	0,7021	0,0000	0,7021	4,9876	02:25	08:35	$0,0670^*$ (0,0076)	$0,7118^*$ (0,0667)	0,7267	0,0768	0,6499	0,5460
	TTS	16:15	09:30	$0,3961^*$ (0,0149)	$0,1722^*$ (0,1626)	0,7130	0,0000	0,7130	1,7532	20:30	08:45	$0,0398^*$ (0,0145)	$0,9586$ (0,0960)	0,7564	0,1330	0,6233	0,6722
	Kernel	01:25	08:30	$0,1074^*$ (0,0067)	$1,0048$ (0,0277)	0,7092	0,0000	0,7092	2,9517	00:25	08:35	$0,0466^*$ (0,0174)	$0,9292$ (0,1028)	0,7352	0,0946	0,6406	0,9634
	Mean	13:23	8:51	0,2922	0,3836	0,7041	0,0000	0,7041	2,8848	12:31	8:38	0,0483	0,8685	0,7190	0,0980	0,6210	0,7023

(j) Statistics for the maximum based on R^2_{marg} over 24 hour trade.

Table 5: Overall Moving Average Seasonal variance forecasting statistics for S&P500 and US30 futures. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

Being able to forecast the dynamics of a time series as caught in R^2_{MAD} is however not sufficient to create reliable forecasts. One can for instance always forecast twice the real value and subsequently get a R^2_{MAD} of 1 whereas the forecasts are clearly biased and therefore unreliable, see Appendix B: Mincer-Zarnowitz regression coefficients. Reviewing MZ α and β statistics from table 5 one can conclude variance forecasts are far from unbiased resulting in big (H)MSPE, starting at triple or more the amount compared to the Mincer-Zarnowitz scaled forecasts, yet swiftly declining as more intra-day data is obtained. Interestingly β is always estimated smaller than one and α always bigger than zero, meaning forecasts for highly volatile days are most often understated whereas forecasts for calm days are in general overstated. Possible reason can be found in extreme observations which, in some numbers, can even adjust GLS estimates to be drawn downward. To see this note that extreme observations only have little influence on the regression slope coefficient as the number of smaller observations is generally far larger. Extreme predictions on the other hand cause the estimates to be biased downward.

Lastly adding overnight information the story remains much the same as for begin-of-the-day variance. Figures are therefore omitted yet are available upon request.

4.1.4 Exponentially Weighted Moving Average Seasonal

The Exponentially Weighted Moving Average was added to the set as to be faster adjusting to changing diurnal patterns as a result of higher or lower return volatility regimes. This could lead to better estimation as the seasonal is to be estimated more accurately yielding more accurate scaling factors and thus forecasts. Confiding table 6 such an approach does produces higher R^2_{MAD} statistics for S&P500 data than is the case for an equally weighted seasonal. Moreover HMSPE decreases. For US30 no significant improvements are observed on R^2_{MAD} and HMSPE actually increases little.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2522^*$ (0,0048)	$0,6398^*$ (0,1047)	0,4882	0,0730	0,4152	3,1544	$0,1445^*$ (0,0052)	$0,6260^*$ (0,1381)	0,3877	0,0978	0,2899	4,4727
	BPV	$0,1599^*$ (0,0026)	$0,4178^*$ (0,0886)	0,4655	0,0701	0,3954	3,3030	$0,0800^*$ (0,0031)	$0,6138^*$ (0,0879)	0,3334	0,0869	0,2465	3,5519
	RR	$0,0953^*$ (0,0051)	$0,5344^*$ (0,0180)	0,6697	0,0690	0,6008	1,5152	$0,0902^*$ (0,0028)	$0,5141^*$ (0,0357)	0,5124	0,0768	0,4356	2,4193
	TTS	$0,1755^*$ (0,0060)	$0,5326^*$ (0,0224)	0,6026	0,0692	0,5334	1,8954	$0,1138^*$ (0,0044)	$0,6624^*$ (0,0674)	0,4761	0,0945	0,3816	3,6123
	Kernel	$0,1726^*$ (0,0063)	$0,4831^*$ (0,0187)	0,6862	0,0755	0,6107	2,2769	$0,1235^*$ (0,0033)	$0,5836^*$ (0,0494)	0,5248	0,0946	0,4302	3,5933
	Mean	0,1711	0,5215	0,5824	0,0714	0,5111	2,4290	0,1104	0,6000	0,4469	0,0901	0,3568	3,5299

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1690^*$ (0,0043)	$0,5858^*$ (0,0274)	0,6780	0,1265	0,5515	1,3651	$0,1149^*$ (0,0071)	$0,7821$ (0,1504)	0,5532	0,1498	0,4034	1,2725
	BPV	$0,1047^*$ (0,0020)	$0,5076^*$ (0,0252)	0,6624	0,1308	0,5316	1,7526	$0,0713^*$ (0,0023)	$0,6308^*$ (0,0885)	0,4762	0,1385	0,3377	1,1790
	RR	$0,0687^*$ (0,0043)	$0,6052^*$ (0,0157)	0,7403	0,1256	0,6147	0,8357	$0,0814^*$ (0,0038)	$0,6000^*$ (0,0462)	0,5948	0,1305	0,4643	0,8283
	TTS	$0,1258^*$ (0,0054)	$0,5782^*$ (0,0176)	0,7582	0,1286	0,6297	1,0038	$0,1017^*$ (0,0053)	$0,7405^*$ (0,0785)	0,5641	0,1481	0,4161	1,1015
	Kernel	$0,1271^*$ (0,0065)	$0,5647^*$ (0,0175)	0,7721	0,1308	0,6413	1,0353	$0,1132^*$ (0,0050)	$0,6698^*$ (0,0633)	0,6409	0,1448	0,4961	1,0718
	Mean	0,1191	0,5683	0,7222	0,1285	0,5938	1,1985	0,0965	0,6846	0,5658	0,1423	0,4235	1,0906

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1295^*$ (0,0057)	$0,5970^*$ (0,0204)	0,7838	0,2430	0,5407	0,6737	$0,0993^*$ (0,0044)	$0,7276^*$ (0,0597)	0,6614	0,2270	0,4344	0,5095
	BPV	$0,0818^*$ (0,0028)	$0,5565^*$ (0,0182)	0,7637	0,2390	0,5247	0,7636	$0,0588^*$ (0,0015)	$0,6656^*$ (0,0480)	0,6147	0,2159	0,3988	0,5085
	RR	$0,0576^*$ (0,0051)	$0,6359^*$ (0,0170)	0,8407	0,2477	0,5930	0,5127	$0,0684^*$ (0,0040)	$0,7018^*$ (0,0448)	0,7057	0,2090	0,4967	0,3448
	TTS	$0,1011^*$ (0,0068)	$0,6338^*$ (0,0198)	0,8336	0,2438	0,5898	0,5301	$0,0869^*$ (0,0048)	$0,7797^*$ (0,0578)	0,6853	0,2254	0,4599	0,4533
	Kernel	$0,0934^*$ (0,0069)	$0,6305^*$ (0,0176)	0,8271	0,2471	0,5799	0,5331	$0,0929^*$ (0,0054)	$0,7646^*$ (0,0566)	0,7182	0,2196	0,4986	0,4354
	Mean	0,0927	0,6108	0,8098	0,2441	0,5656	0,6026	0,0813	0,7278	0,6771	0,2194	0,4577	0,4503

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0887^*$ (0,0063)	$0,7621^*$ (0,0225)	0,8888	0,3933	0,4955	0,2013	$0,0610^*$ (0,0059)	$0,8283^*$ (0,0544)	0,8240	0,3909	0,4332	0,2041
	BPV	$0,0539^*$ (0,0035)	$0,7471^*$ (0,0216)	0,8787	0,3857	0,4930	0,2282	$0,0323^*$ (0,0020)	$0,8294^*$ (0,0368)	0,8046	0,3841	0,4204	0,1937
	RR	$0,0481^*$ (0,0049)	$0,7851^*$ (0,0197)	0,9120	0,4035	0,5085	0,1600	$0,0457^*$ (0,0059)	$0,7983^*$ (0,0538)	0,8486	0,3748	0,4738	0,1430
	TTS	$0,0817^*$ (0,0068)	$0,7789^*$ (0,0227)	0,9095	0,3937	0,5159	0,1680	$0,0569^*$ (0,0065)	$0,8387^*$ (0,0563)	0,8461	0,3912	0,4549	0,1816
	Kernel	$0,0729^*$ (0,0068)	$0,7854^*$ (0,0199)	0,9122	0,3958	0,5164	0,1632	$0,0651^*$ (0,0085)	$0,8155^*$ (0,0656)	0,8447	0,3806	0,4641	0,1777
	Mean	0,0691	0,7717	0,9002	0,3944	0,5059	0,1841	0,0522	0,8220	0,8336	0,3843	0,4493	0,1800

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:05	$0,1604^*$ (0,0050)	$0,5242^*$ (0,0219)	0,7189	0,1592	0,5597	1,4717	08:20	10:10	$0,0654^*$ (0,0055)	$0,8125^*$ (0,0521)	0,8126	0,3640	0,4486	0,2305
	BPV	09:30	10:20	$0,0834^*$ (0,0029)	$0,5375^*$ (0,0187)	0,7647	0,2090	0,5557	1,0434	08:30	10:05	$0,0436^*$ (0,0016)	$0,7420^*$ (0,0353)	0,7766	0,3403	0,4363	0,2227
	RR	09:35	09:55	$0,0729^*$ (0,0041)	$0,6226^*$ (0,0162)	0,7582	0,1074	0,6508	0,8754	08:35	09:20	$0,0719^*$ (0,0039)	$0,7749^*$ (0,0503)	0,7202	0,2090	0,5112	0,2344
	TTS	09:30	10:00	$0,1258^*$ (0,0054)	$0,5782^*$ (0,0176)	0,7582	0,1286	0,6297	1,0038	08:30	10:00	$0,0642^*$ (0,0052)	$0,8605^*$ (0,0494)	0,7939	0,3233	0,4706	0,2052
	Kernel	09:30	10:00	$0,1271^*$ (0,0065)	$0,5647^*$ (0,0175)	0,7721	0,1308	0,6413	1,0353	08:20	09:25	$0,0908^*$ (0,0054)	$0,7705^*$ (0,0546)	0,7419	0,2301	0,5118	0,4027
	Mean	9:31	10:04	0,1139	0,5655	0,7544	0,1470	0,6074	1,0859	8:27	9:48	0,0672	0,7921	0,7690	0,2933	0,4757	0,2591

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	19:30	09:30	$0,1159^*$ (0,0069)	$1,1076^*$ (0,0376)	0,7193	0,0000	0,7193	1,2003	20:40	08:45	$0,0487^*$ (0,0156)	$0,9250$ (0,1048)	0,7260	0,1351	0,5909	0,7084
	BPV	16:30	09:20	$0,0681^*$ (0,0030)	$1,0390$ (0,0308)	0,7008	0,0000	0,7008	1,8952	18:35	08:30	$0,0392^*$ (0,0040)	$0,8176^*$ (0,0563)	0,6507	0,0504	0,6003	0,6214
	RR	18:40	09:25	$0,0899^*$ (0,0042)	$1,1129^*$ (0,0314)	0,7407	0,0000	0,7407	1,1891	02:25	08:35	$0,0670^*$ (0,0076)	$0,7118^*$ (0,0667)	0,7267	0,0768	0,6499	0,5460
	TTS	16:25	08:25	$0,8631^*$ (0,0026)	$-1,5969^*$ (0,1296)	0,7309	0,0000	0,7309	3,2103	20:30	08:45	$0,0398^*$ (0,0145)	$0,9586$ (0,0960)	0,7564	0,1330	0,6233	0,6722
	Kernel	16:30	08:30	$0,7758^*$ (0,0037)	$-1,0639^*$ (0,1779)	0,7505	0,0000	0,7505	3,9307	00:25	08:35	$0,0466^*$ (0,0174)	$0,9292$ (0,1028)	0,7352	0,0946	0,6406	0,9634
	Mean	17:31	9:02	0,3826	0,1197	0,7284	0,0000	0,7284	2,2851	12:31	8:38	0,0483	0,8685	0,7190	0,0980	0,6210	0,7023

(j) Statistics for the maximum based on R^2_{marg} over 24 hour trade.

Table 6: Overall Exponentially Weighted Moving Average Seasonal variance forecasting statistics for S&P500 and US30 futures. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

4.1.5 Fourier Flexible Form

In addition to the diurnal pattern constructed from simple (weighted) average past volatility, we estimate a Fourier Flexible Form which will be fitted to the seasonal pattern. This form is made dependent on begin-of-the-day volatility via σ in (25) as to enable the form to adjust for different diurnal shapes during high and low volatility days. Moreover it is made as to cope with recurring macroeconomic news announcements at 8:30 and 10:00 A.M. EST. Supplemented with earlier discussed parameter values $\{H, J, P, D\}$ this leads to the frontier found in 13 and according statistics in table 7. Note that the FFF seasonal was estimated over open outcry trading hours only, therefore we are unable to use this seasonal to construct forecasts outside of these hours. Alternatively one could try fitting the seasonal on data including overnight information, though we chose not to as 24 hour seasonal estimation would most likely result in worse seasonals during more important open outcry trading hours. Moreover a clear diurnal pattern for overnight information seems absent implying that estimating a pattern would result in arbitrariness.

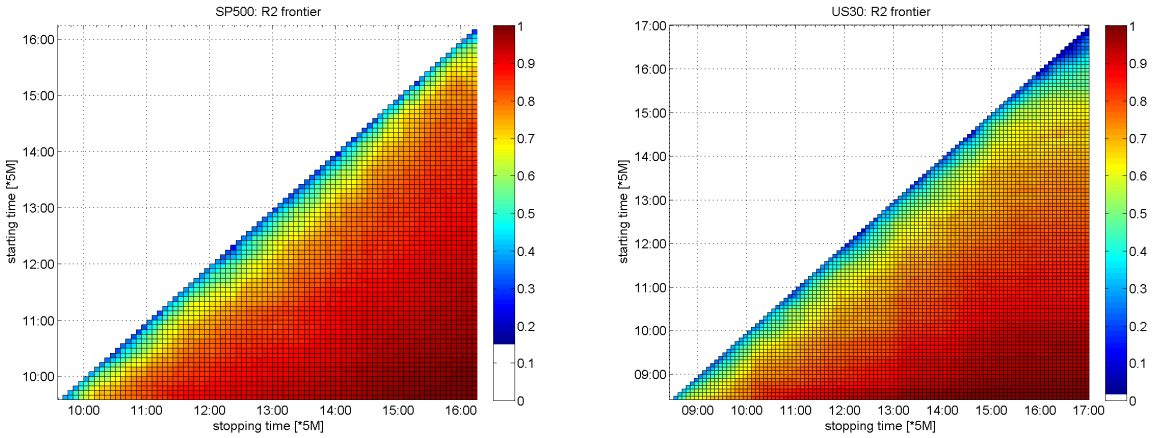

(a) S&P500: R^2_{MAD} **(b)** US30: R^2_{MAD}

Figure 13: Figures display squared correlations of forecasts and estimated daily variance figures, i.e. R^2_{MAD}, for S&P500 as well as US30 futures return data using Realized Variance as volatility proxy and taking H=22. Where H=22 is taken as for the calculations to be feasible. H=252 is computationally to expensive to create such graphs. Note that in both cases half of the graph takes the lowest possible point or is left blank as the start of observations can never overtake the stopping time.

Figures 1 and 2 already demonstrated the important influence of begin-of-the-day volatility on the diurnal pattern. Explicitly allowing the seasonal to be a function hereof has indeed some considerable effects. From table 7 it can be seen that S&P500 R^2_{MAD} statistics on average only improve little or even decrease some as compared to other forecasting methods. Though R^2_{MAD} statistics keep on par with earlier methods, bias decreases considerably as MZ α and β estimates are well closer their optimal values and HMSPE even halves. US30 data, especially using H=22, also benefits seasonal parametrization by FFF. With R^2_{MAD} increasing on average 0.07 over the best alternative and decreased HMSPE, this could well be pointed a significant all-round improvement.

As for the sensitivities to parameters J and P, their influence is generally small as long as $J \geq 1$ and $P \geq 1$. Taking $J = 0$ would undo the advantages to this method, principally relegating FFF to a Seasonal Moving Average method, and thereby obtaining similar performance. Whereas $P = 0$ leads to the elimination of the (co)sine function and thereby a huge part of the ability to fit the seasonal pattern. Other way around, increasing P and J does not yield much advantage. This might be due to the fact that it makes the seasonal somewhat more susceptible to noise and creates additional parameters to be estimated, possibly being redundant as the general seasonal shape is not so complex as to require the additional freedom.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2579^*$ (0,0056)	$0,7978$ (0,1054)	0,4828	0,0730	0,4098	1,0225	$0,1369^*$ (0,0078)	$0,6120^*$ (0,1196)	0,3824	0,0978	0,2846	2,0177
	BPV	$0,1608^*$ (0,0030)	$0,6455^*$ (0,0951)	0,4748	0,0701	0,4047	0,8851	$0,0799^*$ (0,0045)	$0,5733^*$ (0,0807)	0,3790	0,0869	0,2921	1,3385
	RR	$0,0691^*$ (0,0065)	$0,8224^*$ (0,0271)	0,6418	0,0690	0,5728	0,5186	$0,0755^*$ (0,0042)	$0,6229^*$ (0,0403)	0,5376	0,0768	0,4608	0,7617
	TTS	$0,1544^*$ (0,0076)	$0,7562^*$ (0,0301)	0,6159	0,0692	0,5466	0,7291	$0,1098^*$ (0,0069)	$0,5941^*$ (0,0625)	0,5016	0,0945	0,4071	1,8024
	Kernel	$0,1527^*$ (0,0079)	$0,6798^*$ (0,0256)	0,6345	0,0755	0,5590	0,8452	$0,1160^*$ (0,0050)	$0,5465^*$ (0,0456)	0,5351	0,0946	0,4405	1,6783
	Mean	0,1590	0,7403	0,5700	0,0714	0,4986	0,8001	0,1036	0,5898	0,4671	0,0901	0,3770	1,5197

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1588^*$ (0,0051)	$0,7732^*$ (0,0311)	0,6467	0,1265	0,5201	0,5690	$0,1039^*$ (0,0124)	$0,6382^*$ (0,1223)	0,5483	0,1498	0,3985	1,0452
	BPV	$0,1009^*$ (0,0026)	$0,7170^*$ (0,0308)	0,6426	0,1308	0,5117	0,6262	$0,0806^*$ (0,0038)	$0,4043^*$ (0,0698)	0,4637	0,1385	0,3252	0,8512
	RR	$0,0444^*$ (0,0053)	$0,8531^*$ (0,0215)	0,7213	0,1256	0,5958	0,3481	$0,0542^*$ (0,0063)	$0,6826^*$ (0,0479)	0,6344	0,1305	0,5039	0,5041
	TTS	$0,1037^*$ (0,0067)	$0,7764^*$ (0,0231)	0,7142	0,1286	0,5856	0,4783	$0,0815^*$ (0,0095)	$0,6512^*$ (0,0680)	0,6163	0,1481	0,4682	0,9745
	Kernel	$0,0981^*$ (0,0079)	$0,7732^*$ (0,0231)	0,7480	0,1308	0,6172	0,4614	$0,0857^*$ (0,0089)	$0,6328^*$ (0,0583)	0,6526	0,1448	0,5078	0,8540
	Mean	0,1012	0,7786	0,6946	0,1285	0,5661	0,4966	0,0812	0,6018	0,5831	0,1423	0,4407	0,8458

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1129^*$ (0,0071)	$0,7622^*$ (0,0259)	0,7763	0,2430	0,5332	0,3645	$0,0775^*$ (0,0075)	$0,6736^*$ (0,0557)	0,6266	0,2270	0,3996	0,4982
	BPV	$0,0702^*$ (0,0035)	$0,7614^*$ (0,0239)	0,7491	0,2390	0,5101	0,3475	$0,0476^*$ (0,0021)	$0,6657^*$ (0,0366)	0,5952	0,2159	0,3793	0,4534
	RR	$0,0357^*$ (0,0062)	$0,8241^*$ (0,0225)	0,8184	0,2477	0,5707	0,2733	$0,0385^*$ (0,0062)	$0,7745^*$ (0,0457)	0,7268	0,2090	0,5177	0,2738
	TTS	$0,0815^*$ (0,0083)	$0,7908^*$ (0,0253)	0,8016	0,2438	0,5579	0,3097	$0,0600^*$ (0,0081)	$0,7321^*$ (0,0538)	0,7005	0,2254	0,4751	0,4611
	Kernel	$0,0689^*$ (0,0085)	$0,7994^*$ (0,0230)	0,8118	0,2471	0,5647	0,3000	$0,0625^*$ (0,0092)	$0,7294^*$ (0,0555)	0,7039	0,2196	0,4844	0,3896
	Mean	0,0738	0,7876	0,7914	0,2441	0,5473	0,3190	0,0572	0,7150	0,6706	0,2194	0,4512	0,4152

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0685^*$ (0,0077)	$0,8626^*$ (0,0261)	0,8787	0,3933	0,4853	0,1617	$0,0386^*$ (0,0088)	$0,7847^*$ (0,0553)	0,8203	0,3909	0,4294	0,2352
	BPV	$0,0377^*$ (0,0042)	$0,8966^*$ (0,0249)	0,8642	0,3857	0,4785	0,1603	$0,0219^*$ (0,0028)	$0,7895^*$ (0,0357)	0,8123	0,3841	0,4281	0,2293
	RR	$0,0282^*$ (0,0059)	$0,8796^*$ (0,0226)	0,8958	0,4035	0,4922	0,1379	$0,0218^*$ (0,0079)	$0,8324^*$ (0,0547)	0,8667	0,3748	0,4918	0,1525
	TTS	$0,0620^*$ (0,0082)	$0,8630^*$ (0,0259)	0,8903	0,3937	0,4966	0,1437	$0,0311^*$ (0,0095)	$0,8076^*$ (0,0566)	0,8507	0,3912	0,4594	0,2224
	Kernel	$0,0472^*$ (0,0081)	$0,8841^*$ (0,0230)	0,9048	0,3958	0,5090	0,1361	$0,0360^*$ (0,0122)	$0,8100^*$ (0,0672)	0,8537	0,3806	0,4731	0,1852
	Mean	0,0487	0,8772	0,8868	0,3944	0,4923	0,1479	0,0299	0,8048	0,8407	0,3843	0,4564	0,2049

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500								US30							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:15	$0,1266^*$ (0,0070)	$0,7430^*$ (0,0272)	0,7422	0,1963	0,5459	0,4893	08:20	09:50	$0,0622^*$ (0,0076)	$0,7453^*$ (0,0533)	0,7476	0,2929	0,4547	0,2957
	BPV	09:30	10:10	$0,0816^*$ (0,0036)	$0,7508^*$ (0,0272)	0,7161	0,1748	0,5413	0,5071	08:20	10:15	$0,0229^*$ (0,0028)	$0,7857^*$ (0,0359)	0,8041	0,3710	0,4331	0,2401
	RR	09:35	09:55	$0,0485^*$ (0,0049)	$0,8582^*$ (0,0212)	0,7192	0,1074	0,6118	0,3858	08:20	09:15	$0,0404^*$ (0,0069)	$0,7645^*$ (0,0511)	0,7198	0,1968	0,5229	0,2940
	TTS	09:30	09:55	$0,1186^*$ (0,0068)	$0,7796^*$ (0,0249)	0,6999	0,1069	0,5930	0,5082	08:20	09:05	$0,0699^*$ (0,0088)	$0,6815^*$ (0,0597)	0,6888	0,1899	0,4989	0,6368
	Kernel	09:30	10:00	$0,0981^*$ (0,0079)	$0,7732^*$ (0,0231)	0,7480	0,1308	0,6172	0,4614	08:20	08:50	$0,0857^*$ (0,0089)	$0,6328^*$ (0,0583)	0,6526	0,1448	0,5078	0,8540
	Mean	9:31	10:03	0,0947	0,7810	0,7251	0,1432	0,5818	0,4704	8:20	9:27	0,0562	0,7220	0,7226	0,2391	0,4835	0,4641

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

Table 7: Overall Fourier Flexible Form seasonal forecasting statistics for S&P500 and US30 futures. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level. Note that a whole grid of regressions was computationally to expensive using H=252 therefore start has been locked at 09:30.

4.1.6 Random Walk (benchmark)

To check whether forecasts improve upon the most basic of models a Random Walk was added as benchmark. Table 8 outlines the performance statistics which can be obtained by such method.

	S&P500						US30					
	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	$0,1266^*$ (0,0115)	$0,8581^*$ (0,0416)	0,6790	0,0000	0,6790	1,0843	$0,0859^*$ (0,0097)	$0,7880^*$ (0,0636)	0,4968	0,0000	0,4968	2,9248
BPV	$0,0847^*$ (0,0077)	$0,8237^*$ (0,0483)	0,6646	0,0000	0,6646	1,5609	$0,0487^*$ (0,0044)	$0,7504^*$ (0,0563)	0,5057	0,0000	0,5057	2,3720
RR	$0,0968^*$ (0,0083)	$0,8283^*$ (0,0379)	0,7200	0,0000	0,7200	1,8937	$0,0574^*$ (0,0070)	$0,7948^*$ (0,0537)	0,5676	0,0000	0,5676	1,3728
TTS	$0,1090^*$ (0,0107)	$0,8880^*$ (0,0373)	0,6928	0,0000	0,6928	0,9449	$0,0895^*$ (0,0113)	$0,7681^*$ (0,0736)	0,5167	0,0000	0,5167	3,2599
Kernel	$0,1269^*$ (0,0116)	$0,8522^*$ (0,0375)	0,7312	0,0000	0,7312	1,4438	$0,0941^*$ (0,0132)	$0,7721^*$ (0,0762)	0,4841	0,0000	0,4841	3,5566
mean	0,1088	0,8501	0,69752	0,0000	0,69752	1,38552	0,0751	0,7747	0,5142	0,0000	0,5142	2,6972

Table 8: Random Walk variance forecasting statistics for the five volatility measures. Standard errors of the estimates are given between brackets and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

4.1.7 GARCH (1,1) model (benchmark)

A second benchmark, the GARCH(1,1) model is used to put the performance statistics in a more global context. However, although this model is designed for volatility forecasting in the presence of heteroskedasticity, it is known to perform questionable dealing with high frequency volatility measures. Reason is that GARCH is slow in catching up, that is it will take many observations to adjust the conditional variance to a new level whereas IV measures adjust instantly. See

for example Andersen and Bollerslev (1997); Andersen et al. (2003). Indeed for S&P500 data, GARCH performs worse that the RW on R^2_{MAD} and HMSPE statistics however not excessive. For US30 data R^2_{MAD} is again lower than RW but though HMSPE is better. Table 9 presents further figures.

	S&P500						US 30					
	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	$-0,1762^*$ (0,0239)	$1,3346^*$ (0,0485)	0,5497	0,0000	0,5497	2,0392	$0,0530^*$ (0,0099)	$0,7812^*$ (0,0506)	0,4259	0,0000	0,4259	0,7374
BPV	$-0,1048^*$ (0,0144)	$0,7914^*$ (0,0292)	0,5783	0,0000	0,5783	9,1545	$0,0212^*$ (0,0048)	$0,4563^*$ (0,0243)	0,4620	0,0000	0,4620	4,5485
RR	$-0,1213^*$ (0,0170)	$0,9763$ (0,0344)	0,5463	0,0000	0,5463	6,7706	$0,0447^*$ (0,0062)	$0,5702^*$ (0,0318)	0,4619	0,0000	0,4619	1,3606
TTS	$-0,1752^*$ (0,0237)	$1,3427^*$ (0,0481)	0,5294	0,0000	0,5294	1,7447	$0,0508^*$ (0,0099)	$0,7810^*$ (0,0507)	0,4401	0,0000	0,4401	0,7237
Kernel	$-0,1816^*$ (0,0241)	$1,4001^*$ (0,0488)	0,5578	0,0000	0,5578	2,2407	$0,0643^*$ (0,0111)	$0,7894^*$ (0,0568)	0,4405	0,0000	0,4405	0,5712
mean	-0,1518	1,1690	0,5523	0,0000	0,5523	4,3899	0,0468	0,6756	0,4461	0,0000	0,4461	1,5883

Table 9: GARCH(1,1) volatility forecasting statistics for the five volatility measures. Standard errors of the estimates are given between brackets and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence interval.

Note that for GARCH parameter estimation an expanding window is used of at least 252 observations, which is rarther small for GARCH estimation. Simulation studies to the effects of small sample sizes on ML estimates in the GARCH(1,1) model by Ng (2006) find that at least 1000 observations should be used for the likelihood function to converge to the global maximum. However given the data sample (especially the challenging periods near the end) and for reasons of comparability to other forecasting methods this smaller estimation period was chosen.

Summarizing, Begin-of-the-day variance forms a high potential regressor for daily volatility. In itself, it has higher correlation to daily variance then the RW after just 30 minutes of observations, equaling about 13% of daily variance. Compared to the GARCH(1,1), 15-20 minutes of observations even suffice to outperform in terms of correlation (VR=8%). Up-scaling these early variance figures to daily proportions using a seasonal shape or Mincer-Zarnowitz bias adjustment method leads to further improvements in terms of R^2_{MAD} as well as HMSPE. The Fourier Flexible Form can herein generally be addressed best seasonal scaling, especially when HMSPE is concerned. Yet scaling by Mincer-Zarnowitz coefficients seems more effective during the start of the day. In terms of efficiency it stands out that roughly the first 15 minutes of observations for S&P500 and first 20 minutes of observations for US30 are most informative. Using a seasonal scaling however 35 resp. 90 minutes of observations are needed for S&P500 and US30. After this period the increase in squared correlation is outweighed by the extra variance observed in making the new forecast.

Interestingly the first few, 0-10 minutes, often lead to worse correlations than under exclusion of such evidently noisy measurements. To my knowledge now earlier research has been conducted on this subject and no rejection be it conformation on this matter has been found. However, the loss in correlation is generally small.

4.2 Kalman Filter and EM convergence

For a forecast that is reliable, with reliable meaning having small MSPE, high R^2 or R^2_{MAD} is not sufficient. Knowledge of the bias and variance of the estimates is needed to for a complete picture. Throughout the day Mincer-Zarnowitz α and β inevitably become better, reaching resp. 0 and 1 at days end. Yet it becomes well clear from above presented tables that forecasts based on the intraday seasonal are, especially during start of the day, far from unbiased. Frijns and Margaritis assign the reason that α and β deviate from 0 and 1 to realized variance being a noisy measure of actual volatility, perhaps causing a downward bias. However, this reasoning might be wrong as the MZ response variable, explanatory variable and seasonal are all built from the same volatility measure. Here a similar yet less strong bias is found but the reason for this lies in the properties of MZ OLS and GLS regressions. As OLS regressions are highly sensitive to outliers and heteroskedasticity, i.e. errors are not from one and the same distribution, estimates are heavily corrupted. To demonstrate the impact see figure 14 where forecasts are plotted against their real value. On average forecasts seem to create a cone like pattern around the 45^{o} angle (striped black line). The red line depicts the OLS regression line and is heavily susceptible to the few extreme observations. The blue line represents the GLS alternative. This, and the fact that small values occur far more often make it that MZ OLS β estimates are by construction downward biased. To see this note that outliers in the response variable only cause little shifts mainly in the constant term, correctly predicted outliers have big influence on β but typically occur just very rarely whereas outliers in the regressors occur more often and thereby cause serious damage to the β estimate. If, as in our case, there is any suspicion of outliers and/or heteroskedasticity in the MZ errors one should therefore seriously consider outlier and heteroskedasticity robust estimation procedures.

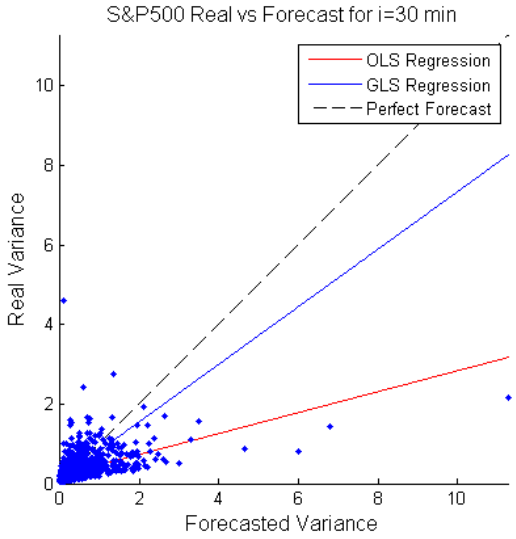

Figure 14: MZ regression estimates get heavily corrupted by outliers and heteroskedasticity. Plotted figure shows S&P500 daily Realized Variance plotted against FFF variance forecasts using RV measure. Additional OLS and GLS regression lines are displayed along a 45^{o} perfect forecast line.

Nonetheless, trusting GLS figures some bias remains. If this bias is implicitly caused by the forecasting method or other consistent source, we might again be able use the fitted Mincer-Zarnowitz bias coefficients and subsequently use them to bias adjust forecasts. In fact using the Kalman Filter we are able to reproduce MZ α and β coefficients and create proper forecasts for these states: α_{t+1} and β_{t+1}. Indeed such adjustments seem possible, generating higher R^2_{MAD} and lower HMSPE.

Before we endeavor such an approach let us first again substantiate the claim with a proof of concept. Again using the earlier mentioned 20 day rolling window Mincer-Zarnowitz regressions, α and β estimates are obtained over the whole sample as approximations to their unobserved states. These estimates are subsequently used to recreate the 21st day forecast, i.e. $RV^{n^*}_{t,scaled} = \alpha_{[t-20,t-1]} + \beta_{[t-20,t-1]} RV^{n^*}_t$, resulting in the graphical adjustments and accompanying statistics found in figure 15. Overall improvements over MA seasonal in terms of bias and variance explained are undeniable, though on par with earlier conducted direct MZ scaling (see section 4.1.2). Note that the difference between these two lies in the additional use of a seasonal scaling before proceeding to MZ bias adjustment.

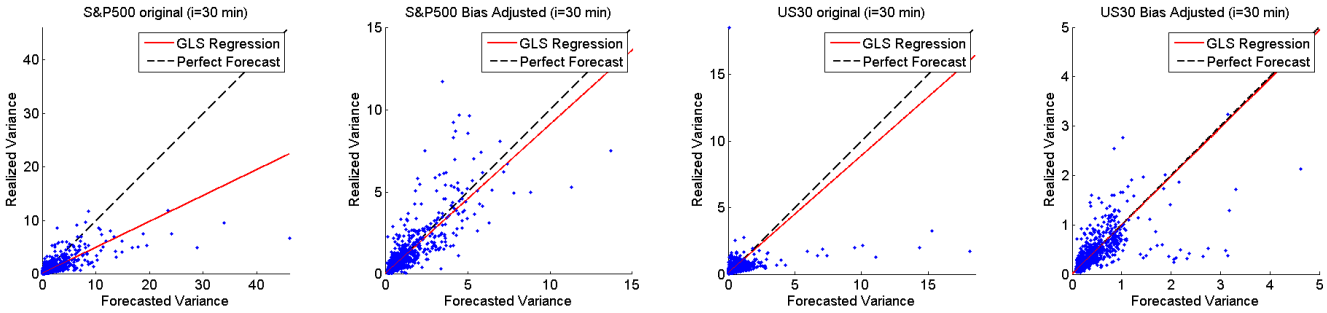

	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	$MSPE$	$HMSPE$
Normal forecasts S&P500	$0,1972^*$ (0,0045)	$0,4866^*$ (0,0261)	0,6344	4,5622	1,6366
Bias adjusted forecasts S&P500	$0,0456^*$ (0,0113)	$0,9052^*$ (0,0320)	0,7101	0,5703	0,7899
Normal forecasts US30	$0,1149^*$ (0,0083)	$0,8670$ (0,1780)	0,5273	1,0696	1,1071
Bias adjusted forecasts US30	$0,0095$ (0,0087)	$0,9869$ (0,0475)	0,7214	0,3159	0,5154

Figure 15: Normal and bias adjusted forecasts and statistics for S&P500 (omitting Sept, Oct and Nov 2008) and US30 after 30 minutes of begin-of-the-day returns. A proof of concept considering the possibility to adjust regular forecasts using Mincer-Zarnowitz α and β estimates. Forecasts are based on AVG Seasonal using the RV measure and H=252.

As the concept seems to work we basically shrink the 20 rolling window to a single observation. Obtaining α and β estimates from regression would in such case be impractical, yet Kalman Filtering with the Expectation Maximization algorithm and encapsulated Maximum Likelihood Estimation helps unveiling the unobserved α and β states. Taking the preferred parametrizations, such that F and Q are diagonal matrices and R a full matrix with $\sigma^2_{ij} = \sigma^2_{ji} = \sigma^2_{ii}$ for $i = 1, ..., n$, and using earlier calculated starting values, stable convergence is obtained for S&P500 (omitting Sept, Oct and Nov 2008) and US30 data series. Note that the great heteroskedasticity during Sept, Oct, Nov 2008 return observations (Highest volatility ever recorded) seriously degrade, even diverge ME parameter estimation, thereby jeopardizing Kalman Filter performance. Reason lies in the extreme observations that do not concord the normality assumption underlying Maximum Likelihood Estimation. Consequently the Newton Raphson algorithm swings wild creating spurious estimates which in turn damage further Kalman Filtering precedure. Therefore S&P500 data series omitting this interval are considered in this section. Figures can subsequently be compared to the S&P500 FFF forecasts statistics also omitting Sept, Oct, Nov 2008 which are to be found in Appendix C. For US30 series a mere single observation was excluded from the sample regarding extreme volatility on March 18, 2009. During this day the Federal Open Market Committee (FOMC) released a press announcement that the Federal Reserve would purchase an additional $750 billion of agency MBS, an additional $100 billion in agency debt, and $300 billion of longer-

term Treasury securities, consequently provoking an extreme upwards jump in the US 30 year treasury bond futures. In terms of the robust R^2_{MAD} and HMSPE this single observation has little influence. Within the Maximum Likelihood estimation such an extreme returns is however highly unlikely and again damages parameter estimation as well as further kalman filtering.

As the likelihood has converged, parameters should also have converged to their ML estimates. A quick check on the estimates of F learns that all AR(1) coefficients are indeed as required between -1 and 1, guarding the stability of the filter. With the conspicuousness though that most φ_{2,n^*} estimates are very close to 1, implying the AR models for beta are rather close to a random walk whereas φ_{1,n^*} estimates are wider spread over the interval $[-1, 1]$, see table 10. Interestingly φ_{2,n^*} estimated values are consistent with first order autocorrelation coefficients found over the 20 day rolling window experiment, and are therefore approximately equal to the starting values to the algorithm. Therefor we recheck the convergence using with random initial values. Approximately the same ending values can then be found.

Subsequently the hidden states α and β can be reviewed for every half an hour throughout the day. The correlation coefficient to daily variance is about 0.4 and visual interpretation of figure 16 lies bare a predominant link between the two as was to be expected. In times of great volatility forecasts fail to perform and consequently yield greater bias as supported by the greater α and more erratic β states.

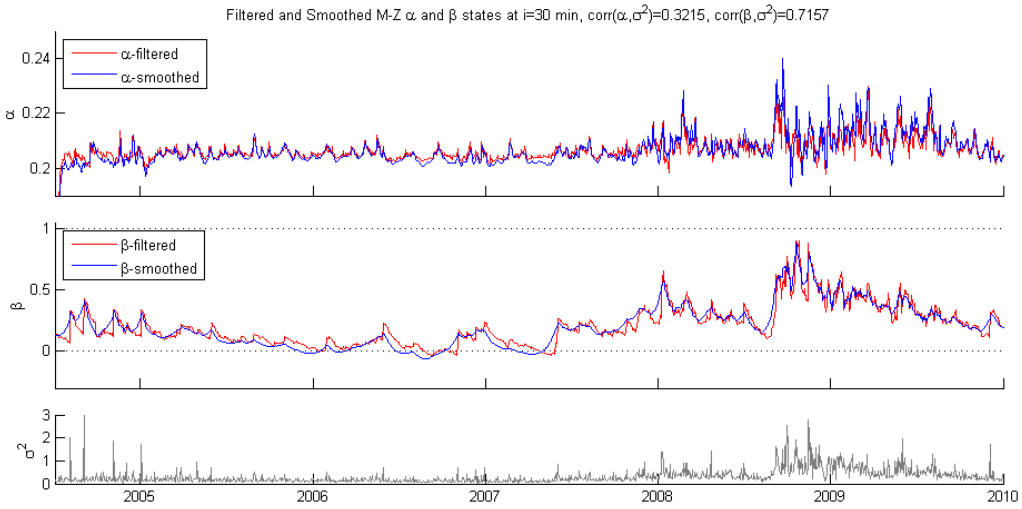

Figure 16: Filtered and smoothed Mincer-Zarnowitz α and β states $\left(\hat{\alpha}_{t|T}, \hat{\beta}_{t|T}\right)$ concerning US30 data on i=30 minutes after start of open outcry markets as obtained from Kalman filtering/smoothing. States are based on FFF forecasts using H=252, the Realized Variance measure and omitting March 18, 2009. Additionally correlations of smoothed states to daily variance are given.

Again using these states, that is the ex-ante predicted states for every time instance $t = 1, ..., T$ - not the ex-post smoothed states, to bias adjust forecasts produces the figures found in table 10. Sub table a-d) displays some statistics on the alpha and beta states. From here it is clear that mean alpha and beta levels come closer to their ending values 0 and 1 as expected. More interesting statistics can be found in sub table e-h). For US30 overall R^2_{MAD} but in particular MSPE and HMSPE seem to experience great benefit from bias adjustments. MSPE/HMSPE is brought back with a major factor 6 at the start naturally declining to a factor 2 at two hours in. Favorably biggest improvements in correlation and bias can thus be found near the start of open outcry markets. As was to be expected since bias is greatest near the start.

For S&P500 smaller but improvements nonetheless can be observed in MSPE and HMSPE but not in R^2_{MAD}. Correlations actually seem to decrease. Reason here fore might lie in the greater volatility of this time series experiences, being inconsistent with the assumption of normality. It would therefore be interesting to check performance under less strict assumptions as possible with a particle filter, imposing a more flexible error distribution.

To visualize the bias adjustments and their forecasting performance figure 17 plots the forecasts against realized values and compares them to the original ones (original in subfig-a and corresponding bias adjusted directly beneath in subfig-b). It can be seen from figure 17 that forecasts are indeed tighter pact around 45^o black dashed line of perfect forecast, implying a decrease of MSPE and HMSPE. As indeed supported by table 10.

Promising as these figures are, it should for a start be noted that this is just a first exploration to Mincer-Zarnowitz bias adjustments. Consequently more research should be conducted to test the sustainability of such results in other timeseries and settings. Though if such results can be confirmed, bias adjusting forecasts in such a manner could be performed in the widest of settings. Second it should be noted that similar, perhaps better performance can be much simpler obtained via direct MZ scaling of begin-of-the-day variance as conducted in section 3.2.2 or through scaling by rolling window MZ regression coefficients as done in figure 15.

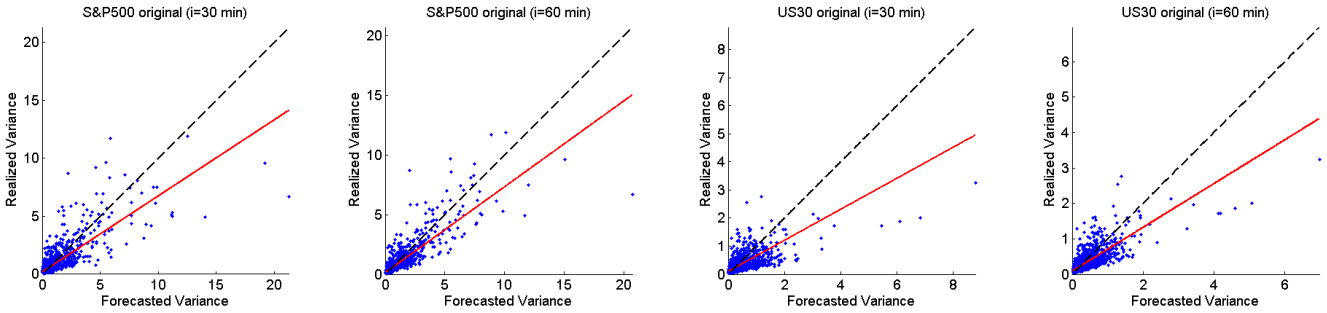

(a) Realized Variance vs Original FFF Variance Forecasts

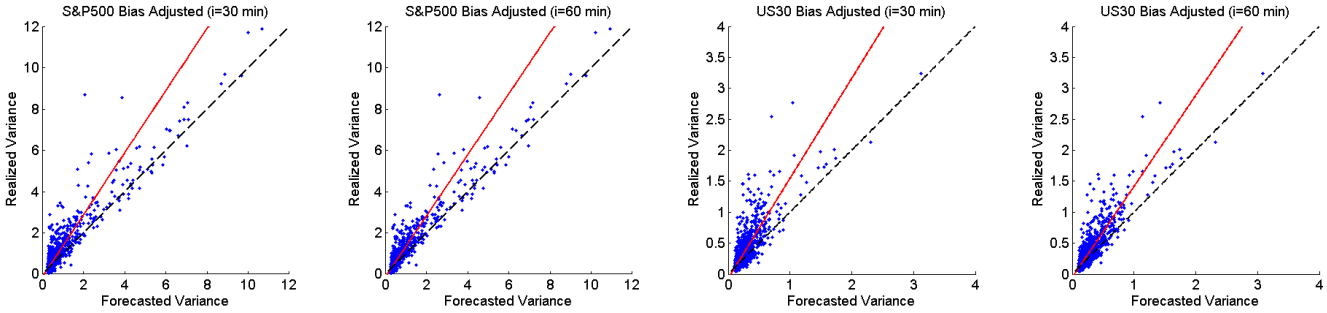

(b) Realized Variance vs Bias Adjusted FFF Variance Forecasts

Figure 17: Figures comparing realized variance vs. original variance forecasts & realized variance vs. bias adjusted variance forecasts. Figures where made for S&P500 and US30 data using FFF forecasts for 30 resp. 60 minutes using the Realized Variance measure. Kalman adjusted forecasts evidently improve upon original forecasts as they are far tighter gathered around the 45^{o} line of perfect forecast.

		S&P500						US30					
		Alpha			Beta			Alpha			Beta		
		C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$	C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$
15 minutes	RV	0,4084	0,2349	0,5337	0,0013	0,9737	0,0532	0,0591	0,7650	0,2507	0,0015	0,9859	0,0959
	BPV	0,3123	-0,0365	0,3013	0,0023	0,9743	0,0814	0,0497	0,6575	0,1445	0,0012	0,9836	0,0763
	RR	0,2677	0,1589	0,3187	0,0050	0,9654	0,1408	0,0608	0,6713	0,1846	0,0008	0,9911	0,0915
	TTS	0,4192	0,0598	0,4460	0,0051	0,9619	0,1322	0,0715	0,7149	0,2501	0,0008	0,9922	0,0852
	Kernel	0,4055	0,1332	0,4681	0,0034	0,9691	0,1107	0,0859	0,6377	0,2367	0,0010	0,9915	0,1232
	Mean	0,3627	0,1101	0,4135	0,0034	0,9689	0,1037	0,0654	0,6893	0,2133	0,0011	0,9889	0,0944

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500						US30					
		Alpha			Beta			Alpha			Beta		
		C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$	C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$
30 minutes	RV	0,5211	-0,0453	0,4985	0,0039	0,9598	0,0918	0,1065	0,5498	0,2363	0,0021	0,9840	0,1239
	BPV	0,2973	-0,0599	0,2805	0,0034	0,9704	0,1136	0,0506	0,6403	0,1402	0,0015	0,9804	0,0829
	RR	0,1982	0,3577	0,3088	0,0050	0,9665	0,1547	0,0617	0,6551	0,1786	0,0011	0,9891	0,1064
	TTS	0,4012	0,0391	0,4175	0,0063	0,9619	0,1622	0,1205	0,4967	0,2391	0,0015	0,9866	0,1063
	Kernel	0,4511	-0,0163	0,4439	0,0049	0,9670	0,1455	0,1308	0,4146	0,2233	0,0023	0,9858	0,1568
	Mean	0,3738	0,0550	0,3898	0,0047	0,9651	0,1335	0,0940	0,5513	0,2035	0,0017	0,9852	0,1153

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500						US30					
		Alpha			Beta			Alpha			Beta		
		C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$	C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$
60 minutes	RV	0,3805	0,1656	0,4560	0,0046	0,9705	0,1529	0,0617	0,7014	0,2063	0,0035	0,9828	0,2049
	BPV	0,2415	0,0601	0,2570	0,0046	0,9737	0,1747	0,0460	0,6461	0,1298	0,0028	0,9788	0,1323
	RR	0,1608	0,4126	0,2739	0,0063	0,9721	0,2296	0,0524	0,6786	0,1627	0,0020	0,9885	0,1693
	TTS	0,3993	-0,0328	0,3865	0,0058	0,9731	0,2160	0,0646	0,6817	0,2027	0,0029	0,9859	0,2033
	Kernel	0,4203	-0,0558	0,3982	0,0068	0,9687	0,2196	0,0773	0,5929	0,1897	0,0032	0,9873	0,2549
	Mean	0,3205	0,1099	0,3543	0,0056	0,9716	0,1986	0,0604	0,6602	0,1783	0,0029	0,9847	0,1930

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500						US30					
		Alpha			Beta			Alpha			Beta		
		C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$	C_1	φ_1	$\overline{\alpha}$	C_2	φ_2	$\overline{\beta}$
120 minutes	RV	0,3673	0,0205	0,3751	0,0077	0,9733	0,2868	0,0529	0,6868	0,1685	0,0062	0,9802	0,3066
	BPV	0,1111	0,5087	0,2260	0,0082	0,9686	0,2666	0,0324	0,6921	0,1049	0,0052	0,9798	0,2507
	RR	0,4250	-0,9100	0,2225	0,0096	0,9739	0,3619	0,0383	0,7203	0,1366	0,0042	0,9846	0,2706
	TTS	0,1560	0,5363	0,3362	0,0087	0,9729	0,3203	0,0521	0,6950	0,1706	0,0045	0,9853	0,2937
	Kernel	0,4350	-0,2220	0,3549	0,0101	0,9680	0,3007	0,0615	0,6065	0,1561	0,0057	0,9846	0,3550
	Mean	0,2989	-0,0133	0,3030	0,0089	0,9713	0,3073	0,0474	0,6801	0,1473	0,0052	0,9829	0,2953

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$-0,1403^*$ (0,0229)	$1,4924$ (0,0566)	0,3911	0,0728	0,3183	0,5142	$-0,0651^*$ (0,0104)	$1,6214^*$ (0,0585)	0,5168	0,0978	0,4190	0,2124
	BPV	$-0,0760^*$ (0,0131)	$1,4498^*$ (0,0548)	0,3568	0,0697	0,2871	0,5334	$-0,0421^*$ (0,0059)	$1,7215^*$ (0,0617)	0,4709	0,0869	0,3840	0,2803
	RR	$-0,0757^*$ (0,0096)	$1,2366^*$ (0,0303)	0,3872	0,0692	0,3180	0,3670	$-0,0601^*$ (0,0082)	$1,6879^*$ (0,0587)	0,4398	0,0768	0,3630	0,1513
	TTS	$-0,1051^*$ (0,0154)	$1,3151^*$ (0,0374)	0,4029	0,0693	0,3336	0,4081	$-0,0828^*$ (0,0103)	$1,7348^*$ (0,0588)	0,4878	0,0945	0,3933	0,1911
	Kernel	$-0,1041^*$ (0,0147)	$1,2790^*$ (0,0339)	0,4030	0,0756	0,3274	0,3494	$-0,0587^*$ (0,0095)	$1,5213^*$ (0,0499)	0,5818	0,0946	0,4872	0,1433
	Mean	-0,1002	1,3546	0,3882	0,0713	0,3169	0,4344	-0,0618	1,6574	0,4994	0,0901	0,4093	0,1957

(i) Bias adjusted FFF forecasting statistics for start of the day up until 15 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$-0,1637^*$ (0,0188)	$1,5052^*$ (0,0463)	0,4728	0,1269	0,3459	0,4706	$-0,0708^*$ (0,0090)	$1,6137^*$ (0,0506)	0,6059	0,1498	0,4561	0,1897
	BPV	$-0,0910^*$ (0,0110)	$1,4672^*$ (0,0457)	0,4658	0,1308	0,3350	0,4914	$-0,0456^*$ (0,0054)	$1,7187^*$ (0,0559)	0,5282	0,1385	0,3897	0,2642
	RR	$-0,0839^*$ (0,0090)	$1,2537^*$ (0,0285)	0,5342	0,1262	0,4080	0,3360	$-0,0662^*$ (0,0074)	$1,7036^*$ (0,0524)	0,5248	0,1305	0,3943	0,1391
	TTS	$-0,1201^*$ (0,0142)	$1,3274^*$ (0,0344)	0,4900	0,1291	0,3609	0,3655	$-0,0856^*$ (0,0094)	$1,7224^*$ (0,0531)	0,5596	0,1481	0,4115	0,1767
	Kernel	$-0,1152^*$ (0,0139)	$1,2912^*$ (0,0322)	0,5114	0,1311	0,3803	0,3151	$-0,0654$ (0,0086)	$1,5272^*$ (0,0447)	0,6398	0,1448	0,4950	0,1317
	Mean	-0,1148	1,3690	0,4948	0,1288	0,3660	0,3957	-0,0667	1,6571	0,5717	0,1423	0,4293	0,1803

(j) Bias adjusted FFF forecasting statistics for start of the day up until 30 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$-0,1672^*$ (0,0159)	$1,4804^*$ (0,0395)	0,6049	0,2443	0,3606	0,4000	$-0,0577^*$ (0,0068)	$1,4703^*$ (0,0379)	0,6922	0,2270	0,4652	0,1524
	BPV	$-0,0956^*$ (0,0089)	$1,4473^*$ (0,0374)	0,6069	0,2400	0,3669	0,4126	$-0,0447^*$ (0,0044)	$1,6569^*$ (0,0456)	0,6252	0,2159	0,4093	0,2325
	RR	$-0,0832^*$ (0,0080)	$1,2425^*$ (0,0258)	0,6667	0,2495	0,4172	0,2724	$-0,0582^*$ (0,0057)	$1,5881^*$ (0,0399)	0,6828	0,2090	0,4738	0,1164
	TTS	$-0,1201^*$ (0,0128)	$1,3133^*$ (0,0315)	0,6362	0,2455	0,3907	0,3033	$-0,0633^*$ (0,0066)	$1,5020^*$ (0,0373)	0,6681	0,2254	0,4427	0,1390
	Kernel	$-0,1185^*$ (0,0121)	$1,2816^*$ (0,0284)	0,6641	0,2486	0,4155	0,2551	$-0,0510^*$ (0,0061)	$1,3782^*$ (0,0314)	0,7325	0,2196	0,5129	0,1021
	Mean	-0,1169	1,3530	0,6358	0,2456	0,3902	0,3287	-0,0550	1,5191	0,6802	0,2194	0,4608	0,1485

(k) Bias adjusted FFF forecasting statistics for start of the day up until 60 min.

		S&P500						US30					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$-0,1563^*$ (0,0111)	$1,4076^*$ (0,0281)	0,7912	0,3955	0,3957	0,2828	$-0,0505^*$ (0,0046)	$1,3548^*$ (0,0256)	0,8434	0,3909	0,4525	0,0960
	BPV	$-0,0911^*$ (0,0061)	$1,3893^*$ (0,0262)	0,8164	0,3876	0,4288	0,3002	$-0,0357^*$ (0,0027)	$1,4633^*$ (0,0283)	0,8353	0,3841	0,4512	0,1443
	RR	$-0,0756^*$ (0,0064)	$1,2168^*$ (0,0212)	0,8328	0,4059	0,4269	0,1930	$-0,0482^*$ (0,0039)	$1,4441^*$ (0,0274)	0,8411	0,3748	0,4663	0,0804
	TTS	$-0,1149^*$ (0,0097)	$1,2805^*$ (0,0244)	0,8181	0,3958	0,4223	0,2116	$-0,0552^*$ (0,0047)	$1,3925^*$ (0,0266)	0,8572	0,3912	0,4660	0,0926
	Kernel	$-0,1111^*$ (0,0091)	$1,2523^*$ (0,0220)	0,8264	0,3978	0,4286	0,1781	$-0,0453^*$ (0,0044)	$1,2943^*$ (0,0227)	0,8809	0,3806	0,5003	0,0664
	Mean	-0,1098	1,3093	0,8170	0,3965	0,4205	0,2331	-0,0470	1,3898	0,8516	0,3843	0,4673	0,0959

(l) Bias adjusted FFF forecasting statistics for start of the day up until 120 min.

Table 10: Forecasting statistics after Mincer-Zarnowitz type bias adjustments to FFF forecasts on S&P500 (omitting Sept, Oct, Nov 2008) and US30 data. Standard errors of the estimates are given between parenthesis if available and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level. Note that S&P500 forecasts can be compared to table C.1-C.4, to be found in Appendix C.

4.3 Forecast Revisions

In general one has to know the Loss Function in order to make statements about a forecasters efficiency. It could for instance be that clients start mistrusting forecasters whom adjust their predictions rather rapidly, seeming overly sensitive to new information. As such, some degree of smoothing behavior, i.e. $\beta > 0$, could be desirable and as such be perfectly rational. Here however we deal with forecasts from econometric models where no such behavior has deliberately been imposed. Our Loss Function is of quadratic form implied by OLS regression. Optimal value of β in (49) is therefore without doubt 0, meaning consecutive forecasts are uncorrelated and no smoothing or overreaction pervades our predictions. To test whether such is indeed the case we look at α and β from (49) and some general statistics of the revisions.

S&P500/US30 forecast revision statistics over open outcry trade for AVG forecasts (H=252)

		Mean	*StdDev*	*Skewness*	*Kurtosis*	$\hat{\alpha}$	$\hat{\beta}$	R^2	AC_1	AC_2	AC_3
S&P500	RV	-0,0177	0,2943	-1,2900	112,7791	$-0,0157^*$ (0,0009)	$0,1125^*$ (0,0030)	0,0127	$0,1125^*$	$0,1095^*$	$0,1067^*$
	BPV	-0,0115	0,1540	-3,2460	113,1707	$-0,0072^*$ (0,0004)	$0,3716^*$ (0,0028)	0,1381	$0,3716^*$	$0,1141^*$	$0,0878^*$
	RR	-0,0109	0,1306	-1,0410	96,1944	$-0,0097^*$ (0,0004)	$0,1130^*$ (0,0030)	0,0128	$0,1130^*$	$0,1210^*$	$0,1096^*$
	TTS	-0,0121	0,2029	-0,1280	104,9098	$-0,0099^*$ (0,0006)	$0,1812^*$ (0,0029)	0,0328	$0,1812^*$	$0,0669^*$	$0,0759^*$
	Kernel	-0,0200	0,2192	-4,0988	114,9233	$-0,0166^*$ (0,0007)	$0,1720^*$ (0,0030)	0,0296	$0,1719^*$	$0,1725^*$	$0,1225^*$
	Mean	-0,0145	0,2002	-1,9608	108,3955	-0,0118	0,1901	0,0452	0,1901	0,1168	0,1005
US30	RV	-0,0018	0,0453	0,4772	135,7472	$-0,0014^*$ (0,0001)	$0,1808^*$ (0,0026)	0,0327	$0,1808^*$	$0,1255^*$	$0,0918^*$
	BPV	-0,0004	0,0203	2,9147	117,5915	$-0,0003^*$ (0,0001)	$0,3228^*$ (0,0025)	0,1042	$0,3228^*$	$0,0640^*$	$0,0523^*$
	RR	-0,0008	0,0234	2,9144	119,5758	$-0,0006^*$ (0,0001)	$0,1606^*$ (0,0026)	0,0258	$0,1606^*$	$0,1021^*$	$0,0690^*$
	TTS	-0,0009	0,0374	3,3053	136,8181	$-0,0007^*$ (0,0001)	$0,2763^*$ (0,0026)	0,0763	$0,2763^*$	$0,0692^*$	$0,0543^*$
	Kernel	-0,0015	0,0390	2,1005	143,9498	$-0,0012^*$ (0,0001)	$0,1883^*$ (0,0026)	0,0355	$0,1883^*$	$0,1278^*$	$0,0961^*$
	Mean	-0,0011	0,0331	2,3424	130,7365	-0,0008	0,2258	0,0549	0,2258	0,0977	0,0727

Table 11: Forecast revision statistics, distribution as well as autocorrelations for H=22. Standard errors of the estimates are given between parenthesis and a star indicates the parameter to be significantly different from 0 on a 95% confidence interval. Revisions are taken over open outcry trade with start of day volatility beginning at the start of open outcry trade and last observation just before closing at 16:15 and 17:00 hours EST for resp. S&P and US30. For the creation of these statistics the 0,1% largest and negative and positive values were deleted from the sample. These solely consist of October 2008 observations, the most volatile week ever recorder in history. The explosive observations found at this time severely alter the autocorrelation function and revisions distribution properties

SP 500/US30 forecast revision statistics over open outcry trade for FFF forecasts (H=252)

		Mean	*StdDev*	*Skewness*	*Kurtosis*	$\hat{\alpha}$	$\hat{\beta}$	R^2	AC_1	AC_2	AC_3
S&P500	RV	-0,0013	0,1761	5,3635	105,3757	$-0,0013^*$ (0,0005)	$0,0204^*$ (0,0030)	0,0004	$0,0204^*$	$0,0234^*$	$0,0434^*$
	BPV	0,0000	0,0822	4,4138	93,8970	0,0000 (0,0002)	$0,3230^*$ (0,0028)	0,1043	$0,3230^*$	$0,0277^*$	$0,0376^*$
	RR	-0,0015	0,0861	4,4998	103,4734	$-0,0014^*$ (0,0003)	$0,0342^*$ (0,0030)	0,0012	$0,0342^*$	$0,0469^*$	$0,0523^*$
	TTS	-0,0009	0,1376	4,8779	105,6767	$-0,0007$ (0,0004)	$0,1273^*$ (0,0030)	0,0162	$0,1273^*$	$0,0085^*$	$0,0297^*$
	Kernel	-0,0059	0,1423	0,9504	88,1547	$-0,0055^*$ (0,0004)	$0,0798^*$ (0,0030)	0,0064	$0,0798^*$	$0,0816^*$	$0,0646^*$
	Mean	-0,0019	0,1249	4,0211	99,3155	-0,0018	0,1169	0,0257	0,1169	0,0376	0,0455
US30	RV	-0,0004	0,0350	8,7044	144,9330	$-0,0004^*$ (0,0001)	$0,0716^*$ (0,0027)	0,0051	$0,0716^*$	$0,0273^*$	$0,0250^*$
	BPV	0,0002	0,0143	7,2583	97,5274	$0,0001^*$ (0,0000)	$0,2941^*$ (0,0025)	0,0865	$0,2941^*$	$0,0247^*$	$0,0130^*$
	RR	-0,0001	0,0164	7,9270	123,0677	$-0,0001$ (0,0000)	$0,1327^*$ (0,0026)	0,0176	$0,1327^*$	$0,0573^*$	$0,0265^*$
	TTS	-0,0004	0,0299	8,9136	153,2527	$-0,0003^*$ (0,0001)	$0,2089^*$ (0,0026)	0,0436	$0,2089^*$	$0,0330^*$	$0,0155^*$
	Kernel	-0,0006	0,0307	9,2171	172,3894	$-0,0006^*$ (0,0001)	$0,1080^*$ (0,0026)	0,0117	$0,1080^*$	$0,0426^*$	$0,0237^*$
	Mean	-0,0003	0,0253	8,4041	138,2341	-0,0002	0,1631	0,0329	0,1631	0,0370	0,0207

Table 12: Forecast revision statistics, distribution as well as autocorrelations for H=252. Standard errors of the estimates are given between parenthesis and a star indicates the parameter to be significantly different from 0 on a 95% confidence interval. Revisions are taken over open outcry trade with start of day volatility beginning at the start of open outcry trade and last observation just before closing at 16:15 and 17:00 hours EST for resp. S&P and US30. For the creation of these statistics the 0,1% largest and negative and positive values were deleted from the sample. These solely consist of October 2008 observations, the most volatile week ever recorder in history. The explosive observations found at this time severely alter the autocorrelation function and revisions distribution properties

From table 11 and 12 it can be seen that revisions are skewed, leptokurtic, zero mean distributed. $\hat{\alpha}$ is most often significant for Seasonal Moving Average forecast revisions but on average fairly small which is in accordance with the idea that there is no reason to expect a systematic bias in the forecast revisions, see Nordhaus (1987). The leptokurtic properties can mainly be addressed to the presence of somewhat extreme observations caused by rapid forecast adjustments at the start of the day. As begin-of-the-day variance still changes significantly the variance of the end-of-day variance estimates remains high, causing high Kurtosis. Interestingly S&P500 and US30 revisions do not agree on the sign of the skewness using the Season Moving Average forecasts, meaning S&P500 revisions have a greater probability mass for being positive. For US30 it is the other way around and for FFF forecasts both are positively skewed.

The efficiency of forecasts, as Nordhaus cites, is deduced from the capability of incorporating all past information. Forecasts are said to be weakly efficient if they minimize $E\left[\varepsilon_t^2|\mathcal{F}_\mathrm{t}\right]$. Where ε_t is the forecast error and $\mathcal{F}_\mathrm{t}$ the set of all past forecasts. Under weak form market efficiency or rational expectations, the expected value of the forecast error is zero which implies all conditional

expectations of forecast revisions should be zero and subsequently β should be zero. In the tables above $\hat{\beta}$ is significant and positive everywhere meaning the models experience an implied smoothing property where Bipower Variation especially stands out. For US30 forecast revisions such significance is however limited to the first few lags. S&P500 revisions on the other hand experience a far more significant cyclical pattern with a peak at the one day lag (81 lags: 81*5min=6:45 uur open outcry trade) which might imply there is still some cyclical behavior left in the S&P500 series not caught by the seasonal pattern. However autocorrelation coefficients for both series are generally low.

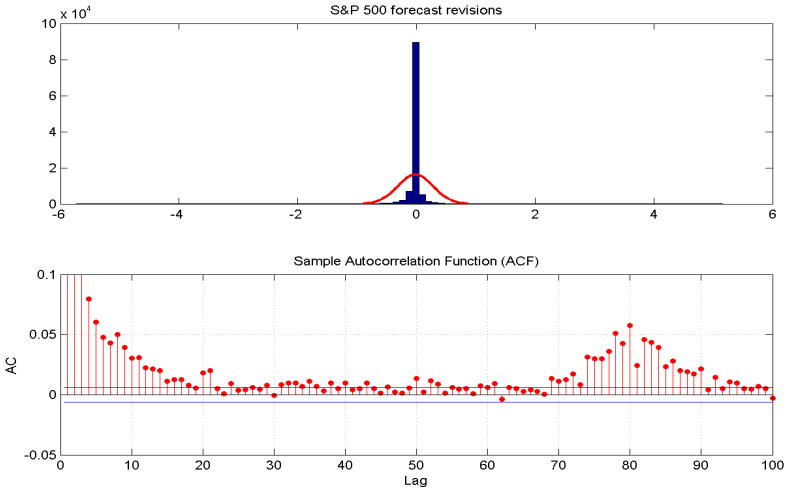

(a) S&P500 revisions histogram & autocorrelation function

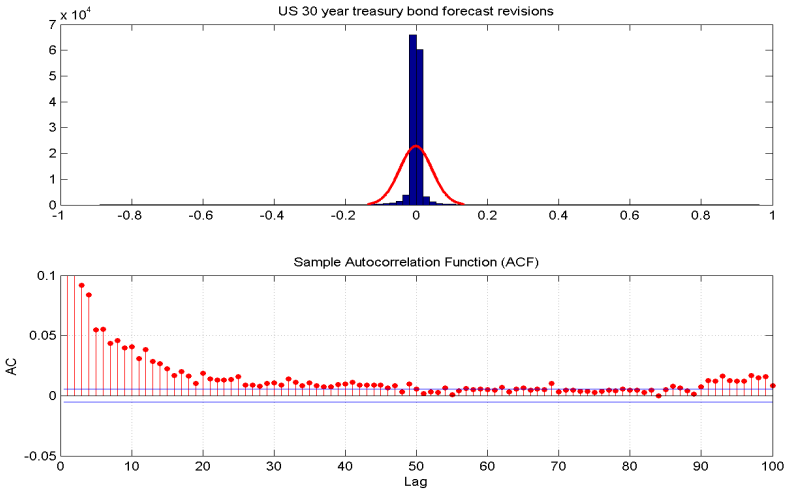

(b) US30 revisions histogram & autocorrelation function

Figure 18: S&P500 and US30 revisions distribution and autocorrelation functions for Seasonal Moving Average forecasts using RV measure.

5 Conclusion

This thesis has reviewed the forecasting ability of begin-of-the-day variance to daily variance in a broad sense. Multiple methods have been used to forecast daily variance and have been tested against two benchmarks. Furthermore these models have been tested on two time series off different origin and using multiple volatility measures as to gain a robust view upon predictability and model performance. Parameter choices therein have been diversified and where possible justified. Furthermore in order to cope with the defiant S&P500 data sample, including the financial crisis and highest volatility ever recorded, robust regressions and performance measures we used.

Summarizing we can conclude that that high frequency return data from begin-of-the-day contains valuable information about end-of-the-day volatility. In itself, it has higher correlation to daily variance then the RW after just 30 minutes of observations, equaling about 13% of daily variance. Compared to the GARCH(1,1) benchmark, 15-20 minutes of observations even suffice to outperform in terms of correlation (VR=8%). Scaling these early variance figures to daily proportions using a seasonal shape leads to further improvements in terms of R^2_{MAD} as well as HMSPE. The Fourier Flexible Form can herein generally be addressed best, especially as HMSPE is concerned, leading to squared correlations of about 0.65-0.70 after half an hour when only 13% of daily volatility of observed. Yet a simple scaling by Mincer-Zarnowitz coefficients outperforms all seasonal scaling methods in terms of R^2_{MAD} and during the first 30 minutes also on HMSPE. Leading to squared correlations of up to 0.75 after 30 minutes of returns for both securities. Additionally adding overnight information to the data set does, however, not improve estimation considerably. Reason can partly be found in the way daily variance is defined (over open outcry returns only) but might also be due to an other type of investors active through night- or daytime.

An interesting new finding during our quest for better predictions lies in Mincer-Zarnowitz regression coefficients. In general these regressions are used to obtain ex-post performance statistics. During this research we took an different approach and tried to model MZ coefficients α and β as unobserved states using Kalman Filtering. Subsequently we engaged in forecasting these values and using them to eliminate forecasting bias which resulted in major improvements over unadjusted forecasts in terms of MSPE and HMSPE. For US30 decreases up to a factor 6 were achieved where, favorably, the biggest gain is to be found near the start of open outcry trade. As was expected since bias is highest during that time. However this is just a first exploration on such adjustments, with no earlier documentation on this matter, further research should reveal

sustainability of such improvements over different samples, securities and settings. Yet if such is confirmed, bias adjustments should be applicable to a wide variety of forecasts. Second, the simpler more practical estimation procedure through direct MZ scaling of begin-of-the-day variance or MZ scaling of AVG/EWMA/FFF forecasts delivered similar and during early start of open outcry trade even better results.

In terms of efficiency it stands that roughly the first 15 minutes of observations for S&P500 and first 20 minutes of observations for US30 are most informative. After this period the increase in squared correlation is outweighed by the extra variance observed in making the new forecast. Interestingly the first few, 0-10 minutes, often lead to worse correlations than under exclusion of such evidently noisy measurements. To my knowledge now earlier research has been conducted on this subject and no rejection be it conformation on this matter has been found. However, the loss in correlation is generally small.

Two last remarks are to be made. First on FFF performance sensitivity where it is to be noted that the estimation sample should not be chosen too short. This could lead to unstable parameter estimation and consequently worse performance. An estimation sample of 252 days is therefore preferred over the alternatively tested 22 days. Second, mutual comparison forecasting performance over different volatility measures leads to the conclusion that they are overall quite comparable. However inherent properties of the volatility estimators do now and then inflict unexpected behavior. Relying on multiple measures is therefore preferable as to reduce the influence of chance.

Further research

Concluding, this research added a broad discussion to the literature. Helping to understand the dynamics of volatility, being possibly useful for risk managing and derivative pricing. Yet a variety of complementing research could be conducted to strengthen (or reject) conclusions found here. Other data samples should be investigated; other security types reviewed and other parametrizations tested. Additionally one could conduct research on the proper R^2 be it R^2_{MAD} correction due to testing against RV instead of IV. Such adjustments are omitted within this research yet would form a worthwhile addition. Another interesting road could be to test similar hypothesis in other, i.e. European, Asian or emerging markets or under different return observations, e.g. midquotes. Furthermore, simulation studies would create a controlled environment to help test the limits and sensitivities of methods employed to a variety of return distributions.

Further elaboration of the Kalman Filter and EM approach conducted in section 3.3.3 should

be considered to test robustness and sensitivities of results obtained here. Additionally more advanced models could be created for α and β; other parametrizations for $\{R, Q\}$ chosen and possibly combined by shrinkage procedures; Analytical expressions to the EM likelihood function under restrictions could perhaps be obtained or another algorithm allowing for non normal error distribution as a Particle Filter might be used.

Results make an attractive impression ratifying additional research with possible future implementation in risk management applications.

Acknowledgements

Special thanks go out to the supervisor (Prof. Dr. D.J.C. van Dijk) and co-reader (Prof. Dr. M.J. McAleer) whom placed worthwhile suggestions and comments that helped improve this thesis. Any remaining errors are my responsibility. All calculations made in this paper are based on software written by the author using the Matlab programming language.

References

Ahoniemi, K., Lanne, M., 2010. Realized volatility and overnight returns. Research Discussion Papers 19/2010, Bank of Finland.

Andersen, T., Bollerslev, T., 1994. Intraday seasonality and volatility persistence in foreign exchange and equity markets (186).

Andersen, T., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4 (2-3), 115–158.

Andersen, T., Bollerslev, T., 1998a. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. International Economic Review 39, 885–905.

Andersen, T., Bollerslev, T., 1998b. Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. Journal of Finance 53 (1), 219–265.

Andersen, T., Bollerslev, T., Diebold, F., 2002. Parametric and nonparametric volatility measurement. NBER Technical Working Papers 0279, National Bureau of Economic Research, Inc.

Andersen, T., Bollerslev, T., Diebold, F., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 1999a. The distribution of exchange rate volatility. NBER Working Papers 6961, National Bureau of Economic Research, Inc.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 2003. Modelling and forecasting realized volatility. Econometrica 71 (2), 579–625.

Andersen, T., Bollerslev, T., Lange, S., 1999b. Forecasting financial market volatility: Sample frequency vis-a-vis forecast horizon. Journal of Empirical Finance 6, 457–477.

Andersen, T., Bollerslev, T., Meddahi, N., 2005. Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities. Econometrica 73 (1), 279–296.

Areal, N., Taylor, S., 2002. The realized volatility of ftse-100 futures prices. Journal of Futures Markets 22 (7), 627–648.

Asai, M., McAleer, M., M.C., M., 2012. Modelling and forecasting noisy realized volatility. Computational Statistics and Data Analysis 56 (1), 217–239.

Baillie, R., Bollerslev, T., 1991. Intra-day and inter-market volatility in foreign exchange rates. Review of Economic Studies 58, 565–585.

Bandi, F., Russell, J., 2005. Comment on "an unbiased measure of realized variance" and "realized variance and market microstructure noise". Tech. rep., The University of Chicago.

Bandi, F., Russell, J., 2006. Market microstructure noise, integrated variance estimators, and the limitations of asymptotic approximations : A solution. Working papers, Graduate School of Business, University of Chicago.

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

Barndorff-Nielsen, O., Hansen, P., Lunde, A., Shephard, N., 2009. Realized kernels in practice: trades and quotes. Econometrics Journal 12 (3), C1–C32.

Barndorff-Nielsen, O., Hansen, P., Lunde, A., Shephard, N., 2011. Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. Journal of Econometrics 162, 149–169.

Barndorff-Nielsen, O., Shephard, N., 2002. Estimating quadratic variation using realized variance. Journal of Applied Econometrics 17 (5), 457–477.

Barndorff-Nielsen, O., Shephard, N., 2003. Power and bipower variation with stochastic volatility and jumps. Economics Papers 2003W17, Economics Group, Nuffield College, University of Oxford.

Battese, G., Griffiths, W., 1980. On r2-statistics for the general linear model with non-scalar covariance matrix. Australian Economic Papers 19 (35), 343–48.

Bedendo, M., Hodges, S., 2004. A parsimonious continuous time model for equity returns: Inferred from high frequency data. International Journal of Theoretical and Applied Finance 07 (08), 997–1030.

Black, F., 1976. Studies of stock price volatility changes. Proceedings of the Business and Economic Statistics Section, American Statistical Association, 177–181.

Blomquist, N., 1980. A note on the use of the coefficient of determination. The Scandinavian Journal of Economics 82 (3), 409–412.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31 (3), 307–327.

Bollerslev, T., 2008. Glossary to arch (2008-49).

Buse, A., 1973. Goodness of fit in generalized least squares estimation. The American Statistician 27 (3), 106–108.

Clements, M., 1997. Evaluating the rationality of fixed-event forecasts. Journal of Forecasting 16, 225–239.

Corsi, F., Mittnik, S., Pigorsch, C., Pigorsch, U., 2008. The volatility of realized volatility. Econometric Reviews 27, 46–78.

Dacorogna, M., MUller, U., Nagler, R., Olsen, R., Pictet, O., 1993. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. Journal of International Money and Finance 12, 413–438.

Danielsson, J., Zigrand, J., 2006. On time-scaling of risk and the square-root-of-time rule. Journal of Banking & Finance 30 (10), 2701–2713.

Dempster, A., Laird, N., Rubin, D., 1977. Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society, Series B 39 (1), 1–38.

Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. Econometrica 50 (4), 987–1007.

Falkenberry, T., 2001. High frequency data filtering. Tick Data.

Fama, E., 1965. The behavior of stock market prices. Journal of Business 38, 34–105.

Fitzgerald, R., 1971. Divergence of the kalman filter. IEEE Transactions on Automatic Control AC-16 (6), 736–747.

Franses, P., Chang, C.-L., McAleer, M., 2011. Analyzing fixed-event forecast revisions. Working Papers in Economics 11/25, University of Canterbury, Department of Economics and Finance.

Frijns, B., Margaritis, D., 2008. Forecasting daily volatility with intraday data. The European Journal of Finance 14:6, 523–540.

Gallant, A., 1984. The fourier flexible form. American Journal of Agricultural Economics 66 (2), 204–208.

Gatheral, J., Oomen, R., 2010. Zero-intelligence realized variance estimation. Finance and Stochastics 14 (2), 249–283.

Hansen, P., Lunde, A., 2004. An unbiased measure of realized variance. Working paper.

Hansen, P., Lunde, A., 2006. Realized variance and market microstructure noise (with discussion). Journal of Business and Economic Statistics 24, 127–218.

Harju, K., Hussain, S., 2006. Intraday seasonalities and macroeconomic news announcements. Working Papers 512, Hanken School of Economics.

Harju, K., Hussain, S., 2008. Intraday return and volatility spillovers across international equity markets. International Research Journal of Finance and Economics 22, 205–220.

Harris, L., 1986. A transaction data study of weekly and intradaily patterns in stock returns. Journal of Financial Economics 16 (1), 99–117.

Harvey, A., 1991. Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press.

Isengildina, O., Irwin, S., Good, D., 2006. Are revisions to usda crop production forecasts smoothed? American Journal of Agricultural Economics 88 (4), 1091–1104.

Jacod, J., Shiryaev, A., 1987. Limit Theorems for Stochastic Processes. Springer.

Jones, R., 1980. Maximum likelihood fitting of arma models to time series with missing observations. Technometrics 22 (3), 389–395.

Kalman, R., 1960. A new approach to linear filtering and prediction problems. Transactions of the ASME–Journal of Basic Engineering 82 (Series D), 35–45.

Ledoit, O., Wolf, M., 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. Journal of Empirical Finance 10 (5), 603–621.

Lee, J., Shevlyakov, G., Shin, V., 2006. Robust estimation of a correlation coefficinet for e-contaminated bivariate normal distributions. Autom. Remote Control 67 (12), 1940–1957.

Mandelbrot, B., 1963. The variation of certain speculative prices. Journal of Business 36, 394–419.

Martens, M., 2001. Forecasting daily exchange rate volatility using intraday returns. Journal of International Money and Finance 20 (1), 1–23.

Martens, M., Dijk, D. v., 2007. Measuring volatility with the realized range. Journal of Econometrics 138 (1), 181–207.

Maybeck, P., 1979. Stochastic models, estimation, and control Volume 1. Academic Press.

Mincer, J., Zarnowitz, V., 1969. The evaluation of economic forecasts. In: Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance. NBER Chapters. National Bureau of Economic Research, Inc, pp. 1–46.

Muller, U., Dacorogna, M., Olsen, R., Pictet, O., Schwarz, M., Morgenegg, C., 1990. Statistical study of foreign exchange rates, empirical evidence of a price change scaling law and intraday analysis. Journal of Banking of Finance 14, 1189–1208.

Munnix, M., Schafer, R., Guhr, T., 2010. Impact of the tick-size on financial returns and correlations. Physica A 389, 4828–4843.

Nelson, D., 1991. Conditional heteroskedasticity in asset pricing: a new approach. Econometrica 59, 347–370.

Newey, W., West, K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55 (3), 703–708.

Ng, H.S., L. K., 2006. How does sample size affect garch models? In: JCIS'06. pp. –1–1.

Nordhaus, W., 1987. Forecasting efficiency: Concepts and applications. The Review of Economics and Statistics 69 (4), 667–674.

Parkinsen, M., 1980. The extreme value method for estimating the variance of the rate of return. The Journal of Business 53 (1), 61–65.

Patton, A. J., Sheppard, K., 2008. Evaluating volatility and correlation forecasts. OFRC Working Papers Series 2008fe22, Oxford Financial Research Centre.

Sangsuk-lam, S., Bullock, T., 1990. Analysis of discrete-time kalman filtering under incorrect noise covariances. IEEE Transactions on Automatic Control 35 (12), 1304–1308.

Schafer, J., Strimmer, K., 2005. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. Statistical Applications in Genetics and Molecular Biology 4 (1), 32.

Schweppe, F., 1965. Evaluation of the likelihood ratios for gaussian signals. IEEE Transactions on Information Theory 11, 61–70.

Schwert, G., 1998. Stock market volatility: Ten years after the crash. NBER Working Papers 6381, National Bureau of Economic Research, Inc.

Shephard, N., Barndorff-Nielsen, O., Lunde, A., 2006. Subsampling realised kernels. Economics Series Working Papers 278, University of Oxford, Department of Economics.

Shevlyakov, G., Smirnov, P., 2011. Robust estimation of the correlation coefficient: An attempt of survey. Austrian Journal of Statistics 40 (1 and 2), 147–156.

Shumway, R., Stoffer, D., 1982. An approach to time series smoothing and forecasting using the em algorithm. J. Time Series Analysis 3 (4), 253–264.

Stoffer, D., Wall, K., 1991. Bootstrapping state space models: Gaussian maximum likelihood estimation and the kalman filter. Journal of the American Statistical Association 86, 1024–1033.

Taylor, S., 1982. Financial returns modelled by the product of two stochastic processes - a study of daily sugar prices 1961-79. Time Series Analysis: Theory and Practice 1, 203–226.

Taylor, S., Xu, X., 1997. The incremental volatility information in one million foreign exchange quotations. Journal of Empirical Finance 4 (4), 317–340.

Thomakos, D., Wang, T., 2003. Realized volatility in the futures markets. Journal of Empirical Finance 10, 321–353.

Thompson, G., 1988. Choice of flexible functional forms: Review and appraisal. Western Journal of Agricultural Economics 13 (2), 169–183.

Wood, R., McInish, T., Ord, J., 1985. An investigation of transactions data for nyse stocks. Journal of Finance 40 (3), 723–39.

Wu, L., 2010. Variance dynamics: Joint evidence from options and high-frequency returns. Journal of the American Statistical Association 100, 1394–1411.

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

Zhou, B., 1996. High frequency data and volatility in foreign exchange rates. Journal of Business & Economic Statistics 14 (1), 45–52.

List of Figures

List of Tables

Appendix A

Microstructure effects

The 10 second trade prices used for this research are exposed to sources of noise, also known as microstructure effects. As these effects play a pivotal role during this research, let us start by shamming some light on their origin and possible future accommodations to prevent such noise in the first place. The best way to gain insight might be through some illustrations. Figure A.1 for instance shows the artificially generated stickiness of prices inherent to the way we organize our data. Meaning the true timing of price changes is different than the timing used during this research. This feature results in negative bias in realized volatility estimation, Bandi and Russell (2005). This can easily be observed as spikes in the log prices can only be dwarfed, not amplified, resulting in lower observed volatility than was actually there.

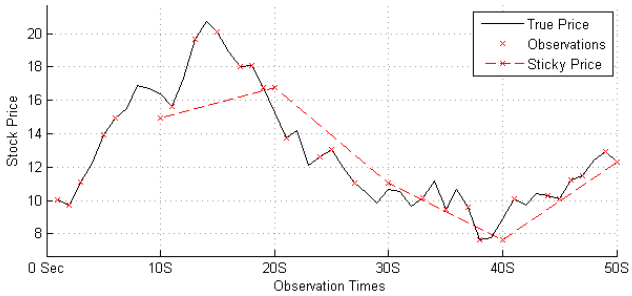

Figure A.1: Stale or Sticky prices cause negative first order autocorrelation in the return series and negative bias in realized volatility estimation.

The bid/ask spread is a second highly influential source of noise and should be circumvented if possible. Gatheral and Oomen (2010) e.g. found a 50% efficiency gain to be obtained through using mid-quote data for RV instead of normal price quotes. The noise induced bias in realized variance can intuitively be explained using figure A.2. Similar to reality, bid and ask prices are observed randomly, showing a more erratic observed price path and thus greater variance than was truly there. Consequently trade data enforces a positive bias on realized variance overestimating true integrated volatility of the underlying price process.

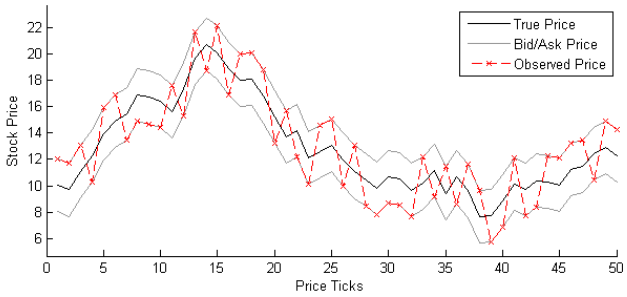

Figure A.2: High frequency mid-quote data is far less noisy than trade data at the same sampling frequency.

Beside the above mentioned there are numerous others effects all imposing their own difficulties on the volatility measures. Realized Range for instance becomes biased due to discrete price observations rather than the potentially higher and lower unobserved prices from a continues process. Realized Volatility on the other hand remains unbiased under this type of noise. Recapitulating, there are many different sources of microstructure noise. As so their true influence might be hard to predict and thus difficult to adjust for. One can however attempt to minimize noise by careful selection of data sources. In practice nonetheless the direction of sparse sampling is often chosen as to mitigate the influence of noise. This approach is further discussed in section 3.1.1.

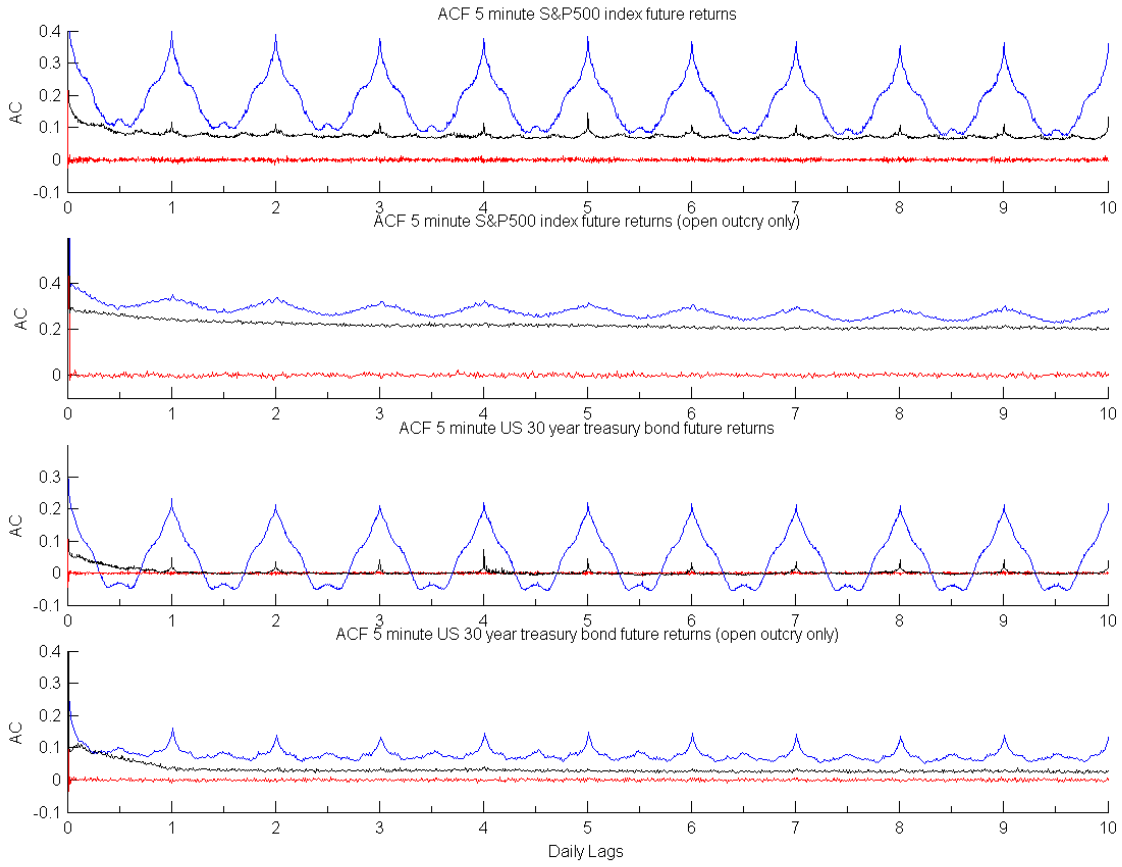

Figure A.3: Autocorrelation functions up to 10 dag lags for (absolute) S&P500 index returns and open outcry only (absolute) S&P500 index returns, as well as (absolute) US 30 year treasury bond futures returns and open outcry only (absolute) US 30 year treasury bond returns. Both autocorrelation functions show a strong cyclical behavior as the result of strong diurnal patterns. Graphs 2 and 4 only take open outcry returns into account resulting in higher overall correlation. Autocorrelations for return observations are depicted in red; absolute returns in blue and deseasonalized absolute returns in black.

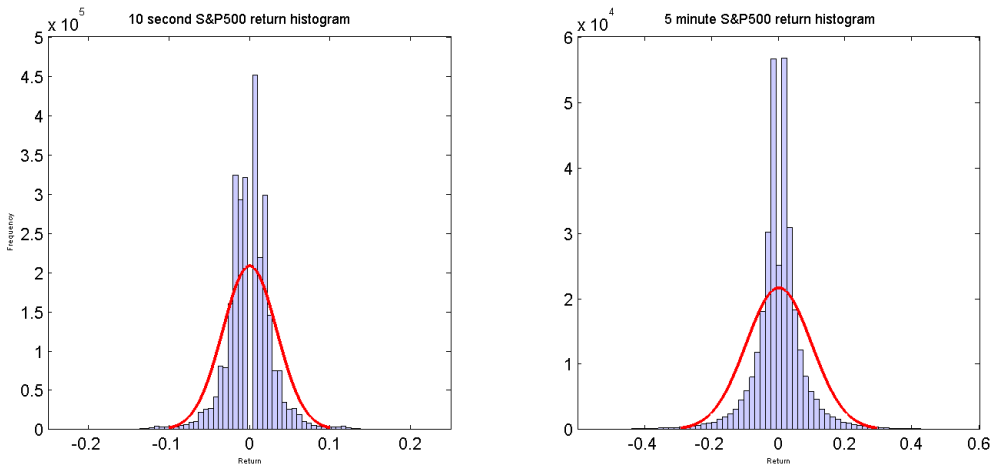

Figure A.4: Histogram of 10 second and 5 minute S&P500 index returns accompanied by a fitted normal distribution with the same mean and variance. As can be seen in the descriptive statistics on page 10 the distributions are quite symmetrical, have fat tails and greater mass in the middle. Note that due unavoidable memory issues within Matlab an abundance of zero returns has been left out of the histogram.

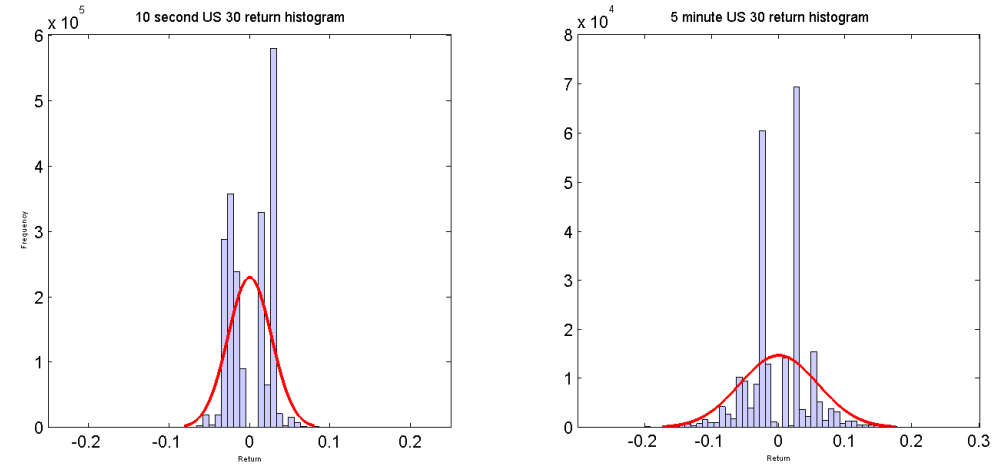

Figure A.5: Histogram of 10 second and 5 minute US 30 year treasury bond returns accompanied by a fitted normal distribution with the same mean and variance. The return distributions of US 30 year treasury bonds do not look as fluent as is the case for S&P500 index returns. This is most probably due to the greater tick-size the US Government uses for US treasury bonds, i.e. 1/32nds of a dollar apposed to 1/100th of a dollar for regular securities on the New York Stock Exchange (NYSE). Furthermore the US bonds only consist of 1 security as apposed to the weighted average of 500 securities for the S&P500 index wearing the effect even further down. For a proper evaluation of the impact of tick-size see: Munnix et al. (2010). Note that due unavoidable memory issues within Matlab an abundance of zero returns has been left out of the histogram.

Positive				Negative		
Day	*Date*	*Return [%]*		*Day*	*Date*	*Return [%]*
Tu	2008/10/28	7,95		Th	2008/10/09	-9,01
Mo	2008/10/13	7,69		We	2008/10/15	-8,34
Th	2008/11/13	5,86		Mo	2008/12/01	-6,64
We	2008/11/26	5,43		Tu	2008/10/07	-6,48
We	2008/01/23	5,37		Mo	2008/09/29	-6,44
We	2008/12/03	4,92		Th	2008/11/20	-5,56
Mo	2009/03/23	4,74		We	2008/11/19	-5,03
Fr	2008/12/05	4,31		Tu	2008/10/14	-4,69
Th	2009/03/12	4,24		Mo	2009/02/23	-4,49
Mo	2008/10/20	4,21		Tu	2009/01/20	-4,05

Table A.1: Ten most extreme positive and negative end-of-the-day return observations for S&P500 index returns as measured from 09:30 until 16:15 over 10 second intervals. If there is a highly likely cause to the event these are stated in the last column. However it seems notoriously difficult to designate a unambiguous cause in terms of unexpected events or news.

Positive				Negative		
	Date	*Return [%]*			*Date*	*Return [%]*
We	2009/03/18	4,47		We	2008/01/23	-2,68
Th	2008/11/20	3,55		Tu	2003/07/15	-2,53
Tu	2008/11/04	2,25		Fr	2004/04/02	-2,28
Tu	2008/12/16	2,23		Fr	2009/01/02	-2,27
Mo	2008/09/29	2,13		Th	2009/05/21	-2,14
Fr	2004/01/09	1,95		Th	2003/07/31	-2,13
Tu	2004/06/15	1,81		Tu	2008/09/30	-2,10
We	2003/08/06	1,77		Mo	2009/06/01	-1,98
Mo	2007/11/26	1,75		We	2009/05/27	-1,97
Mo	2008/09/15	1,74		Fr	2008/10/24	-1,95

Table A.2: Ten most extreme positive and negative end-of-the-day return observations for US30 year treasury bond returns as measured from 08:20 until 17:00. If a likely cause can be determined this is stated in the last column. However it seems notoriously difficult to designate a unambiguous cause in terms of unexpected events, or news.

Appendix B

Realized Variance

Realized variance (RV) is the first and most basic measure of integrated volatility. According to the theory RV can recover the volatility defined by the quadratic variation of a semi martingale price process[24]. Using high frequency returns it can be estimated by $\sum_1^n (Y_{t,i\Delta} - Y_{t,(i-1)\Delta})^2$, resulting in the formula for realized variance.

$$RV_t = \sum_{i=1}^{n} r_{t,i}^2 \tag{B.1}$$

where $r_{t,i}$ is the return $\left(Y_{t,i\Delta} - Y_{t,(i-1)\Delta}\right)$ over period $[i-1, i]$. In an ideal world with no market frictions, i.e. $\varepsilon_{t,i} = 0$ for every t and i, this would be an unbiased, consistent and highly efficient estimator for integrated variance[25]. One would merely have to drive up the sampling frequency to obtain a more accurate estimate and as $n \to \infty$, realized variance converges to IV. Unfortunately reality is not this structured. Practical implementation has to confront the fact that prices are not recorded continuously and markets are not frictionless. Leading RV to be biased and inconsistent resulting in great errors if one was to use RV on high sampling frequencies without correction. To mitigate the bias problems lots of correction methods have been proposed, see for example Andersen and Bollerslev (1998a); Andersen et al. (1999b, 2005); Asai et al. (2012) among many others. Additional problems appear when the log price process also exhibits jumps, as can often be observed in intra-day data. Stated in terms of integrals and using our research specific time notation t, i instead of a general t, the diffusion process of (2) can be written as

$$X_{t,i} = \int_0^1 \mu_{t,i} di + \int_0^1 \sigma_{t,i} dB_{t,i} + J_{t,i} \tag{B.2}$$

Where $J_{t,i} = \sum_{j=1}^{N_i} C_j$ is a finite activity jump process, meaning it has a finite number of jumps, N_i, in any bounded time interval and C_j denotes the size of the jump in question. Calculation of RV over $X_{t,i}$ in (B.2) now estimates not only IV but additionally the quadratic jump components. Quadratic variation thus calculates:

$$QV_t = \int_0^1 \sigma_{t,i}^2 di + \sum_{j=1}^{N_i} C_j^2 \tag{B.3}$$

[24] See Jacod and Shiryaev (1987) and Barndorff-Nielsen and Shephard (2003) for further exposition on quadratic variation (QV) assuming semimartingale properties for log price process.

[25] RV converges at rate $\sqrt{n}$, see Barndorff-Nielsen and Shephard (2002)

Delivering the important finding that realized variance is an inconsistent measure of IV in the presence of jumps. This second component can however be disentangled from integrated variance by consistent estimation of IV as done by Realized Bipower Variation (discussed in the next section).

Having its flaws, the ease of computation makes RV a much used IV estimation technique nonetheless. To deal with the bias in practice, another solution is favorite. Sparse sampling is used as to mitigate microstructure effects. Though this tactic merely limits the influence of noise rather than corrects for it, Hansen and Lunde (2006) advocate it works quite well. Empirical evidence has shown that the optimal trade-off between noise induced bias and variance of the volatility estimates lies somewhere from 5 up to 30 minutes sampling intervals[26]. Zhang et al. (2005) formally prove this by determining the optimal sampling frequency trade-off by means of minimizing the root mean squared error, $RMSE = \left(bias^2 + var\right)^{\frac{1}{2}}$.

Realized Bipower Variation

Due to its quadratic form realized volatility is quite sensitive to outliers and jumps in the return process. A more robust estimator can be found in Realized Bipower Variation (BPV). It is closely related to realized variance in the sense that BPV can be written as a generalization of RV, i.e. taking $r = 2$, $s = 0$ in equation (B.5) boils down to RV. But it has one major advantage over realized variance. By taking the product of subsequent return observations it mitigates the influence of outliers and is far less affected by jumps in the (log) price process. Again this can be shown by extending equation (2) with a finite activity jump process, $J_{t,i}$. i.e.:

$$X_{t,i} = \int_0^1 \mu_{t,i} di + \int_0^1 \sigma_{t,i} dB_{t,i} + J_{t,i} \tag{B.4}$$

Defining a dummy, $d_{t,i}$, to track these jumps

$$d_{t,i} = \begin{cases} 1, & J_{t,i} = J_{t,i+1} = 0 \\ 1, & J_{t,i} \neq 0,\ J_{t,i+1} \neq 0 \\ 0, & \text{elsewhere} \end{cases}$$

The Realized Bipower Variation, normally defined as

[26] Schwert (1998) uses 15-minute returns for construction of daily stock market volatilities, Taylor and Xu (1997) and Andersen et al. (1999a) exploit 5-minute returns for daily exchange rate volatility measurement. Areal and Taylor (2002) also use 5 minute returns for FTSE-100 index futures contracts and Bedendo and Hodges (2004) use 5 minute returns on S&P500 future return series as well. Furthermore Zhang et al. (2005) shows approximately 7.5 minutes to be the optimal weigh-off in terms of RMSE.

$$\left\{\left(\frac{1}{n}\right)^{1-(r+s)/2}\right\}\sum_{i=1}^{n-1}|r_{t,i}|^r|r_{t,i+1}|^s, \qquad r,s\geq 0 \tag{B.5}$$

with n the number of daily increments, $r_{t,i}$ the return over $[i-1,i]$, can now be written as

$$\left\{\left(\frac{1}{n}\right)^{1-(r+s)/2}\right\}\sum_{i=1}^{n-1}|r_{t,i}|^r|r_{t,i+1}|^s d_{t,i} + \left\{\left(\frac{1}{n}\right)^{1-(r+s)/2}\right\}\sum_{i=1}^{n-1}|r_{t,i}|^r|r_{t,i+1}|^s\left(1-d_{t,i}\right) \tag{B.6}$$

Barndorff-Nielsen and Shephard show that the influence of the jump process in the first term converges to zero as $\frac{1}{n}\to 0$ and when $s+r<2$, due to the finite number of jumps. Additionally they prove that these conditions may be relaxed to $max(r,s)<2$ as long as the probability of observing two contiguous jumps converges to zero at rate $\frac{1}{n}$, which is the case taking a finite activity jump process. Next they prove that the second term also converges to 0 as long as $s<2$ and $r<2$, completing the result that realized Bipower Variation is unaffected by the presence of jumps and thus converges to IV.

Taking $r=s=1$, (B.5) simplifies to the realized Bipower Variation as used in this research.

$$BPV_t = \sum_{i=1}^{n-1}|r_{t,i}||r_{t,i+1}|, \qquad r=s=1 \tag{B.7}$$

Note that, as is the case for RV, a sampling frequency has to be chosen. The same argument holds as above, sample to scarce and lose efficiency, sample to frequent and the estimation bias due to microstructure effects could take over.

Realized Range

Another intuitive measure of volatility estimation is the Realized Range (RR) which makes use of the difference in maximum and minimum observed prices during a certain period of time. Properly scaled, Parkinsen (1980) shows that daily high-low range is up to 5 times more efficient than realized variance using daily squared returns. Correspondingly Andersen and Bollerslev (1998a) show that daily high-low range performs similar to realized variance sampled at 4 to 8 times higher frequency. If this result holds for every sampling rate, in theory we'd have an ever more efficient estimator than RV. In this mindset Martens and Dijk (2007) proposed to use high-low range with intra-day data and dubbed the resulting estimator Realized Range. Simulation studies on their part, excluding market frictions, support the finding that RR is unbiased and more efficient in terms of RMSE for every sampling rate from 1 minute up to a day. Adding

noise induced by infrequent trading and bid-ask bounce, however, has the effect that RR becomes biased. Suggesting there is again an optimal trade-off, by simulation found to lay somewhere at $\Delta = 45$ minutes. The bias adjusted realized range on the other hand does not suffer from microstructure effects and thus benefits from greater sampling frequency.

To formalize the setting we again take $P_{t,i\Delta}$ to be the last observed price in the i-th interval of length $\triangle$, as we have already assumed in equation (1). $H_{t,i} = sup_{(i-1)\Delta<j<i\Delta} P_{t,j}$ is the highest observed price or supreme and $L_{t,i} = inf_{(i-1)\Delta<j<i\Delta} P_{t,j}$ is the lowest observed price or infimum. The scaled high-low range estimator is than defined as

$$\frac{(lnH_{t,i} - lnL_{t,i})^2}{4ln2} \tag{B.8}$$

Aggregating over the n daily intervals gives the realized range:

$$RR_t = \frac{1}{4ln2} \sum_{i=1}^{n} (lnH_{t,i} - lnL_{t,i})^2 \tag{B.9}$$

This version of the RR estimator is however somewhat sensitive to macroeconomic noise constituents. As prices are observed about continuously the high price will likely be a ask price whereas the low price will likely be a bid, exactly overestimating the range by the bid/ask spread. On the other hand infrequent trading induces a negative bias due to unobserved intermediary prices. Among other proposed bias-adjustments, Martens and Dijk suggested to multiply by the ratio of daily range, RR_{t-l}^{Daily}, over the previous q days to the realized range over the same previous q days. The idea is that noise is of little concern dealing with daily range. The realized range over n intra-day intervals on the other hand is surely affected. The ratio between the two could thus be used to scale the original estimates. Resulting in a bias-adjusted realized range estimator:

$$RR_t^{Adj} = \left(\frac{\sum_{l=1}^{q} RR_{t-l}^{Daily}}{\sum_{l=1}^{q} RR_{t-l}} \right) RR_t \tag{B.10}$$

What remains is the choice of q. Here we naturally follow Martens and Dijk and use $q = 66$ days or approximately 3 months. Note to this bias adjustment is that it could make disproportional adjustments when dealing with highly volatile data in the sense that sudden high returns average out in a daily accumulation to end-of-day Realized Range but do not with the Daily Range. Consequently the adjustment term in B.10 can get rather large. Our S&P500 index futures sample is typically such a challenging one.

Two Time Scales Estimation

First two volatility estimators have one thing in common. Taking the empirically much used 5 minute intervals, they throw away vast amounts of data just to reduce the influence of microstructure effects. In our case this sparse sampling encompasses throwing away approximately 97% of the data. Pure silliness from a statistical point of view. To quote Zhang et al.: "It is difficult to accept that throwing away data, especially in such quantities, can be an optimal solution". Ultimately we'd like to create an unbiased and consistent estimator which uses all data, assuming this leads to more accurate estimation off course. Zhang et al. (2005) therefore develop a procedure, based on earlier work of Zhou (1996) on scaled Brownian motion plus noise, that uses both a higher and a lower data aggregation: the so called Two Time Scale (TTS) variance estimator. The argument is that low frequency RV estimation reduces bias to IV estimates rather than corrects for it. The relative size of this bias is proportional to the sampling frequency and can be corrected for by consistent estimation of the noise variance in the return process. Now recall that when n goes to infinity equation (6) estimates just that. Thus by combining these we should be able to get an unbiased and consistent estimator of IV[27]. To further improve efficiency Zhang et al. take the average RV over different subsamples as opposed to plain RV with sparse sampling. That is, when one observes N daily prices and samples with intervals of length $\Delta = 1/n$, there are $K = N/n$ different ways to lay these samples on the interval. In our case for example we have 10-second prices ($N = 8640$) and a sparse sampling frequency of 5-minutes ($n = 288$). There are thus 30 different possible subsample grids to base 5 minute RV estimates on. Starting the first returns from 9:30.00-9:35.00, one could also use 9:30.10-9:35.10 or 9:30.20-9:35.20 etc. to compute RV. Averaging these different sparsely sampled RV measures results in the average RV used in TTS estimation. The practical problem what to do with the begin and end of a trading session can be solved by proportionally inflating the daily volatility estimator with a factor $\frac{N}{N-\bar{n}}$. This factor originates from the fact that not all N but rather $N - \bar{n}$ observations are utilized by the K subsumples. In our case this is of lesser importance since we trade around the clock. Days can thus be molted together leaving us with only the begin and end at July 2003 and December 2009. These spots are manually altered by division through half the amount of subsamples.

In mathematical term we pursue the following:

[27] Note that this approach implicitly assumes the error terms to be independent. Empirically a highly questionable property as debated by Hansen and Lunde (2004, 2006) whom find persistence up until a 2 minute lag.

$$TTS_t = \frac{N}{N-\bar{n}}\left(\frac{1}{K}\sum_{k=0}^{K-1} RV_t^{(k)} - \frac{\bar{n}}{N}RV_t^{all}\right) \tag{B.11}$$

with,

$$RV_t^{(k)} = \sum r_{t,i}^2 \qquad for\ i = 1+\frac{k}{K}, 2+\frac{k}{K}, ..., n-1+\frac{k}{K}$$

$$RV_t^{all} = \sum_{j=1}^{N} r_{t,j}^2 \qquad for\ j = 1, 2, ..., N$$

$$\bar{n} = \frac{N-K+1}{K}$$

That is, the two time scales estimator takes the average RV for day t over multiple subsamples and corrects the bias with $\frac{\bar{n}}{N}$ times the realized variance for day t calculated using all available data.

Realized Kernel Estimation

The last discussed and widely accepted estimator of integrated variance we test here is the Realized Kernel (RK). These type estimators, as introduced by Barndorff-Nielsen et al. (2008), are based on the ideas of Hansen and Lunde (2004). They noticed that microstructure noise causes the high frequency intra-day returns to be autocorrelated resulting in biased RV estimation. They figured that the empirical autocorrelation function up to a certain lag H can thus be used to correct the bias in RV through a correction that works in the same way as in which robust covariance estimators of Newey and West (1987) achieve their consistency. The realized kernels are build upon the same principle. The kernel function, $K_t(Y_{t,\Delta})$, consists of RV,$\gamma_0(Y_{t,\Delta})$, which gets corrected by the empirical autocorrelations, $\sum_{h=1}^{H} k\left(\frac{h-1}{H}\right)\{\gamma_h(Y_{t,\Delta}) + \gamma_{-h}(Y_{t,\Delta})\}$, to adjust for market frictions. Formally, the kernel function of Barndorff-Nielsen et al. can thus be written as:

$$K_t(Y_{t,\Delta}) = \gamma_0(Y_{t,\Delta}) + \sum_{h=1}^{H_t} k\left(\frac{h-1}{H_t}\right)\{\gamma_h(Y_{t,\Delta}) + \gamma_{-h}(Y_{t,\Delta})\} \tag{B.12}$$

$$\gamma_h(Y_{t,\Delta}) = \sum_{j=1}^{K}\left(Y_{t,\Delta,j} - Y_{t,\Delta,j-1}\right)\left(Y_{t,\Delta,j-h} - Y_{t,\Delta,j-h-1}\right) = \sum_{j=1}^{K}\left(R_{t,\Delta,j}R_{t,\Delta,j-h}\right) \tag{B.13}$$

Where H_t is the bandwidth parameter which is to be estimated and can be seen as the estimated number of autocorrelations needed to bias adjust the Kernel estimate; $k(x)$ is the chosen kernel weighing function; $\gamma_0(Y_{t,\Delta})$ the realized variance; $\gamma_h(Y_{t,\Delta})$ the h-th order autocovariance of the observed log return series; K equals the size of sparse sampling intervals: N/n; and $Y_{t,\Delta,j}$ is the observed log price level at time j in increment Δ during day t. Note that since we estimate the realized kernel on contaminated Y_t series we get

$$K_t(Y_{t,\Delta}) = K_t(X_{t,\Delta}) + K_t(X_{t,\Delta}, \varepsilon_{t,\Delta}) + K_t(\varepsilon_{t,\Delta}, X_{t,\Delta}) + K_t(\varepsilon_{t,\Delta}) \tag{B.14}$$

Barndorff-Nielsen et al. (2011) show that as $K \to \infty$, $K(\varepsilon_{t,\Delta}) \xrightarrow{p} 0$, $K_t(X_{t,\Delta}) \xrightarrow{p} IV$ and that the dependence between ε_t and X_t is asymptotically unimportant. Making $K_t(Y_{t,\Delta})$ a consistent estimator of IV.

First step for estimation is choosing the weighing function. Here we take the Parzen kernel given by

$$k(x) = \begin{cases} 1 - 6x^2 + 6x^3 & 0 \leq x \leq 0.5 \\ 2\left(1 - x\right)^3 & 0.5 \leq x \leq 1 \\ 0 & x > 1 \end{cases} \tag{B.15}$$

This kernel has some desirable properties: it satisfies the smoothness conditions $k'(0) = k'(1) = 0$ and is guaranteed to produce a non-negative estimate. Furthermore we take a flat-top kernel, i.e. $\frac{h-1}{H}$ - imposing unit weight on the first autocovariance - as these are to produce unbiased and faster converging estimates[28]. For the optimal estimation of bandwidth H we follow the practical estimation procedures laid out by Barndorff-Nielsen et al. (2009). Their preferred bandwidth equals

$$H_t^* = c^* \cdot \xi_t^{4/5} \cdot K^{3/5}, \qquad \text{with } c^* = \left\{ \frac{k"(0)^2}{\int_0^1 k(x)^2 dx} \right\}^{1/5} \text{ and } \xi_t^2 = \frac{\omega_t^2}{\sqrt{n \int_0^n \sigma_t^4 dt}} \tag{B.16}$$

Where an accent denotes the derivative of the function and double accents denote the second derivative, not a transpose. c^* is only dependent of the chosen weighing function and can thus

[28]Smooth Flat-top kernels converge at the optimal rate of $K^{1/4}$, they can however not be guaranteed to produce non-negative volatility estimates. The non-negative kernels of Barndorff-Nielsen et al. (2011) do not have these problems although they have to sacrifice some efficiency. Consequently the convergence rate falls to $K^{1/5}$.

already be calculated. $k"(x)^2 = -12 + 36x$ for $0 \leq x \leq 0.5$ and $k"(x)^2 = 12(1-x)^3$ for $0.5 \leq x \leq 1$. The denominator $\int_0^1 k(x)^2 dx$ equals 0.269 yielding $c^* = \left(\frac{12^2}{0.269}\right)^{1/5} = 3.5134$. Obtaining an estimate for ξ_t is notoriously more difficult. We start by estimating ω_t^2 which is the variance of the error or noise term, $Var(\varepsilon_t)$, previously used in two time scales estimation. Here we take the same approach as during TTS estimation, using all available data to consistently estimate noise variance, with the addition of a little twist. For $\hat{\omega}_t^2$ to be a sensible estimator of ω_t^2, the error terms must be independent of each other through time. That is $E[\varepsilon_{t,i}\varepsilon_{t,i-q}]$ should be 0 for every q. Since there is overwhelming evidence that such is not the case for $q = 1$, see Barndorff-Nielsen et al. (2009), a modification has to be made. Following Barndorff-Nielsen et al. we take

$$\hat{\omega}_{t,z}^2 = \frac{RV_{t,z}^{all}}{2n_{t,z}}, \qquad z = 1, ..., q \tag{B.17}$$

That is, we calculate the realized variance using every q-th trade of day t and divide by twice the number returns n_t[29]. Subsequently varying the starting point of q and taking the average of the ω_t^2 estimates gives the final estimate: $\hat{\omega}_t^2 = \frac{1}{q}\sum_{z=1}^{q} \hat{\omega}_{t,z}^2$. Following Hansen and Lunde (2006); Barndorff-Nielsen et al. (2009) we take $q = 12$ so that subsequent returns in $RV_{t,z}^{all}$ are always 2 minutes apart, better complying the independence assumption[30].

For the estimation of so called integrated quarticity (IQ): $n\int_0^n \sigma_u^4 du$, we pose that this is roughly the same as IV_t^2 when σ_t^2 does not vary too much over the interval $[0, n]$[31]. The square root of IQ can thus be estimated by a preliminary version of $I\hat{V}_t$. Therefore realized variances are used based on sparse sampling with 20 min intervals as to reduce the influence of noise and use subsampling - in the same manner as done for the TTS estimator by shifting the starting points - to further improve efficiency[32]. $\hat{\xi}_t^2$ now becomes

[29]Assuming the log prices to follow an Itô process as stated in equation (2), the conditional mean of $Var(Y_t)$ conditional on the process of X_t equals $Var(X_t) + 2n\omega_t^2$ as stated in equation (6). Taking $\Delta \to 0$, $Var(X_{t,i\Delta/(i-1)\Delta}) \to 0$ stating that the conditional mean of $Var(Y_t)$, which can be estimated by RV_t^{all}, is a consistent estimator of the error variance: $2n\omega_t^2$, giving $\hat{\omega}_t^2 = \frac{RV_t^{all}}{2n_t}$.

[30]Hansen and Lunde (2006) argue that $\hat{\omega}_{t,z}^2$ should be large relative to $Var(X_t)/2n_t$ in estimating ω_t^2. As this is not always the case, especially for highly liquid assets where noise constituents are limited, this leads to an upward bias in $\hat{\omega}_t^2$. This in turn leads to a bigger, more conservative, choice of H. This is not a bad thing as in theory to large a bandwidth does less damage than to small a one. Furthermore larger values of H actually increase robustness to serial dependence in $\varepsilon_{t,i}$

[31]Note that this is a questionable yet necessary assumption as for example the bid-ask spread is typically found to by greater during the opening and closing of open outcry markets, creating a diurnal spread curve. Relaxation of this assumption for the estimation of integrated quarticity is however beyond of the scope of this thesis.

[32]Shephard et al. (2006) elaborate on the possible gains from subsampling for the realized kernel as a whole. They find that it does not improve estimation unless an inappropriate weighing function was chosen. In other cases it could even hurt.

$$\hat{\xi}_t^2 = \frac{\hat{\omega}_t^2}{RV_t^{subsample,sparse(20M)}} \tag{B.18}$$

resulting in the last ingredient to estimate the bandwidth parameter H^*in an optimal way[33].

$$\hat{H}_t^* = c^* \cdot \hat{\xi}_t^{4/5} \cdot K^{3/5} \tag{B.19}$$

Last remarks are placed upon end-effects and a possible solution, jittering. For consistency our estimator we need $K_t(\varepsilon_{t,\Delta}) \xrightarrow{p} 0$ as $n \to \infty$ in equation (B.14). As the first and last observation of a sample are often heavily contaminated with microstructure noise, this property is typically not satisfied making the kernel estimator inconsistent. The problem can be overcome by averaging the first and last observation over its m direct neighbors, also known as jittering. However, theoretically important, in practice optimal choice for m is often 1. Making these effects of no practical importance.

A graph containing all above estimated IV measures has been included. Measures are more or less in line with one another yet react differently to extreme return observations.

[33]This is the optimal solution in an asymptotic Mean Squared Error sense, see Barndorff-Nielsen et al. (2009). Alternative methods that seek optimal finite sample behavior have been proposed by Bandi and Russell (2006).

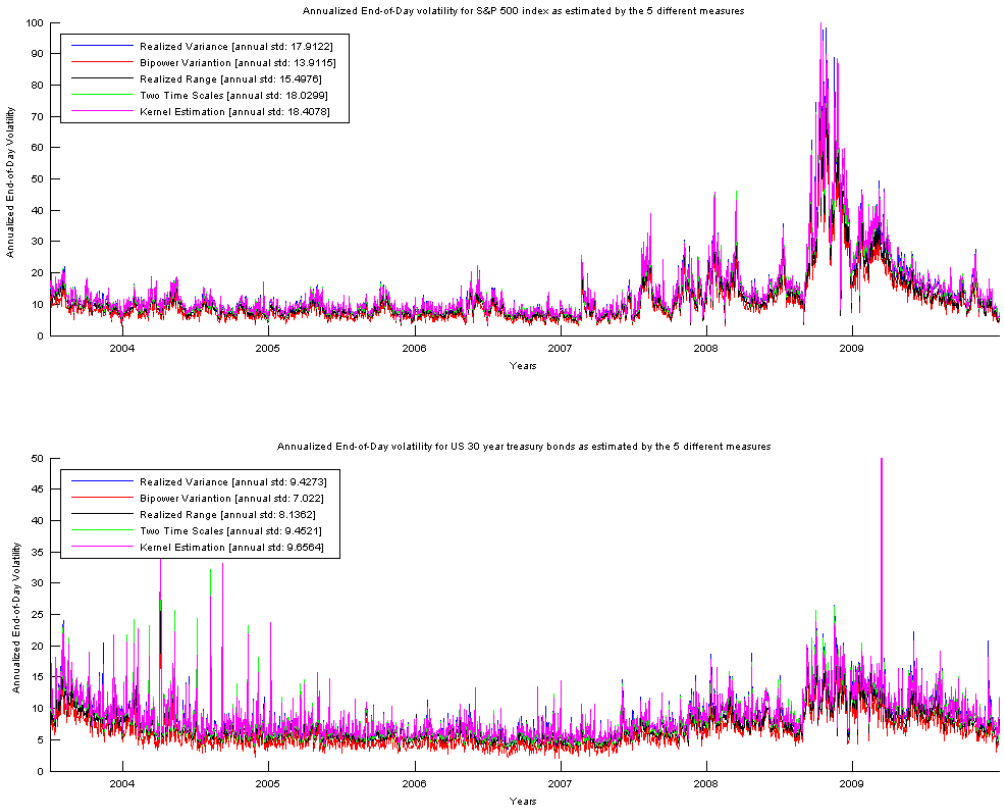

Figure B.1: Integrated Variance estimators with the estimated average annual standard deviation included between brackets. On average the Bipower Variation estimator seems to produces lowest annual volatility in both cases. Further rankings vary over the two securities but overall all figures seem to be in quite a close range. Most notable is the high peaks and annual standard deviation of the S&P500 as measured by the Realized Range. As pointed out earlier the bias does some seemingly disproportionate adjustment through the high volatility period.

Seasonal simulation experiment

To confirm the suspicion that the Fourier Flexible Form OLS regression should be altered as compared to Andersen and Bollerslev (1997) a little simulation experiment is set up where the return components from the return decomposition, $r_i = E[r_i] + \sigma_i s_i Z_i$, are already known. To simulate these return paths we take: $E[r_i] = 0$, $\sigma_i = [0.5,\ 0.1,\ 0.01]$, $Z_i \sim N(0,1)$ and s_i follows the parametrization:

$$ln\left(s_i^2\right) = f(\theta; i) = \mu_0 + \mu_1 \frac{i}{N_1} + \mu_2 \frac{i^2}{N_2} + \gamma \cdot cos\left(\frac{i2\pi}{N}\right) + \delta \cdot sin\left(\frac{i2\pi}{N}\right) \tag{B.20}$$

with $\theta = \{\mu_0 = -5,\ \mu_1 = 7,\ \mu_2 = -3,\ \gamma = 2,\ \delta = -1\}$ and where $N = 100$ and N_1 and N_2 are normalizing constants as noted in section 3.2.6. Seasonal patterns are now estimated via regression $x_i = f(\theta; i) + \varepsilon_i$ with the two different forms for x_i

1. Bollerslev $x_i \equiv 2 \cdot ln\left(|r_i - E[r_i]|\right) - ln\left(\sigma^2\right) + ln(N) = ln\left(s_i^2\right) + ln\left(Z_i^2\right) - E[ln(Z_i^2)]$ (B.21)

2. Adjusted $x_i \equiv 2 \cdot ln\left(|r_i - E[r_i]|\right) - ln\left(\sigma^2\right) + ln(N) - E[ln(Z_i^2)] = ln\left(s_i^2\right) + ln\left(Z_i^2\right) - E[ln(Z_i^2)]$ (B.22)

with $\varepsilon_i = ln\left(Z_i^2\right) - E[ln(Z_i^2)]$. Figure B.2 shows some graphical illustrations from the simulation experiment. It can clearly be seen that the unaltered seasonal pattern consistently underestimates the true seasonal pattern. The altered version on the other hand shows an unbiased course through the simulated daily seasonal as set out beforehand.

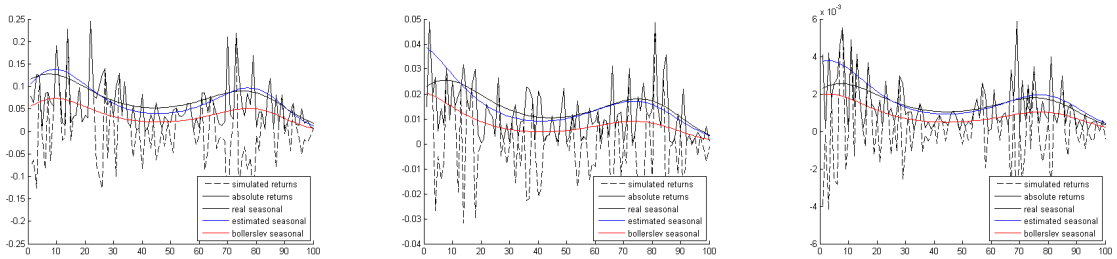

(a) $\sigma = 0.5$ **(b)** $\sigma = 0.1$ **(c)** $\sigma = 0.01$

Figure B.2: Seasonal simulation experiment: the unadjusted seasonal is clearly biased, yielding to low a seasonal pattern regardless of the choice for σ. The second approach with x_i as in (B.22) gives an unbiased estimate.

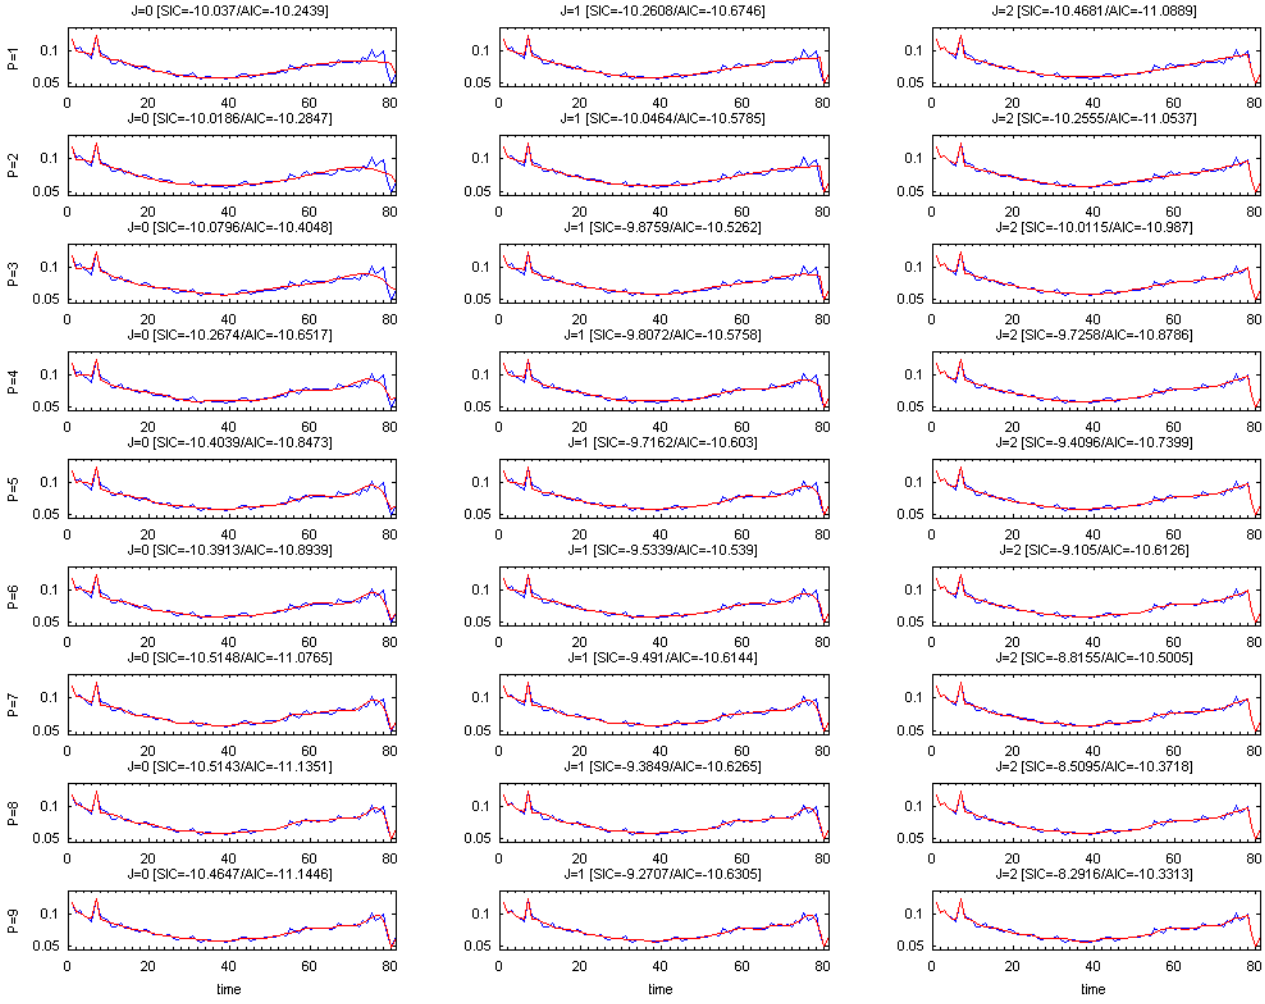

Figure B.3: Fitted Flexible Fourier Form on mean absolute S&P 500 index return data for different values of J and P accompanied by SIC and AIC values. The fit gradually improves as the amount of parameters increase. Note however that it is not the objective to track the ex-post volatility to our best capabilities but to make sensible forecasts. Tracking with to many parameters would induce noise in the forecast and enlarge parameter estimation uncertainty.

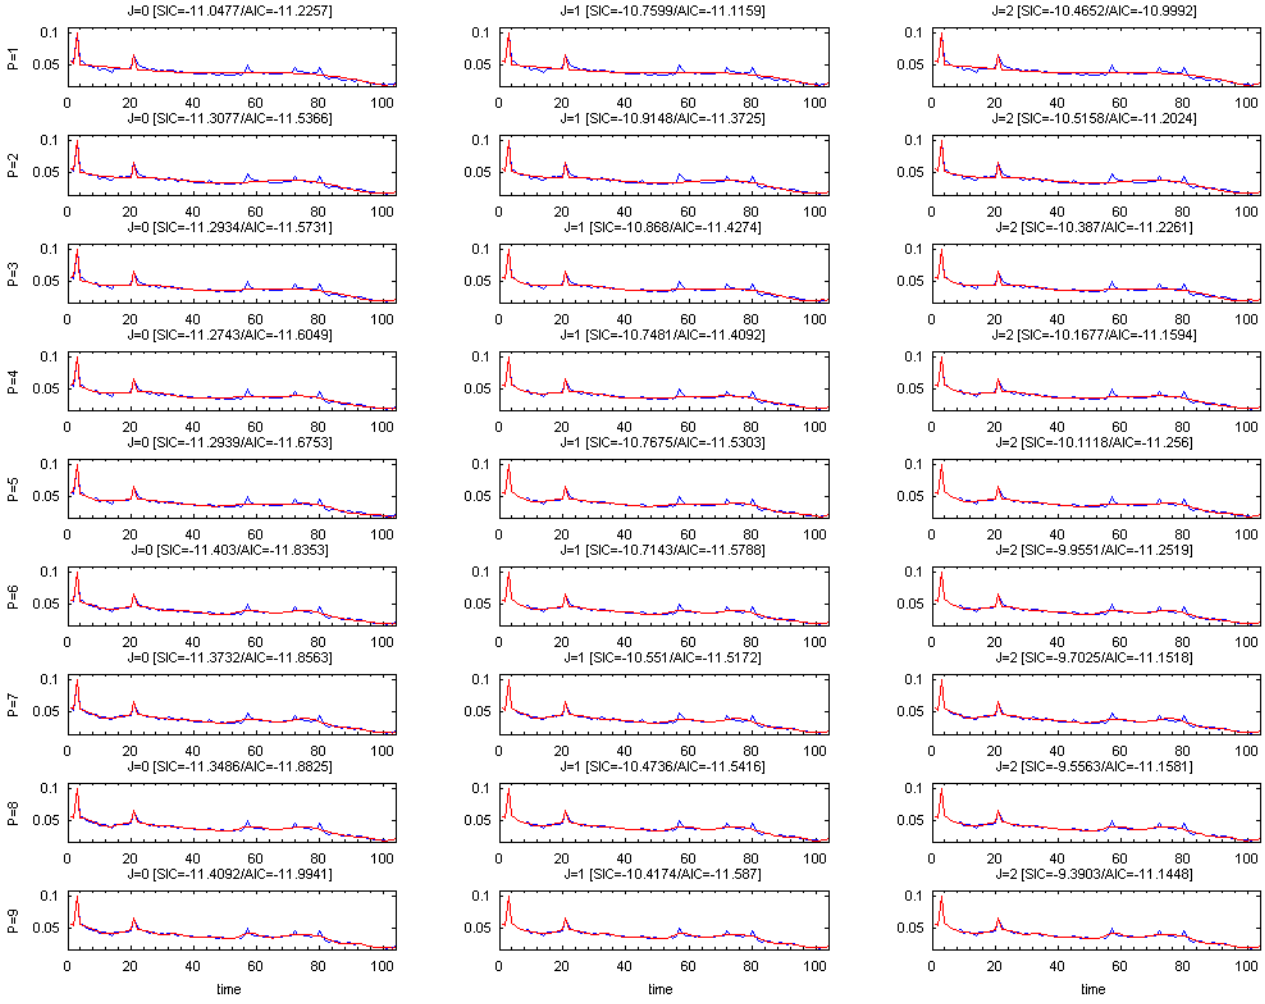

Figure B.4: Fitted Flexible Fourier Form on mean absolute US 30 year treasury bond return data for different values of J and P accompanied by SIC and AIC values. The fit gradually improves as the amount of parameters increase. Note however that it is not the objective to track the ex-post volatility to our best capabilities but to make sensible forecasts. Tracking with to many parameters would induce noise in the forecast and enlarge parameter estimation uncertainty.

Mincer-Zarnowitz regression coefficients

The Mincer-Zarnowitz regression is generally a simple ordinary least squares regression of the forecast on the (realized) true value, taking the form $R_{realization} = \alpha + \beta P_{prediction} + \varepsilon_t$. In our case however it is replaced by the more robust GLS counterpart. Al the same, conjunction of the estimates $\hat{\alpha}$ and $\hat{\beta}$ gives us some information about how accurate the prediction was. Aim is to have $\alpha = 0$ and $\beta = 1$ as we'd than have a perfect prediction in the sense that it is unbiased. Both requirements should therefore simultaneously be tested using a F-test. Too visualize the setting take the following diagram from Mincer and Zarnowitz (1969). Additional it gives clear insight in how the much used Mean Squared Error is composed, $MSE = E\left[(R-P)^2\right] = \frac{1}{n}\sum_{i=1}^{n}\left(\varepsilon_i^2\right)$.

LPF	Line of perfect forecast, $R_{realization}$
RL	Regression line, $\hat{\alpha} + \hat{\beta} P_{prediction}$
$\overline{R}$	Mean Realization
$\overline{P}$	Mean Prediction
E	Mean Point
$\hat{\alpha}$	Estimated error constant
$\hat{\beta}$	Estimated slope error

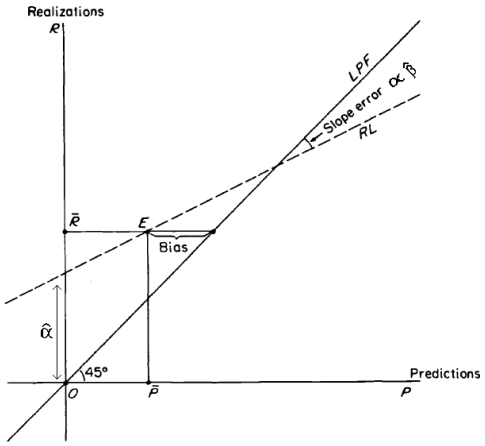

Figure B.5: Mincer-Zarnowitz regression coefficients - Prediction/Realization diagram.

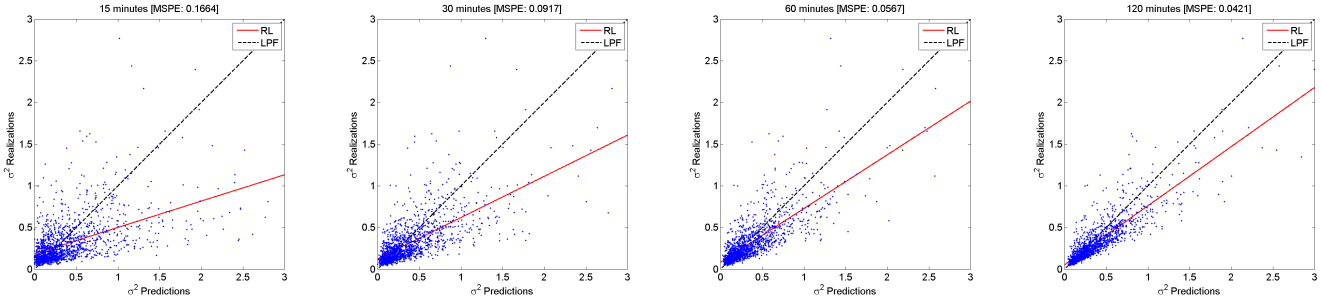

Figure B.6: Mincer-Zarnowitz diagram example for variance forecasts using FFF variance estimation on US 30 year bond futures return data. Notice that the bias and therefore MSPE tends to decline as one moves further from start of open outcry trade, i.e. 15, 30, 60, 120 minutes in.

R^2

As mentioned the coefficient of determination, R^2, from Mincer-Zarnowitz regressions should be adjusted for the fact that volatility in itself cannot be observed. Consequently the R^2 from MZ-regressions has to be adjusted. To see this, take the following infeasible and feasible Mincer-Zarnowitz regressions:

$$MZ^{IV}: \ IV_t = \gamma + \delta I\tilde{V}_t + \varepsilon_t \tag{B.23}$$

$$MZ^{RV}: \ RV_t = \alpha + \beta I\tilde{V}_t + \varepsilon_t \tag{B.24}$$

Where RV can again be replaced by any other volatility measure. Instead of the infeasible first regression we would like to perform, we are bound to use the feasible second one. The R^2 which we subsequently calculate is given by

$$R^2_{RV} = \frac{\hat{\beta}^2 \sum_{t=H}^{T} \left(I\tilde{V}_t - \overline{I\tilde{V}}\right)^2}{\sum_{t=H}^{T} \left(RV_t - \overline{RV}\right)^2}, \text{ instead of: } \quad R^2_{IV} = \frac{\hat{\delta}^2 \sum_{t=H}^{T} \left(I\tilde{V}_t - \overline{I\tilde{V}}\right)^2}{\sum_{t=H}^{T} \left(IV_t - \overline{IV}\right)^2} \tag{B.25}$$

Where $\hat{\delta}$ and $\hat{\beta}$ are OLS estimates from regression (B.23) and (B.24) respectively, and $\overline{I\tilde{V}}$, $\overline{IV}$ and $\overline{RV}$ are the averages of respectively $I\tilde{V}_t$, IV_t and RV_t. Thus to obtain R^2_{IV} we need to multiply R^2_{RV} by the factor

$$\frac{\sum_{t=1}^{T} \left(RV_t - \overline{RV}\right)^2}{\sum_{t=1}^{T} \left(IV_t - \overline{IV}\right)^2} \cdot \left(\frac{\hat{\delta}}{\hat{\beta}}\right)^2 \tag{B.26}$$

to accommodate for the bias. First part in this formula equals the adjustment proposed by Andersen et al. (2005) creating a partly adjusted R^2, second part, $\left(\hat{\delta}/\hat{\beta}\right)^2$, was proposed by Asai et al. (2012) creating what they call the fully corrected R^2.

Theoretically sound, empirically though, there are still a few obstacles to overcome. The needed IV cannot be observed and consequently δ cannot be estimated. To resolve this issue let us start with the estimation of the first part of equation (B.26) as suggested by Andersen et al. They argue that realized variance, or any other measure, can be written as the integrated variance plus an error term: $RV_t = IV_t + \varepsilon_t$. The Variance of RV thus consists of both, variances for the two constituents and a covariance term, $Var(RV_t) = Var(IV_t) + Var(\varepsilon_t) + 2Cov(IV_t, \varepsilon_t)$ where $Var(\varepsilon_t) = 2nE\left[RQ_t\right] + O(n)$, n is the number of daily increments, RQ is the realized quarticity

and $O(n)$ is of order n, generally to be neglected. Taking the results of Andersen et al. (2002) the covariance term is generally small and therefore also negligible. Hence the variability of IV will on average be overstated by $2nE\left[RQ_t\right]$. In particular they state that, ignoring the $O(n)$ term, the R^2 is on average underestimated by the factor $Var(RV_t)/\left(Var(RV_t) - 2\frac{1}{n}E\left[RQ_t\right]\right)$ equaling the first part of equation (B.26). Furthermore they state that RQ_t can consistently be estimated by $n \cdot \frac{1}{3} \cdot \sum_{i=1}^{n} r_{t,i}^4$ completing the result in a practical sense. Note must however be made that for their results to be confidently accurate the sample frequency, n, is recommended to be greater than 48. That is one measurement every 8:26 minutes for a normal open outcry 'S&P500' day, once every 10:50 minutes for the 'US30' day or once every 30 minutes for a 24 hour day. All well within our range.

For the second part one can follow the procedure of Asai et al. (2012). They countered that the partly corrected R^2 of Andersen et al. in turn overestimates the true R^2, sometimes even exceeding one. They show algebraically that the second part of (B.26) is formally needed to correct this. However δ cannot be estimated from regression (B.23) as IV is still unobserved. To work around this, take again $RV_t = IV_t + \varepsilon_t$ and write the estimated MZ-coefficient as

$$\hat{\delta} = \frac{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)\left(IV_t - \overline{IV_t}\right)}{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)^2} \text{ and replace } IV_t \text{ by } RV_t - \varepsilon_t$$

$$\hat{\delta} = \underbrace{\frac{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)\left(RV_t - \overline{RV_t}\right)}{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)^2}}_{\hat{\beta}} - \frac{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)\varepsilon_t}{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)^2} + \underbrace{\frac{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)\overline{\varepsilon}_t}{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)^2}}_{0} \quad \rightarrow$$

$$\hat{\delta} = \hat{\beta} - \frac{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)\varepsilon_t}{\sum\left(\tilde{IV}_t - \overline{\tilde{IV}}\right)^2} \tag{B.27}$$

Where the average error, $\overline{\varepsilon}_t$, is assumed zero. This leaves us with a last unknown, ε_t. The procedure is as follows: (1) create estimates for β through equation (B.24) and calculate the regression R^2; (2) multiply R^2 with the factor $Var(RV_t)/\left(Var(RV_t) - 2\frac{1}{n}\left[n \cdot \frac{1}{3} \cdot \sum_{i=1}^{n} r_{t,i}^4\right]\right)$ to obtain the partly adjusted R^2; (3) create estimates for ε_t through conducting filtering techniques[34]; (4) replace ε_t by its estimate $\hat{\varepsilon}_t$ in equation (B.27); (5) multiply the partly adjusted R^2 by the estimate of $\left(\hat{\delta}/\hat{\beta}\right)^2$ creating the fully corrected R^2.

[34] Asai et al. (2012) recommend particle filters when volatility models are considered since the distribution of ε is highly likely non Gaussian. The use and implementation of such filters for this purpose is however beyond the scope of this thesis.

The Kalman Filter

The Kalman Filter (KF) as first developed by Rudolf Emil Kálmán in the article Kalman (1960) is a highly popular recursive estimation algorithm used to obtain estimates of unobserved states through noisy measurements. Its popularity comes from its recursive nature making it computationally attractive; the fact that it uses all available information using an optimal weighing between signals; the relative ease of implementation and its excellent performance. In our setting the filter is used to model and predict the unobserved state variables α and β as measured through a Mincer-Zarnowitz type regression.

Before we jump into the specifics of implementation lets begin with some background information on the Kalman Filter, creating a sense of what is actually going on and why it works so well. To do this we take an easy to visualize example, take for instance a car and its position. The cars position at time t is not known with certainty and can only be measured with some error. Starting in state 1, the car drives a bit. Due to general laws of physical we are able to predict where the car will be at time $t+1$. This new state, divided from state 1 by the mean velocity of the car multiplied by the time interval, is not known exactly due to some error is the evolution of states. Think of this as wheels that might be slipping, the speed not being displayed 100% accurate and so on. In short there is a newly predicted state, the conditional mean, surrounded by insecurity caught in a state estimation covariance creating the whole conditional probability density function (pdf). Luckily we then get a new measurement from the cars on-line GPS system. The difference from the GPS measurement and the earlier predicted measurement is used to obtain a measure for the insecurity surrounding this new signal. A weighing function, called the Kalman gain, is then used to combine the two precarious signals, the estimation and the new measurement, to one optimal position estimate having smaller variance than either of the two signals. Graphically this process can be illustrated by figure B.7.

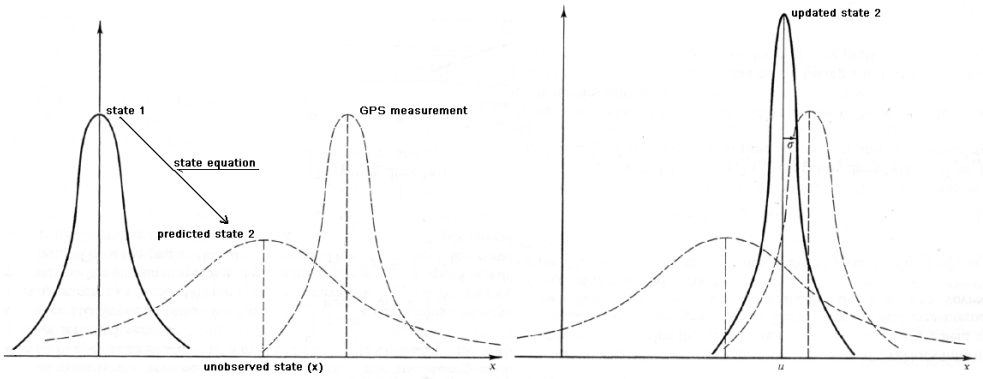

Figure B.7: The figure gives a graphical representation on how the recursive Kalman Filter algorithm calculates the optimal new state by joining a noisy measurement with a precarious state prediction. From the updated state 2 the next cycle of the filter is entered.

Additional pleasing feature of the KF algorithm is that all estimated probability density functions are conditional on the past in a recursive way. Making the mathematics tractable and computationally fast as just one new data point has to be reviewed every time. For a very intuitive and more thorough exposition of the Kalman Filter one could read the first chapter of Maybeck (1979).

Now that the general idea is somewhat clear let us proceed with the more specific situation at hand. Before we initiate state estimation some definitions and equations have to be set. First and foremost are the state equations, governing the evolution of states through time, and the measurement equation revealing how these states are connected to potentially noisy measurements. As this is a first approach in modeling these states a simple but potentially restrictive first order autoregressive model, 'AR(1)', is proposed for both states. These can therefore be written as:

$$\alpha_{t,n^*} = c_{\alpha,n^*} + \varphi_{1,n^*}\alpha_{t-1,n^*} + \zeta_{t,n^*} \text{ for } n^* = 1, ..., n \text{with } \zeta_t \sim N(0, Q_\zeta) \tag{B.28}$$

$$\beta_{t,n^*} = c_{\beta,n^*} + \varphi_{2,n^*}\beta_{t-1,n^*} + \eta_{t,n^*} \text{ for } n^* = 1, ..., n \text{with } \eta_t \sim N(0, Q_\eta) \tag{B.29}$$

Or more conveniently abbreviated in state-space notation as:

$$x_t = C + Fx_{t-1} + \omega_t \text{ with } \omega_t \sim N \begin{pmatrix} 0, & Q_\zeta & 0 \\ 0, & 0 & Q_\eta \end{pmatrix} \tag{B.30}$$

with

$$C\,(2n\times 1)=\begin{bmatrix} c_{\alpha,1}\\ \vdots\\ c_{\alpha,n}\\ c_{\beta,1}\\ \vdots\\ c_{\beta,n}\end{bmatrix},\qquad F\,(2n\times 2n)=\begin{bmatrix} \varphi_{1,1} & 0 & \cdots & 0\\ 0 & \ddots & & \\ & & \varphi_{1,n} & \ddots & \vdots\\ \vdots & & \ddots & \varphi_{2,1} & \\ & & & & \ddots & 0\\ 0 & & \cdots & & 0 & \varphi_{2,n}\end{bmatrix},$$

$$\omega_t\,(2n\times 1)=\begin{bmatrix} \zeta_{t,1}\\ \vdots\\ \zeta_{t,n}\\ \eta_{t,1}\\ \vdots\\ \eta_{t,n}\end{bmatrix},\qquad Q\,(2n\times 2n)=\begin{bmatrix} \sigma^2_{\zeta,1} & \cdots & \sigma_{\zeta,1}\sigma_{\zeta,n} & 0 & \cdots & 0\\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots\\ \sigma_{\zeta,n}\sigma_{\zeta,1} & \cdots & 0\sigma_{\zeta,n}\sigma_{\zeta,n} & 0 & \cdots & 0\\ 0 & \cdots & 0 & \sigma^2_{\eta,1} & \cdots & \sigma_{\eta,1}\sigma_{\eta,n}\\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots\\ 0 & \cdots & 0 & \sigma_{\eta,n}\sigma_1 & \cdots & \sigma_{\eta,n}\sigma_{\eta,n}\end{bmatrix}$$

Where x_t is a vector, containing unobserved state $\{\alpha_{t,n^*},\beta_{t,n^*}\}'$and ω_t is an i.i.d. random normal error term with mean zero and block-diagonal covariance matrix $\{Q_\zeta,Q_\eta\}$. This process noise covariance Q can be seen as a measure of the uncertainty in the state dynamics during the time interval between measurement updates. See for example figure B.7, the variance of the predicted state 2 is greater than was the case for (updated) state 1 due to this extra insecurity.

Second is the measurement equation which relates the measurements to the estimated process states. For this purpose we take equation (31) and rewrite this to a state-space form.

$$RV_t=\alpha_{t,n^*}+\beta_{t,n^*}\cdot I\tilde{V}_{t,n^*}+v_{t,n^*}\text{ for }n^*=1,...,n \tag{B.31}$$

$$\text{with }I\tilde{V}_{t,n^*}=RV_t^{n^*}\frac{\sum_{i=1}^{n}s_{t,i}}{\sum_{i=1}^{n^*}s_{t,i}}\text{ and }v_t\sim N\,(0,R)$$

Put in state-space form:

$$z_t=H_tx_t+v_t\text{ with }v_t\sim N\,(0,R) \tag{B.32}$$

$$H_t\,(n \times 2n) = \begin{bmatrix} 1 & 0 & 0 & \tilde{IV}_{t,1} & 0 & 0 \\ 0 & \ddots & 0 & 0 & \ddots & 0 \\ 0 & 0 & 1 & 0 & 0 & \tilde{IV}_{t,n} \end{bmatrix},\; x_t\,(2n \times 1) = \begin{bmatrix} \alpha_{t,1} \\ \vdots \\ \alpha_{t,n} \\ \beta_{t,1} \\ \vdots \\ \beta_{t,n} \end{bmatrix},$$

$$v_t\,(n \times 1) = \begin{bmatrix} v_{t,1} \\ \vdots \\ v_{t,n} \end{bmatrix},\; R\,(n \times n) = \begin{bmatrix} \sigma^2_{v,1} & \cdots & \sigma_{v,1}\sigma_{v,n} \\ \vdots & \ddots & \vdots \\ \sigma_{v,n}\sigma_{v,1} & \cdots & \sigma_{v,n}\sigma_{v,n} \end{bmatrix}$$

Where $RV_t^{n^*}$ is the measurement of a volatility measure on day t up to time n^*; $\tilde{IV}_{t,n^*}$ is the predicted end-of-the-day volatility, $s_{t,i}$ is the deterministic seasonal pattern from FFF or (Exponentially Weighted) Moving Average seasonal; $v_{t,i}$ is assumed a mean zero, i.i.d. random normal variable with covariance matrix R assumed independent of ω_t; H the transition matrix and z_t the end-of-the-day variance measurements as obtained by various volatility measures like RV. Essentially equation (B.30) and (B.32) already summarize the equations needed for the Kalman Filter, the rest is more straightforward yet insightful and will be worked out below.

Next section reveals the working of updates and predictive iterations by the Kalman Filter. It is actually nothing more than a mathematical exposition of what was already made clear in figure B.7. Lets start with a state prediction, or one step ahead forecast. By equation (B.30) it is easily shown that the predicted state (conditional mean), conditional on the current state is given by

$$\hat{x}_{t|t-1} = C + F\hat{x}_{t-1|t-1} \tag{B.33}$$

The predicted covariance of the state estimate equals

$$P_{t|t-1} = FP_{t-1|t-1}F' + Q \tag{B.34}$$

That is the former corrected covariance pre- and post-multiplied by the transition matrix, increased by the variance from state equation insecurity. So now we have predicted the conditional mean and the covariance of the state, thereby obtaining the whole conditional pdf of the forecasts for a one-step ahead future states. Next to the algorithm is the updating step. We start with the measurement residual, being the difference between the measurement forecast from the measurement equation and the actual measurement $z_t = RV_t$. Subsequent the residual covariance is

measured by (B.36).

$$e_t = z_t - H_t \hat{x}_{t|t-1} \tag{B.35}$$

$$S_t = H_t P_{t|t-1} H_t' + R \text{ with } R = cov\left(v_{t,i}\right) \tag{B.36}$$

Above equations mainly yield insight in the accuracy of the measurements. These quantities will be needed to assess how much weight must be attached to future measurement predictions. In the formula for the actual weigh-off, the so called Kalman gain, the inverse of the covariance of measurement residuals is therefore taken.

$$K_t = P_{t|t-1} H_t' S_t^{-1} \tag{B.37}$$

This function thus directs optimal weight to the new signal/measurement, optimal in the sense of smallest MSPE, conditional on its past performance as measured by residual covariance. That is, the bigger the residual covariance, the smaller the Kalman gain and the less weight is put to the new measurement update. If on the other hand S_t is small, there is strong belief in the new measurement and high weight is directed toward its value. In the extreme case when the estimation error covariance explodes the Kalman gain converges to zero and visa versa. With that all the ingredients to create an updated state estimate are obtained. The update can be calculated as

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + K_t e_t \tag{B.38}$$

and the updated covariance estimate becomes

$$P_{t|t} = \left(I - K_t H_t\right) P_{t|t-1} \tag{B.39}$$

Completing the circle for a next recursive step. However in order to initiate the recursions one must first obtain initial values and as research in other directions of science yet on the Kalman Filter has shown that poor estimates of the input noise statistics may seriously degrade Kalman performance and even provoke divergence of the filter, see Fitzgerald (1971); Sangsuk-lam and Bullock (1990), these must be selected sensibly. Therefore the first year of observations is taken as estimation sample to obtain initial values. For the problem at hand this will be done in the

following way.

By use of an arbitrarily chosen 20 day moving window for Mincer-Zarnowitz regressions we approximated time series for alpha and beta. Associated means can then be obtained. Last value of these time series can be used as initial states, μ_0. The time series themselves can be used to estimate an AR(1) model. thereby one can obtain initial values for C and F based on the first year and obtain Q_ζ and Q_η from regression residuals. Once α and β time series are obtained these can also be used to gain proxy's for the measurement error from equation (B.32) and subsequently produce the measurement error covariance R_0. Last remaining initial value is Σ_0 which we initiate at $1e^{-3}I_{[n\times n]}$. The choice of this parameter is of somewhat arbitrary yet of lesser concern as the algorithm adjusts it and it is found to converge as long as the initial value is not taken to be zero.

Summarizing we get:

Summing Kalman Filter and Kalman Smoother variables:	
x_t	$2n \times 1$ - State vector
z_t	$n \times 1$ - Observation vector
H_t	$n \times 2n$ - Linear transition matrix
C	$2n \times 1$ - vector of unconditional state constants
F	$2n \times 2n$ - Process noise covariance matrix
Q	$2n \times 2n$ - Process noise covariance matrix
R	$n \times n$ - Measurement noise covariance matrix

Table B.1: Kalman Filter and Smoother variables

Kalman Filter Equations	
Model forecast step/prediction-step:	$\hat{x}_{t\|t-1} = C + F\hat{x}_{t-1\|t-1}$
	$P_{t\|t-1} = FP_{t-1\|t-1}F' + Q$
Data assimilation step/update-step:	$e_t = z_t - H_t\hat{x}_{t\|t-1}$
	$S_t = H_tP_{t\|t-1}H_t' + R$
	$K_t = P_{t\|t-1}H_t'S_t^{-1}$
	$\hat{x}_{t\|t} = \hat{x}_{t\|t-1} + K_te_t$
	$P_{t\|t} = (I - K_tH_t)\, P_{t\|t-1}$

Table B.2: Kalman Filter Equations. An recursive algorithm to reveal unobserved time series from noisy measurements and the state equation. The filter is optimal in a mean squared prediction error (MSPE) kind of sense.

Kalman Smoother Equations
$u_{t-1} = H_t' S_t \left(z_t - H_t \hat{x}_{t\|t-1}\right) + L_t' u_t$ with $L_t = F\left(I - K_t H_t\right)$
$U_{t-1} = H_t' S_t H_t + L_t' U_t L_t$
$\hat{x}_{t\|T} = \hat{x}_{t\|t-1} + P_{t\|t-1} u_{t-1}$
$P_{t\|T} = P_{t\|t-1} + P_{t\|t-1} U_{t-1} P_{t\|t-1}$
$P_{t,t-1\|T} = \left(I - P_{t\|t-1} U_{t-1}\right) L_{t-1} P_{t-1\|t-2}$
with final calculations for t=0:
$\hat{x}_{0\|T} = \hat{x}_0 + P_0 F' u_0$
$P_{0\|T} = P_0 - P_0 F' U_0 F P_0$
$P_{1,0\|T} = \left(I - P_{0\|1} U_0\right) F P_0$

Table B.3: Kalman Smoothing Equations. Basically smoothing the filter estimated states through a backward recursion of the filter. Estimating the optimal states, in the sense of mean squared error (MSE), having already observed all the data, i.e. $E\left[x_t | Z_T\right]$. Note that we made use of the following starting conditions: $u_T = 0$ and $U_T = 0$. Furthermore $P_{t,t-1|T}$ has been included being the lag-one covariance smoother needed for the Expectation Maximization algorithm.

Unconstrained EM M-step derivations

Complete data likelihood function:

$$-2lnL\left(X,Z;\theta\right)=ln\left(\left|\Sigma_0\right|\right)+\left(x_0-\mu_0\right)'\Sigma_0^{-1}\left(x_0-\mu_0\right)+$$
$$T\cdot ln\left(\left|Q\right|\right)+\sum_{t=1}^{T}\left(x_t-C-Fx_{t-1}\right)'Q^{-1}\left(x_t-C-Fx_{t-1}\right)+$$
$$T\cdot ln\left(\left|R\right|\right)+\sum_{t=1}^{T}\left(z_t-H_tx_t\right)'R^{-1}\left(z_t-H_tx_t\right)$$

Conditional expectation of likelihood function under current parameters:

$$Q_{function}\left(\theta^j|Z_T,\theta^{j-1}\right)=E\left[-2lnL\left(X,Z;\theta\right)|Z_T,\theta^{j-1}\right]$$

$$Q_{function}\left(\theta^j|Z_T,\theta^{j-1}\right)=ln\left(\left|\Sigma_0\right|\right)+tr\left(\Sigma_0^{-1}\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}\right]\right)+$$
$$T\cdot ln\left(\left|Q\right|\right)+tr\left(Q^{-1}\left[S_{11}-S_1C'-S_{10}F'+T\left(CC'\right)-CS_1'+CS_0'F'-FS_1'+FS_0C'+FS_{00}F'\right]\right)+$$
$$T\cdot ln\left(\left|R\right|\right)+tr\left(R^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t-H_tx_{t|T}\right)\left(z_t-H_tx_{t|T}\right)'+H_tP_{t|T}H_t'\right)\right]\right)$$

with

$$S_0=\sum_{t=1}^{T}\left(x_{t-1|T}\right)$$

$$S_1=\sum_{t=1}^{T}\left(x_{t|T}\right)$$

$$S_{00}=\sum_{t=1}^{T}\left(x_{t-1|T}x_{t-1|T}'+P_{t-1|T}\right)$$

$$S_{10}=\sum_{t=1}^{T}\left(x_{t|T}x_{t-1|T}'+P_{t,t-1|T}\right)$$

$$S_{11}=\sum_{t=1}^{T}\left(x_{t|T}x_{t|T}'+P_{t|T}\right)$$

Where we used $x'Ax=tr\left(x'Ax\right)=tr\left(Axx'\right)$ which holds as long as x is a vector or scalar; $\sum tr\left(Axx'\right)=tr\sum\left(Axx'\right)$ and $E\left[\sum_{t=1}^{T}(x_tx_t')|Z_n;\theta\right]=\sum_{t=1}^{T}E\left[x_tx_t'|Z_n;\theta\right]=\sum_{t=1}^{T}\left(P_{t|T}+x_{t|T}x_{t|T}'\right)$. The latter property follows from the covariance definition: $cov\left(x,x\right)=E\left[\left(x-E\left(x\right)\right)\left(x-E\left(x\right)\right)'\right]=E\left[xx'\right]-E\left[x\right]E\left[x\right]'\rightarrow E\left(xx'\right)=cov\left(x,x\right)+E\left(x\right)E\left(x\right)'$ and results in the extra conditional covariance terms in the Q-function. To see this we work out the last 'R' term line of the Complete data likelihood function. The rest takes some paperwork but follows in much the same way.

$$E\left[T\cdot ln\left(\left|R_{\left(\theta^{j-1}\right)}\right|\right)+\sum_{t=1}^{T}\left(z_t-H_tx_t\right)'R_{\left(\theta^{j-1}\right)}^{-1}\left(z_t-H_tx_t\right)|Z_T,\theta^{j-1}\right]=$$

$$E\left[T\cdot ln\left(\left|R_{\left(\theta^{j-1}\right)}\right|\right)+tr\left(R_{\left(\theta^{j-1}\right)}^{-1}\sum_{t=1}^{T}\left(z_t-H_tx_t\right)\left(z_t-H_tx_t\right)'\right)|Z_T,\theta^{j-1}\right]=$$

$$E\left[T\cdot ln\left(\left|R_{\left(\theta^{j-1}\right)}\right|\right)+tr\left(R_{\left(\theta^{j-1}\right)}^{-1}\sum_{t=1}^{T}\left(z_tz_t'-z_tx_t'H_t'-H_tx_tz_t'+H_tx_tx_t'H_t'\right)\right)|Z_T,\theta^{j-1}\right]=$$

now noting that $E\left[x_tx_t'|Z_T,\theta^{j-1}\right]=cov\left(x_{t|T},x_{t|T}\right)+E\left[x_t|Z_T,\theta^{j-1}\right]E\left[x_t|Z_T,\theta^{j-1}\right]'=P_{t|T}+x_{t|T}x_{t|T}'$

$$T\cdot ln\left(\left|R_{\left(\theta^{j-1}\right)}\right|\right)+tr\left(R_{\left(\theta^{j-1}\right)}^{-1}\sum_{t=1}^{T}\left(z_tz_t'-z_tx_{t|T}'H_t'-H_tx_{t|T}z_t'+H_tx_{t|T}x_{t|T}'H_t'+H_tP_{t|T}H_t'\right)\right)=$$

$$T\cdot ln\left(\left|R_{\left(\theta^{j-1}\right)}\right|\right)+tr\left(R_{\left(\theta^{j-1}\right)}^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t-H_tx_{t|T}\right)\left(z_t-H_tx_{t|T}\right)'+H_tP_{t|T}H_t'\right)\right]\right)= \quad \text{(B.40)}$$

Yielding the last line from the Q-function. Note that we put some extra emphasis on the conditionality of parameter R on the last iteration, θ^{j-1}.

In order to obtain the minimum of the $Q_{function}$ one has to calculate the first order partial derivatives with respect to the parameters $\theta=\{\mu_0,\Sigma_0,C,F,Q,R\}$, and zero this partial derivative to gain the optimal analytical solution.

$$\frac{\delta Q_{function}}{\delta\mu_0}=\frac{\delta\left[tr\left(\Sigma_0^{-1}\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'\right)\right]}{\delta\mu_0}=$$

$$\Sigma_0^{-1}\left(x_{0|T}-\mu_0\right)+\left(\Sigma_0^{-1}\right)'\left(x_{0|T}-\mu_0\right)=$$

$$2\Sigma_0^{-1}\left(x_{0|T}-\mu_0\right)=0$$

$$\mu_0=x_{0|T}$$

$$\frac{\delta Q_{function}}{\delta\Sigma_0}=\frac{\delta\left[ln\left(|\Sigma_0|\right)+tr\left(\Sigma_0^{-1}\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}\right]\right)\right]}{\delta\Sigma_0}=$$

$$\left(\Sigma_0^{-1}\right)'-\left(\Sigma_0^{-1}\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}\right]\Sigma_0^{-1}\right)'=$$

$$\Sigma_0'\left(\Sigma_0^{-1}\right)'\Sigma_0'-\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}\right]'=$$

$$\Sigma_0'-\left[\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}\right]'=0$$

$$\Sigma_0=\left(x_{0|T}-\mu_0\right)\left(x_{0|T}-\mu_0\right)'+P_{0|T}$$

where the symmetry property of the covariance matrix is used, i.e. $\Sigma_0'=\Sigma_0$

$$\frac{\delta Q_{function}}{\delta C} = \frac{\delta\left[tr\left(Q^{-1}\left[-S_1C' + T\left(CC'\right) - CS_1' + CS_0'F' + FS_0C'\right]\right)\right]}{\delta C} =$$
$$-\frac{\delta\left[tr\left(Q^{-1}S_1C'\right)\right]}{\delta C} + \frac{\delta\left[tr\left(Q^{-1}T\left(CC'\right)\right)\right]}{\delta C} - \frac{\delta\left[tr\left(Q^{-1}CS_1'\right)\right]}{\delta C} + \frac{\delta\left[tr\left(Q^{-1}CS_0'F'\right)\right]}{\delta C} + \frac{\delta\left[tr\left(Q^{-1}FS_0C'\right)\right]}{\delta C} =$$
$$-Q^{-1}S_1 + 2TQ^{-1}C - Q^{-1}S_1 + Q^{-1}FS_0 + Q^{-1}FS_0 =$$
$$-2Q^{-1}S_1 + 2TQ^{-1}C + 2Q^{-1}FS_0 =$$
$$-Q^{-1}S_1 + TQ^{-1}C + Q^{-1}FS_0 =$$

$$-S_1 + TC + FS_0 = 0 \qquad \rightarrow \qquad C = T^{-1}\left(S_1 - FS_0\right)$$

$$C = T^{-1}\left(S_1 - \left(S_{10} - T^{-1}S_1S_0'S_{00}^{-1}\right)\left(I - T^{-1}S_0S_0'S_{00}^{-1}\right)^{-1}S_0\right)$$

$$\frac{\delta Q_{function}}{\delta F} = \frac{\delta\left[tr\left(Q^{-1}\left[-S_{10}F' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right]\right)\right]}{\delta F} =$$
$$-\frac{\delta\left[tr\left(Q^{-1}S_{10}F'\right)\right]}{\delta F} + \frac{\delta\left[tr\left(Q^{-1}CS_0'F'\right)\right]}{\delta F} - \frac{\delta\left[tr\left(Q^{-1}FS_{10}'\right)\right]}{\delta F} + \frac{\delta\left[tr\left(Q^{-1}FS_0C'\right)\right]}{\delta F} + \frac{\delta\left[tr\left(Q^{-1}FS_{00}F'\right)\right]}{\delta F} =$$
$$-Q^{-1}S_{10} + Q^{-1}CS_0' - Q^{-1}S_{10} + Q^{-1}CS_0' + \left(S_{00}F'Q^{-1}\right)' + Q^{-1}FS_{00} =$$
$$-2Q^{-1}S_{10} + 2Q^{-1}CS_0' + 2Q^{-1}FS_{00} =$$
$$S_{10} + CS_0' + FS_{00} = 0$$

$$S_{10} + CS_0' + FS_{00} = 0 \qquad \rightarrow \qquad F = \left(S_{10} - CS_0'\right)S_{00}^{-1}$$

$$F = \left(S_{10} - T^{-1}S_1S_0'S_{00}^{-1}\right)\left(I - T^{-1}S_0S_0'S_{00}^{-1}\right)^{-1}$$

where the symmetry property of the S_{00} matrix is used, i.e. $S_{00}' = S_{00}$

$$\frac{\delta Q_{function}}{\delta Q} =$$
$$\frac{\delta T \cdot ln\left(|Q|\right) + tr\left(Q^{-1}\left[S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right]\right)}{\delta Q} =$$
$$T\left(Q^{-1}\right)' - \left(Q^{-1}\left[S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right]Q^{-1}\right)' =$$
$$T\left(QQ^{-1}Q\right)' - \left(QQ^{-1}\left[S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right]Q^{-1}Q\right)' =$$
$$TQ' - \left(S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right)' = 0$$

$$Q = T^{-1}\left(S_{11} - S_1C' - S_{10}F' + T\left(CC'\right) - CS_1' + CS_0'F' - FS_{10}' + FS_0C' + FS_{00}F'\right)$$

$$\frac{\delta Q_{function}}{\delta R} = \frac{\delta\left[T \cdot ln\left(|R|\right) + tr\left(R^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t - H_t x_{t|T}\right)\left(z_t - H_t x_{t|T}\right)' + H_t P_{t|T} H_t'\right)\right]\right)\right]}{\delta R} =$$

$$T \cdot \left(R^{-1}\right)' - \left(R^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t - H_t x_{t|T}\right)\left(z_t - H_t x_{t|T}\right)' + H_t P_{t|T} H_t'\right)\right] R^{-1}\right)' =$$

$$T \cdot \left(RR^{-1}R\right)' - \left(RR^{-1}\left[\sum_{t=1}^{T}\left(\left(z_t - H_t x_{t|T}\right)\left(z_t - H_t x_{t|T}\right)' + H_t P_{t|T} H_t'\right)\right] R^{-1}R\right)' =$$

$$T \cdot R' - \left(\sum_{t=1}^{T}\left(\left(z_t - H_t x_{t|T}\right)\left(z_t - H_t x_{t|T}\right)' + H_t P_{t|T} H_t'\right)\right)' = 0$$

$$R = T^{-1}\sum_{t=1}^{T}\left(\left(z_t - H_t x_{t|T}\right)\left(z_t - H_t x_{t|T}\right)' + H_t P_{t|T} H_t'\right)$$

To obtain the above calculated partial derivatives of the Q-function a few matrix derivative rules were used: $tr\left(A+B\right) = tr\left(A\right) + tr\left(B\right)$, $\frac{\delta ln|X|}{\delta X} = \left(X^{-1}\right)'$, $\frac{\delta A'XA}{\delta X} = AA'$, $\frac{\delta X'AX}{\delta X} = AX + A'X$, $\frac{\delta tr\left(X^{-1}B\right)}{\delta X} = -\left(X^{-1}BX^{-1}\right)'$, $\frac{\delta tr(XA)}{\delta X} = A'$, $\frac{\delta tr(X'A)}{\delta X} = A$,$\frac{\delta tr(AXBX)}{\delta X} = B'X'A' + A'X'B'$.

Newton Raphson Algorithm

Simple custom made Newton Raphson algorithm using the marginal gradients (Jacobian) and fixed step-size. The second order partial numerical derivative (Hessian) was replaced by an easier decreasing step-size to save computational time. Boundaries are taken account for and the likelihood is checked for imaginary parts and positive definiteness of covariance matrices. Furthermore steps are only taken when the likelihood f(ϑ) decreases thus minimizing $-2lnL\left(X, Z; \theta\right)$.

Algorithm. *pseudocode - Newton Raphson Algorithm*

```
while abs ( f(ϑinit) - f(ϑ) ) > tolerance
    for i=1 : parameters
        f(ϑ)=f(ϑo)
        calculate gradient: grad(i)=( f(ϑo(i) + step ) - f(ϑ) ) / step
        if grad(i)≠imaginary & gradient < 0
            if ( ϑo(i) + step ) < upper-boundary(i)
                if ( fϑnew < fϑ & fϑnew≠imaginary & covar = positive-definite)
                    ϑo(i)=ϑo(i)+step
                end
                step = step/2
            end
        elseif grad(i)≠imaginary & grad > 0
            if ( ϑo (i) - step ) > lower-boundary(i)
                if ( fϑnew < fϑ & fϑnew≠imaginary & covar = positive-definite)
                    ϑo=ϑo-step
                end
            end
            step = step/2
        end
    end
    iteration = iteration+1
end
```

Restrictions and Parametrizations

By laws of statistics it is clear that estimation insecurity, that is, the standard errors of the estimates, increase as more parameters need to be estimated on the same sample. Other way around also holds, decreasing standard errors for less parameters. Main objective of restricting and parametrizing the covariance matrix is therefore to dwarf standard errors and hopefully enhance estimation precision, that is lower MSPE, bearing in mind that parametrization also induces some extra bias. During this research a few trials were put to the test:

1. As $i \to n$ the covariance of v_t, caught in matrix R, decreases towards zero as a simple result of having already obtained all data with which the volatility measure itself was constructed (see B.8). R could therefore be parametrized by a logistic function $\left(1 - \frac{1}{1+exp-(\gamma_1+\gamma_2 n^*)} + e_{n^*}\right)$ over the diagonals of R. Noting that R must be symmetric this reduces the number of estimates from $\frac{1}{2}n\,(1+n)$ to $2n$ as we only estimate γ_1 and γ_2 instead of $\sigma_{v,i}\sigma_{v,i}$. However problems occur with this method as the parameterized covariance matrix is no longer bound to be positive definite. Consequently yielding obscure likelihood values. Second more general problem to the procedure is that as errors are highly correlated to direct neighbors (up to 0.98) eigenvalues are close to one. Consequently the ratio between maximum and minimum singular value explodes, Yielding a ill-conditioned matrix.

$$R\,(n \times n) = \begin{bmatrix} \sigma^2_{1,1}\,(\gamma_1, \gamma_2) & \sigma_2\sigma_1\,(\gamma_3, \gamma_4) & \cdots & \sigma_n\sigma_1\,(\gamma_{2n-1}, \gamma_{2n}) \\ \sigma_1\sigma_2\,(\gamma_3, \gamma_4) & \searrow & \searrow & \vdots \\ \vdots & \searrow & \searrow & \vdots \\ \sigma_1\sigma_n\,(\gamma_{2n-1}, \gamma_{2n}) & \cdots & \cdots & 0 \end{bmatrix}$$

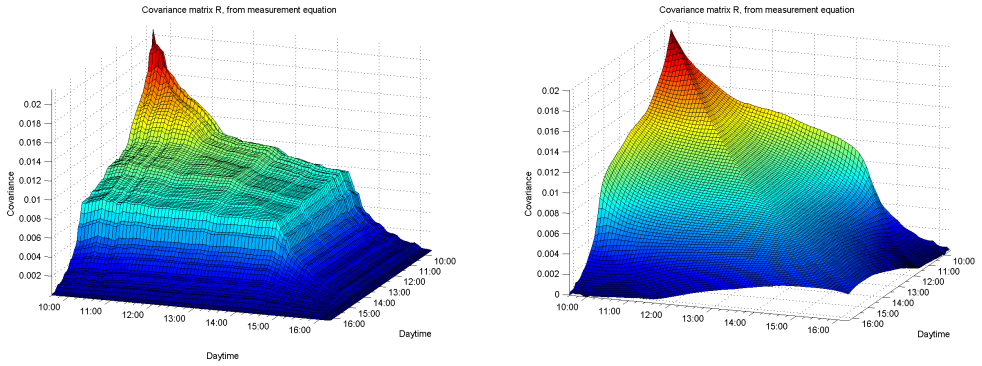

(a) Measurement Equation covariance matrix, R

(b) Fitted covariance matrix

Figure B.8: Reducing covariance of v_t as $i \to n$ for S&P500 data using FFF forecasting and RV as measure. Figure a) exhibits a noticeable ridge at approximately 14:15 p.m. Reason for this pattern remains ambiguous. Due to this ridge the fitted covariance matrix in figure b) is somewhat troubled.

2. Second trial comprehended estimating variances and complementing rows and columns with their diagonal element. Thereby reducing the number of parameters from $\frac{1}{2}n\,(1+n)$ to n. The resulting matrix could be ill-conditioned yet reducing the dimensionality such problems can be circumvented. Furthermore as long as the covariances decrease as $i \to n$, correlation boundaries are satisfied. The positive definite covariance property however cannot be guaranteed but could be checked upon initiation and withheld during estimation.

3. Third possibility is to reduce to covariance matrix to a diagonal one. This reduces the number of parameters from $\frac{1}{2}n\,(1+n)$ to n and is bound to produce all positive eigenvalues as variances need to be positive. Thereby protecting the sufficient positive definite property (which includes positive determinant).

$$R\,(n \times n) = \begin{bmatrix} \sigma^2_{i=1} & 0 & \cdots & 0 \\ 0 & \sigma^2_{i=2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma^2_{i=n} \end{bmatrix}$$

4. Fourth option to reduce insecurity and improve conditioning is by simply diminishing the dimensions of the problem. As the measurement errors are highly correlated, mounting to 0.98, taking errors time-wise further apart could help lower dependence and singularity problems during estimation.

5. Fifth option to reduce standard errors would be to obtain estimates for the unconstrained as well as a constrained covariance matrix and take an optimally weighted average of both by a procedure called shrinkage, see Ledoit and Wolf (2003) and the well readable exposition by Schafer and Strimmer (2005). This procedure has some advantages as the resulting shrinkage covariance estimate will automatically be positive definite. Moreover even if the target (constrained covariance matrix) is strongly misspecified, shrinkage will lead to a reduction of MSE; and, as analytical expressions can be obtained for the optimal shrinkage intensity, the procedure is computationally inexpensive. However a further exposition on the possible gains of such a method is left for later research.

As a workable solution to this first survey it was chosen to reduce dimensionality from 5 minute intervals to half an hour, making the matrix better conditioned. The diagonal is then estimated and the rest of the matrix complemented following the second bullet, subsequently yielding undermentioned parametrization. Furthermore an extra effort is made to obtain good initial values (better than random values between the boundaries) so that MLE converges more easily.

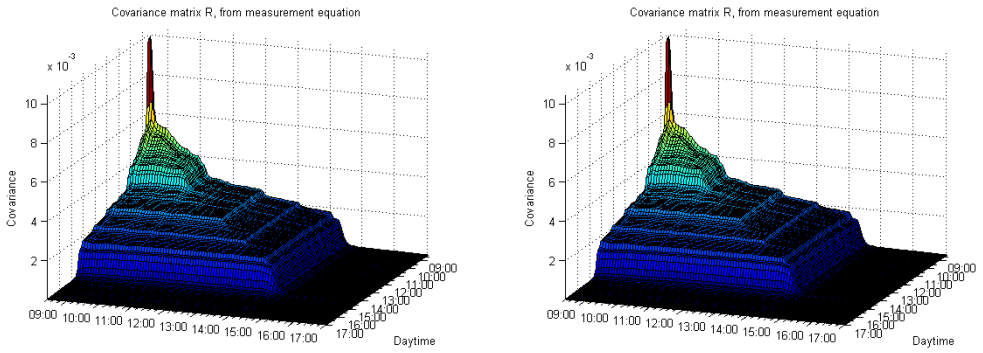

(a) Initial **(b)** Parametrized

Figure B.9: Parametrization of R based on first year estimation sample for US 30 year treasury bond futures using the RV measure and FFF forecasts.

For Q similar restrictions can be imposed though a logit function would not make sense as no exponential decay is present. Furthermore it is noted that the third option reduces the assumed SUTSE (seemingly uncorrelated timeseries estimation) model to a mere multivariate AR(1) setting, removing all mutual correlations. For Q it was chosen to reduce the dimensionality similar to R and further restricted to a diagonal matrix:

$$Q\left(2n \times 2n\right) = \begin{bmatrix} \sigma^2_{\zeta,1} & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \searrow & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma^2_{\zeta,n} & 0 & \cdots & 0 \\ 0 & \cdots & 0 & \sigma^2_{\eta,1} & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \searrow & \downarrow \\ 0 & \cdots & 0 & 0 & \cdots & \sigma^2_{\eta,n} \end{bmatrix} \tag{B.41}$$

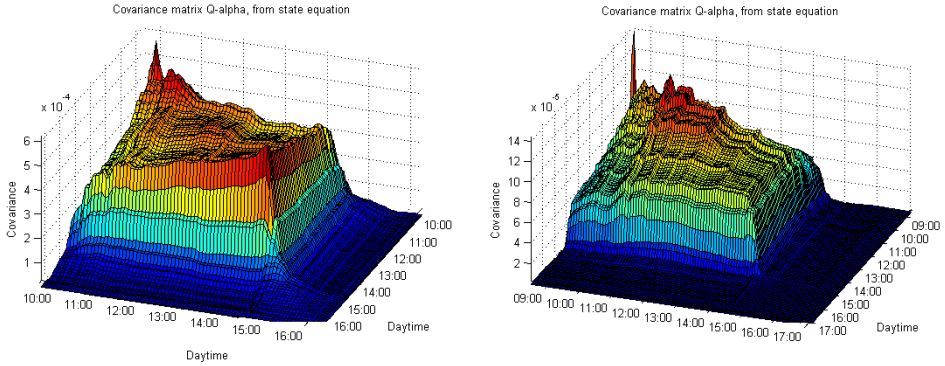

(a) Covariance state equation α: $Q_{\zeta}\left(n^*\right)$ for resp. S&P500 and US30

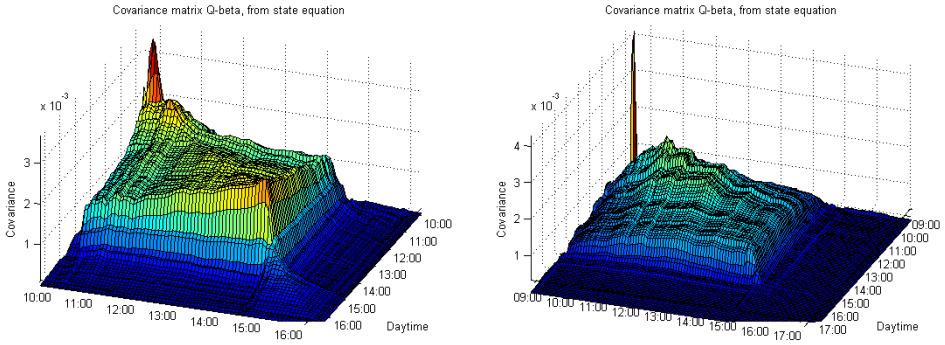

(b) Covariance state equation β: $Q_{\eta}\left(n^*\right)$ for resp. S&P500 and US30

Figure B.10: α- and β-block of blockdiagonal covariance matrix Q of state equation residual ω_t as $i \to n$ as estimated on first year estimation sample. Graphs display a far less stylized pattern than covariance matrix R.

Appendix C

Annualized standard deviation of S&P data omitting Sept, Oct, Nov 2008

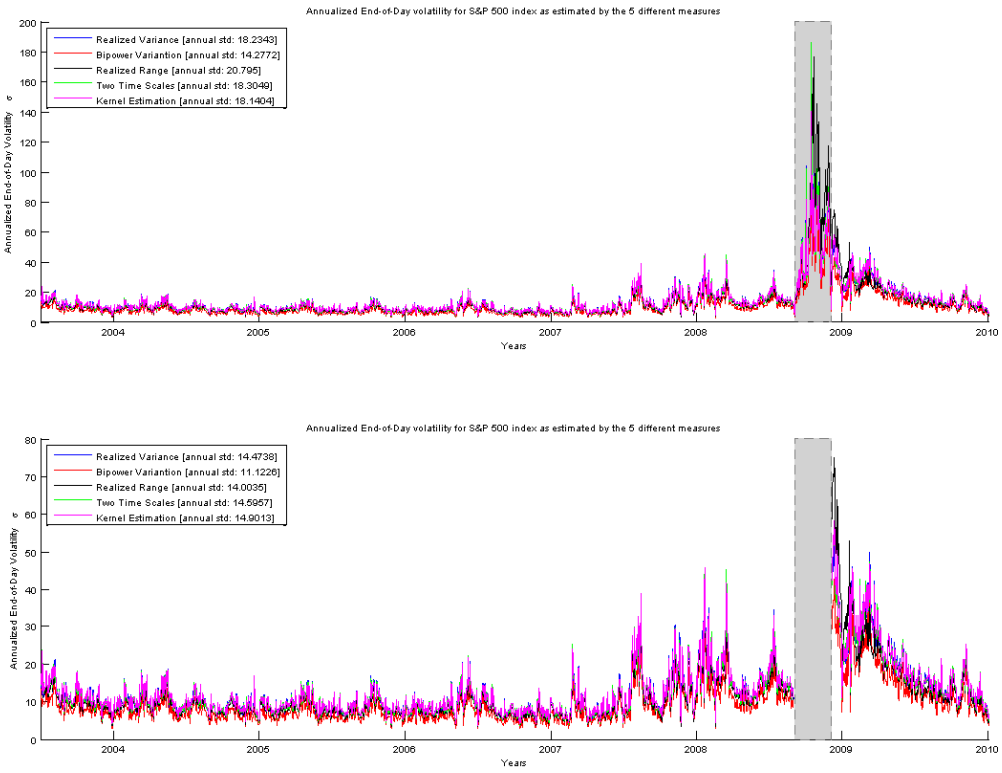

Figure C.1: Annualized standard deviation of S&P500 data by the 5 volatility measures. First including, second excluding Sept, Oct, Nov 2008 data. Note the scale on which graphs are plotted for a sense on the extremity of differences.

		S&P500 - BOD						S&P500 - MA					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2599^*$ (0,0050)	$15,0971^*$ (2,6750)	0,4471	0,0728	0,3742	0,8645	$0,2597^*$ (0,0050)	$0,6029^*$ (0,1039)	0,4436	0,0728	0,3708	3,7650
	BPV	$0,1659^*$ (0,0028)	$10,0676^*$ (2,5910)	0,4499	0,0697	0,3802	0,8695	$0,1645^*$ (0,0028)	$0,3828^*$ (0,0949)	0,4368	0,0697	0,3671	3,7793
	RR	$0,1095^*$ (0,0054)	$12,8351^*$ (0,4868)	0,6209	0,0692	0,5517	0,8681	$0,1089^*$ (0,0054)	$0,4839^*$ (0,0184)	0,6150	0,0692	0,5458	1,8574
	TTS	$0,1910^*$ (0,0066)	$12,1705^*$ (0,6047)	0,5843	0,0693	0,5150	0,8690	$0,1910^*$ (0,0066)	$0,4804^*$ (0,0240)	0,5778	0,0693	0,5086	2,3068
	Kernel	$0,1869^*$ (0,0069)	$11,1777^*$ (0,4965)	0,5762	0,0756	0,5006	0,8578	$0,1862^*$ (0,0069)	$0,4416^*$ (0,0198)	0,6054	0,0756	0,5298	2,8391
	Mean	0,1826	12,2696	0,5357	0,0713	0,4643	0,8658	0,1821	0,4783	0,5357	0,0713	0,4644	2,9095

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500 - BOD						S&P500 - MA					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1972^*$ (0,0045)	$6,2328^*$ (0,3316)	0,6270	0,1269	0,5001	0,7701	$0,1972^*$ (0,0045)	$0,4866^*$ (0,0261)	0,6344	0,1269	0,5075	1,6366
	BPV	$0,1159^*$ (0,0020)	$5,9500^*$ (0,3228)	0,5925	0,1308	0,4617	0,7634	$0,1156^*$ (0,0020)	$0,4388^*$ (0,0241)	0,6113	0,1308	0,4805	2,0244
	RR	$0,0787^*$ (0,0053)	$7,3887^*$ (0,2329)	0,7103	0,1262	0,5841	0,7672	$0,0807^*$ (0,0053)	$0,5580^*$ (0,0178)	0,7081	0,1262	0,5819	1,0492
	TTS	$0,1432^*$ (0,0064)	$6,6651^*$ (0,2470)	0,6621	0,1291	0,5330	0,7638	$0,1435^*$ (0,0064)	$0,5216^*$ (0,0194)	0,6610	0,1291	0,5319	1,2799
	Kernel	$0,1411^*$ (0,0076)	$6,6851^*$ (0,2455)	0,7027	0,1311	0,5716	0,7601	$0,1429^*$ (0,0075)	$0,5175^*$ (0,0192)	0,7030	0,1311	0,5718	1,3106
	Mean	0,1352	6,5843	0,6589	0,1288	0,5301	0,7649	0,1360	0,5045	0,6636	0,1288	0,5347	1,4601

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500 - BOD						S&P500 - MA					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1377^*$ (0,0062)	$3,4633^*$ (0,1343)	0,7536	0,2443	0,5092	0,5852	$0,1384^*$ (0,0061)	$0,5543^*$ (0,0216)	0,7443	0,2443	0,4999	0,8122
	BPV	$0,0855^*$ (0,0029)	$3,3970^*$ (0,1206)	0,7450	0,2400	0,5051	0,5914	$0,0858^*$ (0,0028)	$0,5186^*$ (0,0185)	0,7363	0,2400	0,4964	0,8870
	RR	$0,0575^*$ (0,0059)	$3,8437^*$ (0,1175)	0,7731	0,2495	0,5236	0,5721	$0,0624^*$ (0,0056)	$0,6026^*$ (0,0185)	0,7924	0,2495	0,5429	0,6509
	TTS	$0,1075^*$ (0,0078)	$3,6733^*$ (0,1336)	0,7857	0,2455	0,5403	0,5803	$0,1103^*$ (0,0075)	$0,5896^*$ (0,0213)	0,7991	0,2455	0,5536	0,6780
	Kernel	$0,0960^*$ (0,0081)	$3,7504^*$ (0,1220)	0,7681	0,2486	0,5195	0,5749	$0,1017^*$ (0,0078)	$0,5931^*$ (0,0194)	0,7905	0,2486	0,5419	0,6843
	Mean	0,0968	3,6256	0,7651	0,2456	0,5195	0,5808	0,0997	0,5717	0,7725	0,2456	0,5269	0,7425

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500 - BOD						S&P500 - MA					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0873^*$ (0,0075)	$2,3080^*$ (0,0781)	0,8644	0,3955	0,4690	0,3838	$0,0923^*$ (0,0071)	$0,7302^*$ (0,0247)	0,8692	0,3955	0,4737	0,2302
	BPV	$0,0503^*$ (0,0039)	$2,3733^*$ (0,0731)	0,8388	0,3876	0,4512	0,3936	$0,0535^*$ (0,0037)	$0,7262^*$ (0,0224)	0,8482	0,3876	0,4606	0,2519
	RR	$0,0409^*$ (0,0058)	$2,3973^*$ (0,0674)	0,8863	0,4059	0,4804	0,3665	$0,0499^*$ (0,0055)	$0,7568^*$ (0,0215)	0,8829	0,4059	0,4770	0,1891
	TTS	$0,0783^*$ (0,0080)	$2,3433^*$ (0,0778)	0,8752	0,3958	0,4794	0,3807	$0,0852^*$ (0,0076)	$0,7438^*$ (0,0246)	0,8851	0,3958	0,4893	0,1987
	Kernel	$0,0641^*$ (0,0080)	$2,4183^*$ (0,0693)	0,8908	0,3978	0,4930	0,3770	$0,0757^*$ (0,0075)	$0,7556^*$ (0,0218)	0,8995	0,3978	0,5017	0,1944
	Mean	0,0642	2,3680	0,8711	0,3965	0,4746	0,3803	0,0713	0,7425	0,8770	0,3965	0,4805	0,2129

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500 - BOD							S&P500 - MA						
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:10	$0,1642^*$ (0,0056)	$4,6235^*$ (0,2061)	0,6990	0,1790	0,5200	0,6855	09:30	10:10	$0,1661^*$ (0,0057)	$0,4805^*$ (0,0220)	0,7062	0,1790	0,5272	1,5233
	BPV	09:30	10:10	$0,0982^*$ (0,0030)	$4,7086^*$ (0,2049)	0,6949	0,1753	0,5196	0,6911	09:30	10:10	$0,0989^*$ (0,0030)	$0,4638^*$ (0,0206)	0,7108	0,1753	0,5355	1,6619
	RR	09:35	09:55	$0,0805^*$ (0,0048)	$11,4820^*$ (0,3486)	0,6929	0,1079	0,5850	0,8443	09:30	09:55	$0,0844^*$ (0,0053)	$0,5385^*$ (0,0175)	0,7011	0,1079	0,5932	1,2194
	TTS	09:30	10:10	$0,1218^*$ (0,0069)	$5,0187^*$ (0,1800)	0,7430	0,1786	0,5643	0,6830	09:30	10:25	$0,1139^*$ (0,0076)	$0,5707^*$ (0,0213)	0,7923	0,2302	0,5621	0,7827
	Kernel	09:30	09:50	$0,1637^*$ (0,0076)	$9,0314^*$ (0,3726)	0,6677	0,0957	0,5720	0,8217	09:30	09:55	$0,1488^*$ (0,0074)	$0,4976^*$ (0,0188)	0,6986	0,1134	0,5853	1,5838
	Mean	9:31	10:03	0,1257	6,9728	0,6995	0,1473	0,5522	0,7451	9:30	10:07	0,1224	0,5102	0,7218	0,1612	0,5607	1,3542

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500 - BOD							S&P500 - MA						
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	16:15	09:05	$0,7912^*$ (0,0025)	$-3,8403^*$ (0,4010)	0,6561	0,0000	0,6561	0,7569	16:15	09:05	$0,8079^*$ (0,0033)	$-1,3901^*$ (0,2152)	0,6785	0,0000	0,6785	2,9327
	BPV	04:15	08:30	$0,0706^*$ (0,0025)	$6,6582^*$ (0,2397)	0,6410	0,0000	0,6410	0,7397	21:55	08:30	$0,0551^*$ (0,0022)	$0,9885$ (0,0270)	0,6814	0,0000	0,6814	1,8063
	RR	17:10	08:35	$0,1117^*$ (0,0038)	$3,1440^*$ (0,1080)	0,6865	0,0000	0,6865	0,7290	16:15	09:20	$0,5445^*$ (0,0008)	$-1,2374^*$ (0,1256)	0,6820	0,0000	0,6820	2,1514
	TTS	16:15	09:30	$0,3518^*$ (0,0114)	$0,3962^*$ (0,2951)	0,6739	0,0000	0,6739	0,6470	16:15	09:30	$0,4091^*$ (0,0160)	$0,0692^*$ (0,1799)	0,6883	0,0000	0,6883	1,7176
	Kernel	16:15	09:05	$0,7630^*$ (0,0029)	$-3,0111^*$ (0,4563)	0,6831	0,0000	0,6831	0,7428	16:15	09:30	$0,7806^*$ (0,0038)	$-1,1906^*$ (0,2743)	0,6837	0,0000	0,6837	2,1769
	Mean	14:02	8:57	0,4177	0,6694	0,6681	0,0000	0,6681	0,7231	17:23	9:11	0,5194	-0,5521	0,6828	0,0000	0,6828	2,1570

(j) Statistics for the maximum based on R^2_{marg} oven 24 hour trade.

Table C.1: Overall Begin-of-the-day (BOD) and moving average (MA) seasonal forecasting statistics for S&P500 (omitting Sept, Oct, Nov 2008) futures taking H=252. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

		S&P500 - EWMA						S&P500 - FFF					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
15 minutes	RV	$0,2570^*$ (0,0050)	$0,6191^*$ (0,1085)	0,4483	0,0728	0,3754	3,0168	$0,2622^*$ (0,0057)	$0,7361^*$ (0,1071)	0,4473	0,0728	0,3745	0,9763
	BPV	$0,1612^*$ (0,0027)	$0,4026^*$ (0,0921)	0,4465	0,0697	0,3768	3,1579	$0,1619^*$ (0,0031)	$0,5955^*$ (0,0980)	0,4679	0,0697	0,3982	0,8130
	RR	$0,1023^*$ (0,0050)	$0,5087^*$ (0,0178)	0,6599	0,0692	0,5907	1,4600	$0,0803^*$ (0,0062)	$0,7647^*$ (0,0262)	0,6177	0,0692	0,5485	0,5251
	TTS	$0,1824^*$ (0,0062)	$0,5095^*$ (0,0235)	0,5907	0,0693	0,5215	1,8367	$0,1639^*$ (0,0077)	$0,7091^*$ (0,0310)	0,5878	0,0693	0,5186	0,7132
	Kernel	$0,1789^*$ (0,0065)	$0,4644^*$ (0,0194)	0,6535	0,0756	0,5780	2,2012	$0,1635^*$ (0,0080)	$0,6358^*$ (0,0262)	0,6002	0,0756	0,5247	0,8487
	Mean	0,1764	0,5009	0,5598	0,0713	0,4885	2,3345	0,1664	0,6882	0,5442	0,0713	0,4729	0,7753

(a) Statistics for volatility measured from start of the day up until 15 min.

		S&P500 - EWMA						S&P500 - FFF					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
30 minutes	RV	$0,1979^*$ (0,0047)	$0,4919^*$ (0,0277)	0,6611	0,1269	0,5342	1,3224	$0,1860^*$ (0,0055)	$0,6563^*$ (0,0314)	0,6248	0,1269	0,4979	0,5630
	BPV	$0,1132^*$ (0,0020)	$0,4569^*$ (0,0246)	0,6346	0,1308	0,5038	1,6768	$0,1086^*$ (0,0025)	$0,6433^*$ (0,0296)	0,5960	0,1308	0,4652	0,6103
	RR	$0,0783^*$ (0,0050)	$0,5763^*$ (0,0174)	0,7351	0,1262	0,6089	0,8083	$0,0539^*$ (0,0059)	$0,8078^*$ (0,0235)	0,6942	0,1262	0,5680	0,3610
	TTS	$0,1371^*$ (0,0060)	$0,5470^*$ (0,0192)	0,7345	0,1291	0,6054	0,9919	$0,1159^*$ (0,0074)	$0,7287^*$ (0,0250)	0,6781	0,1291	0,5490	0,4906
	Kernel	$0,1394^*$ (0,0071)	$0,5367^*$ (0,0189)	0,7596	0,1311	0,6284	1,0051	$0,1131^*$ (0,0084)	$0,7247^*$ (0,0247)	0,7019	0,1311	0,5707	0,4723
	Mean	0,1332	0,5218	0,7050	0,1288	0,5761	1,1609	0,1155	0,7121	0,6590	0,1288	0,5302	0,4994

(b) Statistics for volatility measured from start of the day up until 30 min.

		S&P500 - EWMA						S&P500 - FFF					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
60 minutes	RV	$0,1327^*$ (0,0058)	$0,5790^*$ (0,0213)	0,7723	0,2443	0,5280	0,6555	$0,1205^*$ (0,0072)	$0,7206^*$ (0,0268)	0,7543	0,2443	0,5099	0,3633
	BPV	$0,0833^*$ (0,0028)	$0,5407^*$ (0,0189)	0,7612	0,2400	0,5212	0,7406	$0,0739^*$ (0,0035)	$0,7219^*$ (0,0244)	0,7359	0,2400	0,4960	0,3429
	RR	$0,0616^*$ (0,0052)	$0,6189^*$ (0,0179)	0,8259	0,2495	0,5764	0,5012	$0,0445^*$ (0,0063)	$0,7833^*$ (0,0235)	0,8033	0,2495	0,5538	0,2783
	TTS	$0,1063^*$ (0,0070)	$0,6131^*$ (0,0209)	0,8176	0,2455	0,5721	0,5264	$0,0910^*$ (0,0085)	$0,7513^*$ (0,0267)	0,7961	0,2455	0,5506	0,3150
	Kernel	$0,0984^*$ (0,0072)	$0,6140^*$ (0,0187)	0,8075	0,2486	0,5589	0,5234	$0,0791^*$ (0,0088)	$0,7632^*$ (0,0244)	0,7808	0,2486	0,5322	0,3039
	Mean	0,0964	0,5932	0,7969	0,2456	0,5513	0,5894	0,0818	0,7481	0,7741	0,2456	0,5285	0,3207

(c) Statistics for volatility measured from start of the day up until 60 min.

		S&P500 - EWMA						S&P500 - FFF					
		$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
120 minutes	RV	$0,0916^*$ (0,0066)	$0,7482^*$ (0,0239)	0,8861	0,3955	0,4907	0,1904	$0,0757^*$ (0,0079)	$0,8299^*$ (0,0276)	0,8739	0,3955	0,4784	0,1587
	BPV	$0,0557^*$ (0,0036)	$0,7328^*$ (0,0228)	0,8550	0,3876	0,4673	0,2134	$0,0421^*$ (0,0042)	$0,8617^*$ (0,0260)	0,8499	0,3876	0,4623	0,1555
	RR	$0,0513^*$ (0,0051)	$0,7695^*$ (0,0209)	0,9078	0,4059	0,5019	0,1502	$0,0347^*$ (0,0061)	$0,8480^*$ (0,0239)	0,8889	0,4059	0,4829	0,1347
	TTS	$0,0857^*$ (0,0071)	$0,7613^*$ (0,0241)	0,8995	0,3958	0,5037	0,1604	$0,0693^*$ (0,0085)	$0,8325^*$ (0,0275)	0,8897	0,3958	0,4939	0,1410
	Kernel	$0,0769^*$ (0,0070)	$0,7707^*$ (0,0212)	0,9006	0,3978	0,5027	0,1548	$0,0548^*$ (0,0084)	$0,8566^*$ (0,0245)	0,8895	0,3978	0,4916	0,1332
	Mean	0,0722	0,7565	0,8898	0,3965	0,4933	0,1738	0,0553	0,8458	0,8784	0,3965	0,4818	0,1446

(d) Statistics for volatility measured from start of the day up until 120 min.

		S&P500 - EWMA								S&P500 - FFF							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
Open Outcry maximum	RV	09:30	10:05	$0,1721^*$ (0,0052)	$0,4831^*$ (0,0227)	0,7303	0,1602	0,5701	1,4418	09:30	10:15	$0,1378^*$ (0,0072)	$0,6898^*$ (0,0283)	0,7263	0,1976	0,5287	0,4936
	BPV	09:30	10:10	$0,0940^*$ (0,0029)	$0,4918^*$ (0,0206)	0,7103	0,1753	0,5350	1,3797	09:30	10:10	$0,0861^*$ (0,0036)	$0,7058^*$ (0,0277)	0,6949	0,1753	0,5195	0,5015
	RR	09:35	09:55	$0,0825^*$ (0,0046)	$0,5924^*$ (0,0177)	0,7353	0,1079	0,6274	0,8479	09:40	10:05	$0,0540^*$ (0,0047)	$0,7860^*$ (0,0219)	0,7477	0,1595	0,5882	0,4733
	TTS	09:30	10:00	$0,1371^*$ (0,0060)	$0,5470^*$ (0,0192)	0,7345	0,1291	0,6054	0,9919	09:30	09:55	$0,1300^*$ (0,0073)	$0,7325^*$ (0,0265)	0,6754	0,1071	0,5682	0,5158
	Kernel	09:30	10:00	$0,1394^*$ (0,0071)	$0,5367^*$ (0,0189)	0,7596	0,1311	0,6284	1,0051	09:30	09:55	$0,1192^*$ (0,0083)	$0,7074^*$ (0,0245)	0,6950	0,1134	0,5817	0,5438
	Mean	9:31	10:02	0,1250	0,5302	0,7340	0,1407	0,5933	1,1333	9:32	10:04	0,1054	0,7243	0,7079	0,1506	0,5573	0,5056

(i) Statistics for the maximum based on R^2_{marg} over open outcry trade.

		S&P500 - EWMA								S&P500 - FFF							
		start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
24 hour maximum	RV	16:15	09:30	$0,7995^*$ (0,0030)	$-1,7073^*$ (0,2166)	0,7030	0,0000	0,7030	2,1892								
	BPV	16:15	09:20	$0,0692^*$ (0,0030)	$1,0234$ (0,0321)	0,6955	0,0000	0,6955	1,8518								
	RR	16:30	09:30	$0,5424^*$ (0,0008)	$-1,4512^*$ (0,1193)	0,7260	0,0000	0,7260	1,8322	**Infeasible due to seasonal**							
	TTS	16:15	09:00	$0,9459^*$ (0,0359)	$-2,2840$ (7,8013)	0,7159	0,0000	0,7159	2,3531								
	Kernel	16:30	09:15	$0,7751^*$ (0,0036)	$-1,1818^*$ (0,2212)	0,7143	0,0000	0,7143	2,4973								
	Mean	16:21	9:19	0,6264	-1,1202	0,7109	0,0000	0,7109	2,1447								

(j) Statistics for the maximum based on R^2_{marg} oven 24 hour trade.

Table C.2: Overall Exponentially Weighted Moving Average (EWMA) and Fourier Flexible Form (FFF) Seasonal variance forecasting statistics for S&P500 futures taking H=252. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

	S&P500 - MZS - 15 min						S&P500 - MZS - 30 min					
	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	$0,0528^*$ (0,0130)	$0,9011^*$ (0,0360)	0,6462	0,0731	0,5731	0,9136	$0,0454^*$ (0,0114)	$0,9061^*$ (0,0321)	0,7042	0,1273	0,5770	0,7783
BPV	$0,0244^*$ (0,0071)	$0,9233^*$ (0,0338)	0,6236	0,0699	0,5537	1,3280	$0,0195^*$ (0,0063)	$0,9358^*$ (0,0306)	0,6801	0,1312	0,5489	0,9264
RR	$0,0379^*$ (0,0080)	$0,8974^*$ (0,0291)	0,6819	0,0694	0,6125	0,6336	$0,0318^*$ (0,0077)	$0,9076^*$ (0,0285)	0,7142	0,1265	0,5877	0,5708
TTS	$0,0470^*$ (0,0118)	$0,9054^*$ (0,0321)	0,6705	0,0695	0,6009	0,6129	$0,0470^*$ (0,0113)	$0,8947^*$ (0,0315)	0,7493	0,1296	0,6197	0,5414
Kernel	$0,0497^*$ (0,0121)	$0,8987^*$ (0,0311)	0,6452	0,0758	0,5694	0,5935	$0,0563^*$ (0,0117)	$0,8805^*$ (0,0307)	0,7058	0,1315	0,5743	0,5344
Mean	0,0424	0,9052	0,6535	0,0715	0,5819	0,8163	0,0400	0,9049	0,7107	0,1292	0,5815	0,6703

(a) Statistics for volatility measured from start of the day up until 15 min and 30 min.

	S&P500 - MZS - 60 min						S&P500 - MZS - 120 min					
	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	$0,0287^*$ (0,0091)	$0,9357^*$ (0,0270)	0,7848	0,2446	0,5402	0,4290	$0,0293^*$ (0,0079)	$0,9263^*$ (0,0244)	0,8719	0,3961	0,4758	0,2763
BPV	$0,0170^*$ (0,0054)	$0,9337^*$ (0,0275)	0,7612	0,2403	0,5209	0,5655	$0,0167^*$ (0,0047)	$0,9299^*$ (0,0247)	0,8622	0,3884	0,4738	0,3794
RR	$0,0594^*$ (0,0097)	$0,8255^*$ (0,0376)	0,8138	0,2498	0,5641	0,3506	$0,0390^*$ (0,0071)	$0,8809^*$ (0,0285)	0,8800	0,4065	0,4735	0,2440
TTS	$0,0401^*$ (0,0102)	$0,9110^*$ (0,0298)	0,7920	0,2458	0,5462	0,3413	$0,0635^*$ (0,0108)	$0,8563^*$ (0,0331)	0,8719	0,3965	0,4755	0,2404
Kernel	$0,0744^*$ (0,0131)	$0,8424^*$ (0,0362)	0,7926	0,2488	0,5438	0,3401	$0,0571^*$ (0,0105)	$0,8732^*$ (0,0300)	0,8876	0,3984	0,4892	0,2377
Mean	0,0439	0,8897	0,7889	0,2459	0,5430	0,4053	0,0411	0,8933	0,8747	0,3972	0,4776	0,2756

(b) Statistics for volatility measured from start of the day up until 60 min and 120 min.

	S&P500 - MZS - Open Outcry maximum								S&P500 - MZS - 24 h maximum							
	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	start	stop	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	09:30	09:40	$0,5289^*$ (0,0291)	$-0,1557^*$ (0,0833)	0,6687	0,0519	0,6168	1,1440	18:40	08:45	$0,1181^*$ (0,0216)	$0,7847^*$ (0,0608)	0,7024	0,0000	0,7024	1,1580
BPV	09:30	09:35	$0,3331^*$ (0,0197)	$-0,2268^*$ (0,0961)	0,6286	0,0275	0,6010	1,8479	16:15	09:05	$0,0462^*$ (0,0093)	$0,8487^*$ (0,0446)	0,6661	0,0000	0,6661	2,1285
RR	09:30	09:40	$0,0615^*$ (0,0101)	$0,8298^*$ (0,0368)	0,6609	0,0478	0,6130	0,7954	22:55	09:15	$0,0385^*$ (0,0092)	$0,9060^*$ (0,0338)	0,7154	0,0000	0,7154	0,9420
TTS	09:30	10:00	$0,0470^*$ (0,0113)	$0,8947^*$ (0,0315)	0,7493	0,1296	0,6197	0,5414	23:40	09:30	$0,0505^*$ (0,0124)	$0,9084^*$ (0,0344)	0,6946	0,0000	0,6946	0,8404
Kernel	09:35	09:55	$0,0346^*$ (0,0101)	$0,9282^*$ (0,0265)	0,7183	0,1136	0,6047	0,5886	20:10	09:15	$0,0956^*$ (0,0202)	$0,8352^*$ (0,0529)	0,7056	0,0000	0,7056	0,9258
Mean	9:31	9:46	0,2010	0,4540	0,6852	0,0741	0,6110	0,9835	20:20	9:10	0,0698	0,8566	0,6968	0,0000	0,6968	1,1989

(e) Statistics for the maximum based on R^2_{marg} over open outcry trade.

Table C.3: Overall Mincer Zarnowits scaled (MZS) variance forecasting statistics for S&P500 futures, taking a rolling window of length 20 days. Standard errors of the estimates are given between parenthesis and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

	Random Walk						GARCH(1,1)					
	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE	$\hat{\alpha}$	$\hat{\beta}$	R^2_{MAD}	VR	R^2_{marg}	HMSPE
RV	$0,1254^*$ (0,0119)	$0,8378^*$ (0,0442)	0,6439	0,0000	0,6439	0,8659	$-0,1570^*$ (0,0249)	$1,2822^*$ (0,0514)	0,5315	0,0000	0,5315	1,2399
BPV	$0,0822^*$ (0,0080)	$0,8091^*$ (0,0509)	0,6472	0,0000	0,6472	1,2637	$-0,0898^*$ (0,0150)	$0,7524^*$ (0,0310)	0,5247	0,0000	0,5247	5,7043
RR	$0,0905^*$ (0,0086)	$0,8263^*$ (0,0399)	0,7091	0,0000	0,7091	1,1653	$-0,1053^*$ (0,0176)	$0,9347$ (0,0364)	0,5263	0,0000	0,5263	2,3114
TTS	$0,1093^*$ (0,0111)	$0,8672^*$ (0,0396)	0,6655	0,0000	0,6655	0,6869	$-0,1589^*$ (0,0248)	$1,2969^*$ (0,0511)	0,5276	0,0000	0,5276	0,9109
Kernel	$0,1162^*$ (0,0120)	$0,8557^*$ (0,0394)	0,7085	0,0000	0,7085	0,9132	$-0,1649^*$ (0,0249)	$1,3531^*$ (0,0515)	0,5204	0,0000	0,5204	0,7670
mean	0,1047	0,8392	0,6748	0,0000	0,6748	0,9790	-0,1352	1,1239	0,5261	0,0000	0,5261	2,1867

Table C.4: Random Walk and GARCH(1,1) variance forecasting statistics for the five volatility measures. Standard errors of the estimates are given between brackets and a star is appointed to estimates of α and β significantly different from respectively 0 and 1 on a 95% confidence level.

Wesley Ballering

Tweede Keucheniusstraat 17-3

1051 VP Amsterdam

tel: +31 (0) 6 81494645

mail: wesleyballering@gmail.com