

Laakkonen, Helinä

Working Paper

Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method

Bank of Finland Research Discussion Papers, No. 23/2007

Provided in Cooperation with:

Bank of Finland, Helsinki

Suggested Citation: Laakkonen, Helinä (2007) : Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method, Bank of Finland Research Discussion Papers, No. 23/2007, ISBN 978-952-462-395-7, Bank of Finland, Helsinki, <https://nbn-resolving.de/urn:nbn:fi:bof-20140807284>

This Version is available at:

<https://hdl.handle.net/10419/212081>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Helinä Laakkonen

Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method



EUROJÄRJESTELMÄ
EUROSYSTEMET

Bank of Finland Research
Discussion Papers
23 • 2007

**Suomen Pankki
Bank of Finland
PO Box 160
FI-00101 HELSINKI
Finland
☎ + 358 10 8311**

<http://www.bof.fi>



Helinä Laakkonen*

Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method

The views expressed are those of the author and do not necessarily reflect the views of the Bank of Finland.

* Finnish Doctoral Programme in Economics and University of Jyväskylä, School of Business and Economics, e-mail: helina.laakkonen@cc.jyu.fi.

I would like to thank Markku Lanne, Timo Teräsvirta, Giampiero M. Gallo, Heikki Kauppi, Stefano Schiavo and Katja Ahoniemi for their constructive comments. The suggestions of participants in EEA-ESEM 2007, INFINITI 2007, the EFMA 2007 ‘Merton H. Miller’ Doctoral Seminar, the Annual Meeting of the Finnish Society for Economic Research 2007, the Workshop on Quantitative Finance 2007 in Venice, Volatility workshop 2007 in Stockholm, the FDPE Workshop on Econometrics and Computational Economics, and the Bank of Finland research workshop in Helsinki are also highly appreciated. In addition, I am grateful to Jouko Vilmunen for the opportunity to collect the announcements data during my internship at the Bank of Finland. Finally, I would like to thank the Yrjö Jahnesson Foundation and the Finnish Foundation for Advancement of Securities Markets for their financial support.

This paper can be downloaded without charge from <http://www.bof.fi> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1032876.

<http://www.bof.fi>

ISBN 978-952-462-394-0
ISSN 0785-3572
(print)

ISBN 978-952-462-395-7
ISSN 1456-6184
(online)

Helsinki 2007

Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method

Bank of Finland Research
Discussion Papers 23/2007

Helinä Laakkonen
Monetary Policy and Research Department

Abstract

Filtering intraday seasonality in volatility is crucial for using high frequency data in econometric analysis. This paper studies the effects of filtering on statistical inference concerning the impact of news on exchange rate volatility. The properties of different methods are studied using a 5-minute frequency USD/EUR data set and simulated returns. The simulation results suggest that all the methods tend to produce downward-biased estimates of news coefficients, some more than others. The study supports the Flexible Fourier Form method as the best for seasonality filtering.

Keywords: high-frequency, volatility, macro announcements, seasonality

JEL classification numbers: C22, C49, C52, E44

Valuuttamarkkinoiden volatiliteetti, makrotalouden uutiset ja päivänsisäisen kausivaihtelun puhdistusmenetelmän valinta

Suomen Pankin keskustelualoitteita 23/2007

Helinä Laakkonen
Rahapolitiikka- ja tutkimusosasto

Tiivistelmä

Käytettäessä tiheäfrekvenssisia tuottoaineistoja on tulosten luotettavuuden vuoksi tärkeää puhdistaa volatiliteetin päivänsisäinen kausivaihtelu ennen varsinaista analyysia. Tutkimuksessa tarkastellaan puhdistusmenetelmän valinnan vaikutuksia tuloksiin, kun tutkitaan makrotalouden uutisten vaikutusta valuuttakurssien volatiliteettiin. Eri menetelmien ominaisuuksia tutkitaan käyttämällä 5 minuutin dollari/euro-tuottoaineistoa ja simuloituja tuottoja. Simulointitulosten mukaan uutismuuttujien vaikutukset estimoituvat puhdistuksen vuoksi aina liian pieniksi. Estimointiharhan suuruus riippuu käytetystä menetelmästä. Joustavien Fourier-muotojen menetelmä soveltuu lisäksi tulosten mukaan parhaiten päivänsisäisen kausivaihtelun puhdistamiseen.

Avainsanat: päivänsisäinen, volatiliteetti, makrouutiset, kausivaihtelu

JEL-luokittelu: C22, C49, C52, E44

Contents

Abstract.....	3
Tiivistelmä (abstract in Finnish).....	4
1 Introduction.....	7
2 Intraday seasonality.....	9
2.1 Data.....	9
2.2 Seasonality filtering.....	10
3 Seasonality after filtering.....	11
3.1 Flexible Fourier Form method.....	11
3.2 Locally Weighted Scatterplot Smoothing method.....	14
3.3 The Intraday Average Observation Model.....	15
3.4 The impact of news on volatility.....	16
4 Properties of the filtered returns.....	21
5 Simulation study.....	22
5.1 Returns.....	23
5.2 Properties of the simulated returns.....	25
5.3 Simulation results.....	26
6 Conclusions.....	28
References.....	29

1 Introduction

High frequency data sets have been used extensively since the development of electronic trading systems in the early 1990s. One area of research in which intraday data sets have been widely used concerns the impact of news on exchange rate volatility.¹ Since markets react to new information very quickly, intraday data sets are more informative and provide more accurate results than do daily data sets. Nevertheless, these kinds of data sets have some special features that need to be scrutinized for the sake of reliable results. One crucial issue is the filtering of intraday seasonality of volatility, which is caused by differences in trading times in the global foreign exchange markets. This seasonality causes periodical U-shape patterns in autocorrelation functions of volatility, and therefore it has to be filtered out of intraday data used for statistical inference.²

Usually seasonality is considered as nuisance in time series econometrics, and filtering is something one needs to do before starting the actual work with the data. However, filtering might cause at least three kinds of problems that researchers rarely consider. First, most filtering methods are capable of removing some of the seasonality, but not all of it. The periodicity in volatility that remains after filtering may affect the other results of the study. Second, the filtering might generate something new, for example extra jumps, in returns. Third, in filtering out seasonality, something important might be filtered out as well. For example, if one is studying the impact of news on volatility, the filter should not eliminate the news effects.

To our knowledge, filtering methods have been compared in only two papers so far: Martens et al (2001) conclude that explicitly modeling the intraday volatility component improves out-of-sample forecasts. However, they compare only the two step Flexible Fourier Form model (FFF) of Andersen and Bollerslev (1997) to the computationally costly periodical GARCH model, which includes the FFF. Ben Omrane and Bodt (2005) compare the Intradaily Average Observation method (IAOM) of Bauwens et al (2005) and a kernel smoothing method to a method that uses self-organizing maps (SOM) to capture seasonality. The authors conclude that while the IAOM and kernel smoothing methods are capable of estimating the deterministic part of seasonality, they both fail to capture the stochastic part. The new method that they propose, SOM, seems to capture also stochastic cycles.

Compared to those two papers, in this paper we concentrate more carefully on the possible consequences that filtering might have. We consider three commonly used methods, each of which belongs to one of three filtering method categories, set out by Ben Omrane and Bodt (2005). These are the Flexible Fourier Form method (FFF) of Andersen and Bollerslev (1997), which uses sinusoids as exogenous variables to capture the periodicity; the Locally Weighted Scatterplot Smoothing method (LOWESS) of Cleveland (1979), which adopts a weighted time trend estimation for deseasonalizing; and the

¹DeGennaro and Schrieves (1997), Andersen et al (2003), Bauwens et al (2005), Dominguez and Panthaki (2006), Laakkonen (2007) among others.

²Evidence of the intraday volatility pattern has been shown eg by Andersen and Bollerslev (1998), Melvin and Yin (2000), Cai et al (2001), and Bauwens et al (2005).

Intradaily Average Observations Model (IAOM) of Bauwens et al (2005) in which the intraday volatility estimate is computed by averaging the squared returns (per each intraday interval and separately per each weekday) over the whole sample period.

We use high-frequency returns to study the first two problems that filtering might cause: 1) we compare autocorrelation functions of different filtered absolute returns to see if there is periodicity left in volatility after filtering and study the news effects with different filtered returns to see whether the filtering affects the magnitudes of the news coefficients, 2) we compare the key statistical figures of the original and filtered returns to study whether filtering adds some new elements to the return process.

Our data contain six years of five minute United States dollar against euro (USD/EUR) quotes from 1 Jan. 1999 to 31 Dec. 2004. This data set is longer than those used in most of the earlier literature. As the filtering methods usually perform better when the seasonality is regular (Ben Omrane and Bodt, 2005), the six year data set complicates the deseasonalizing, since it is very likely that the intraday pattern has developed over the years, especially since the data also cover the early years of the euro. The macro announcements are obtained from the Bloomberg World Economic Calendar (WECO) and contain all macro news from the USA, Germany and Japan.

Furthermore, we use simulated returns to study the third problem that filtering might cause: 3) we use the properties of the USD/EUR data to simulate returns, which comprise the daily volatility component, intraday volatility component, news component and random shocks. We deseasonalize the simulated returns, estimate the news effects with the filtered returns and compare the estimated news coefficient values to the simulated ones to see if the estimated news effects are biased.

The results suggest that the methods have some differences, but all the methods are capable of filtering the periodicity in volatility only if the estimation is done by dividing data into sufficiently short subsets. This indicates that the intraday volatility seasonality is time varying. The filtering method and the selected subset length also affect the magnitudes of the news coefficients: when the subset is shortened, the news coefficients increase when the filtering is done by using the FFF method, and decrease, when the LOWESS or IAOM methods are used for filtering.

According to the results of the simulation study, if the returns are not filtered at all, the US and European news coefficients are too large and the Asian news coefficients too small. Also, all the methods tend to produce downward biased estimates for the news coefficients, ie to filter out part of the news effects, some worse than others. While the LOWESS is performing the best in terms of getting rid of periodicity in volatility, it also seems to filter out more news effects than do the other two filters. The IAOM performs almost as well as the FFF, but in the case of regularly announced news it produces very downward biased estimates compared to the FFF method. Therefore, the study supports the FFF model as the best for seasonality filtering.

The paper is divided into six sections. Section 2 presents the properties of USD/EUR data and summarizes briefly the filtering methods used in the literature. The three problems that filtering might cause are studied in sections 3, 4 and 5. Section 6 concludes the study.

2 Intraday seasonality

In this section we present the properties of the USD/EUR data, summarize the filtering methods used in the literature and present the general idea behind in all of the methods.

2.1 Data

The original data contain quotes at 5-minute intervals on the USD/EUR exchange rate during 1 Jan. 1999 – 31 Dec. 2004 and is obtained from Olsen and Associates.³ To study the intraday seasonality of volatility, we use absolute returns as a measure of volatility⁴ and expose the average intraday volatility pattern by computing the mean absolute return per every five-minute interval in 24 hours (Figure 2.1a). The level of volatility during a day depends on the trading times in different markets: the Far East markets open around 23:00 GMT and cause a small increase in volatility; the European markets open around 7:00 GMT and volatility increases more; and the US markets open around 14:00 GMT after which volatility reaches its highest level. It is noteworthy that there are two spikes in the average volatility pattern: most US macro figures are announced at 13:30 GMT and 15:00 GMT and, as can be seen, they cause significant increases in volatility.

When this pattern is repeated every day, it causes a U-shape pattern in the autocorrelation of volatility. Figure 2.1b presents the autocorrelation coefficients of absolute returns for 1500 five-minute lags, ie the autocorrelogram for five days. As can be seen, the U-shape pattern is repeated almost identically every day. Therefore, before using these kind of data sets the returns have to be deseasonalized.

³Weekends and certain holidays were excluded, and daylight saving times in the USA and Europe were taken into account, as is standard in the literature.

⁴Absolute returns have been widely used as a volatility measure in the literature. A literature review of the use of absolute returns is provided by Granger and Sin (1999).

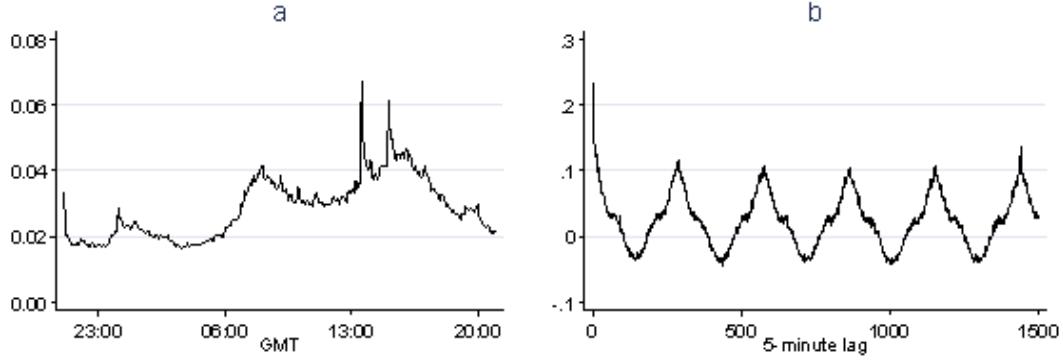


Figure 2.1 Intraday seasonality of volatility

Figure 2.1a graphs the average USD/EUR volatility during 24 hours (mean absolute return over each five minute interval). Volatility pattern is caused by different trading times in the global foreign exchange markets. Figure 2.1b graphs the five day autocorrelogram of the five minute USD/EUR absolute returns.

2.2 Seasonality filtering

Even though high-frequency data sets have been used for a more than a decade now, research in seasonality filtering is still very active. Many different methods have been proposed, none of which has become standard in the literature. Ben Omrane and Bodt (2005) provide a good review and introduce a taxonomy of filtering methods used in the literature.

The first category of filtering methods uses a linear projection to model the seasonality component: Andersen and Bollerslev (1997) use sinusoids as exogenous variables, while Degennaro and Schrieves (1997) include hourly dummies in the volatility regression to capture the seasonality. The second category uses smoothing techniques: Cleveland (1979) adopts a weighted time trend estimation, Engle and Russell (1998) use a cubic splines technique, whereas Veredas et al (2002) and Ben Omrane and Bodt (2005) employ kernel estimators for deseasonalizing. The third category uses the average (absolute/squared) returns per intraday interval to compute the filter: Dacorogna et al (1993) introduced the ϑ -time transformation, which deseasonalizes volatility by expanding periods with high average market activity and contracting periods with low average market activity; Melvin and Yin (2000) divide each observation by the mean number of quotes for each hour of the business week during a subset; while Bauwens et al (2005) compute the intraday volatility estimate by averaging the squared returns per 5-minute interval over the whole sample period (separately for the different weekdays).

We consider three commonly used methods, each of which belongs to one of the above categories (linear projection, smoothing and average observation). We tried to consider methods that are widely used in the news literature.

However, we think the main differences between the filters are between the three categories rather than within any of the categories.

The general idea of seasonality filtering in all of the methods is as follows: the model produces an estimate for the intraday volatility, denoted $s_{t,n}$. This estimate $\hat{s}_{t,n}$ is normalized so that the mean of the normalized seasonality estimate equals one⁵

$$\tilde{s}_{t,n} = \frac{T \cdot \hat{s}_{t,n}}{\sum_{t=1}^{\lfloor T/N \rfloor} \sum_{n=1}^N \hat{s}_{t,n}} \quad (2.1)$$

where T is the total number of observations and N is the number of observations during one day (288 for 5-minute intervals in the 24 hour market). The original returns $R_{t,n}$ are then divided by the normalized estimate $\tilde{s}_{t,n}$ to get the filtered returns $\tilde{R}_{t,n} = \frac{R_{t,n}}{\tilde{s}_{t,n}}$. Since the mean of $\tilde{s}_{t,n}$ equals one (and $\tilde{s}_{t,n}$ is always positive), the consequences of filtering are that volatility is increased in the low-volatility periods and decreased in the high-volatility periods. Other than that, the returns should remain the same.

3 Seasonality after filtering

In this section we study the ability of different methods to filter the seasonality in intraday volatility. As is common in the literature that compares different filtering methods (Martens et al, 2002, Ben Omrane and Bodt, 2005), we study graphically whether there is periodicity left in the filtered absolute returns. The filter performs the better, the less the periodicity left in the volatility. In addition, we examine whether the remaining seasonality affects the results of the further study with filtered returns. For this, we look at the impact of news on the volatility of different filtered returns. The results apply only for the used data set, which is very representative of the data sets used in the literature, however.

3.1 Flexible Fourier Form method

Of the linear projection techniques, we consider the Flexible Fourier Form method (FFF) introduced in this context by Andersen and Bollerslev (1997).⁶ Andersen and Bollerslev (1997) state that, since the variability during a day is so systematic, the intraday dynamics of absolute returns can be estimated by using different frequencies of sine and cosine functions. The volatility of the return process $R_{t,n}$ is divided into three components: the daily volatility component (σ_t divided by $N^{1/2}$ where N is the number of observations in one

⁵In the equations, n is an index for 5-minute intraday interval and t for day.

⁶The FFF method is one of the most widely used filtering methods in the news literature. It has been used for example in the studies of Cai et al (2001), Andersen et al (2003), Dominquez and Panthaki (2006) and Laakkonen (2007).

day, ie 288 for 5-minute intervals in the 24 hour market), the intraday volatility component $s_{t,n}$, and the random error term $Z_{t,n}$

$$R_{t,n} = E(R_{t,n}) + \frac{\sigma_t}{N^{1/2}} s_{t,n} Z_{t,n} \quad (3.1)$$

The expected return $E(R_{t,n})$ is then replaced by the mean return \bar{R} , and the absolute demeaned returns are used as the measure of volatility $|R_{t,n} - \bar{R}|$. The daily volatility component is eliminated by dividing the volatility measure by $\frac{\hat{\sigma}_t}{N^{1/2}}$, where $\hat{\sigma}_t$ is the GARCH(1,1) estimate of daily volatility. After elimination of the daily component, squaring and taking logs, equation 3.1 becomes

$$2 \ln \frac{|R_{t,n} - \bar{R}|}{\hat{\sigma}_t/N^{1/2}} = 2 \ln(s_{t,n}) + 2 \ln(Z_{t,n}) \quad (3.2)$$

There are two components left on the right-hand side of equation 3.2. The first is the component for the intraday volatility, which can be modeled using the trigonometric functions; and the other component is the error term, which includes the extra volatility in the markets, such as the volatility caused by new information. Equation 3.3

$$\begin{aligned} f_{t,n} = & c + \delta_1 \frac{n}{N_1} + \delta_2 \frac{n^2}{N_2} + \sum_{k=1}^D \lambda_k I_k(t, n) \\ & + \sum_{p=1}^P \left(\delta_{c,p} \cos \left(\frac{p2\pi}{N} n \right) + \delta_{s,p} \sin \left(\frac{p2\pi}{N} n \right) \right) + \varepsilon_{t,n} \end{aligned} \quad (3.3)$$

where $f_{t,n} = 2 \ln \frac{|R_{t,n} - \bar{R}|}{\hat{\sigma}_t/N^{1/2}}$, presents the Flexible Fourier Form model. Besides the sinusoids,⁷ the model contains the intercept c , the error term $\varepsilon_{t,n}$ and the normalizing factors $\frac{n}{N_1}$ and $\frac{n^2}{N_2}$, where $N_1 = (N+1)/2$ and $N_2 = (N+1)(N+2)/6$. The model also contains the indicator variables $I_k(t, n)$, which are used to control for holiday effects, weekday effects, Monday effects etc. The estimate for intraday volatility $\hat{s}_{t,n}$ is then obtained as $\hat{s}_{t,n} = \exp(\hat{f}_{t,n}/2)$, where $\hat{f}_{t,n}$ are the fitted values of the FFF model.

If the FFF model is estimated once by using the whole data, it is assumed that the intraday pattern remains the same every day. Unfortunately, this might not be the case. For example, the trading hours of European markets caused much higher volatility in the early years of the euro than they do today. Therefore, besides estimating the FFF by using the whole data set we estimated the model by using subsets of different length. The subset FFF model was re-estimated a) yearly b) quarterly and c) weekly. Figure 3.1 presents the autocorrelation coefficients of the filtered absolute returns

⁷The value $P = 9$ was chosen for the USD/EUR data according to AIC and BIC.

compared to the raw absolute returns, when seasonality is filtered out by the FFF model.

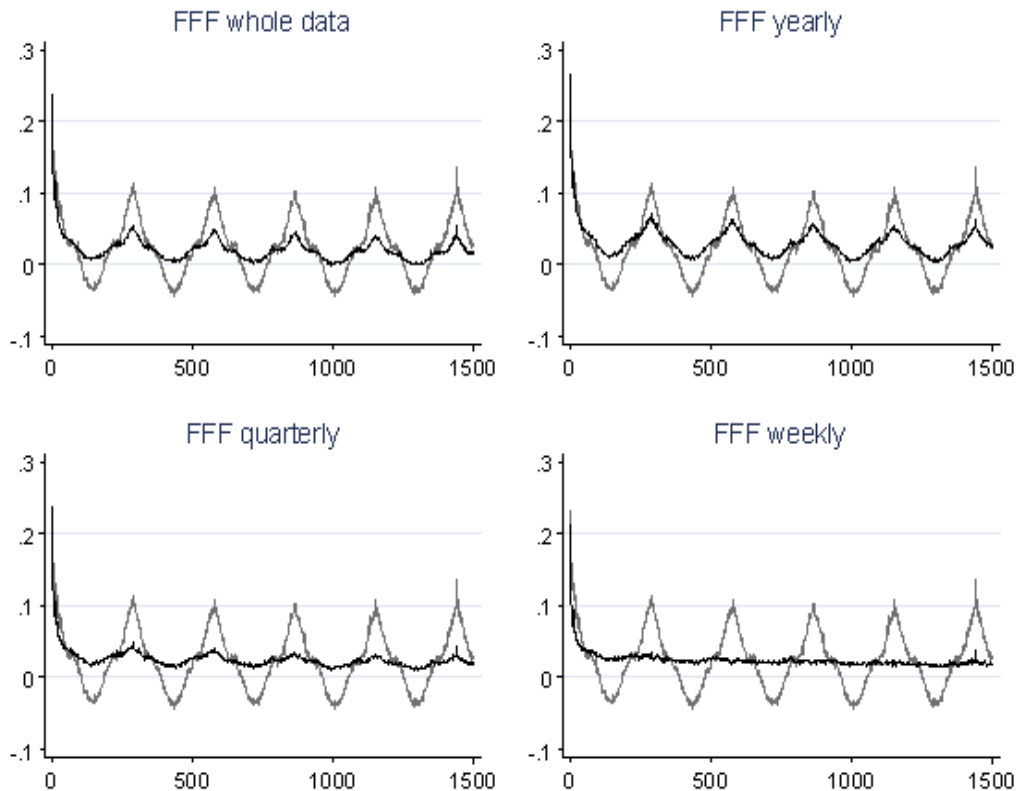


Figure 3.1 Autocorrelograms of raw and filtered absolute returns.

The figures graph the five day autocorrelogram of the raw and filtered 5-minute absolute USD/EUR returns. Intraday periodicity was filtered by using the Flexible Fourier Form method. The FFF model was estimated by using the whole data (upper left corner) and re-estimating the model yearly (upper right corner), quarterly (lower left corner) and weekly (lower right corner).

As can be seen from the first figure, if the whole data is estimated at once, the model is clearly not capable of filtering out all the periodicity. There is a lot of periodicity left in volatility also after filtering. The situation is not better even when the model is re-estimated every year. However, when the model is estimated separately for each quarter, the filter performs better and if the model is estimated separately for each week, there is no periodicity left in the autocorrelation function of absolute returns. Therefore, since estimating the model in subsets seems to improve filtering, it seems that seasonality of volatility is not constant in time.

3.2 Locally Weighted Scatterplot Smoothing method

Of the smoothing techniques, we consider the Locally Weighted Scatterplot Smoothing method (LOWESS) introduced by Cleveland (1979). Even though the smoothing techniques are often used in the statistics and in the seasonality filtering literature (especially in the Autoregressive Conditional Duration models), they are not so common in the literature studying news effects. The basic idea in the LOWESS method is to create a new variable $\hat{s}_{t,n}$ which contains the corresponding smoothed value for each point of the series $|R_{t,n}|$. The smoothed values are obtained by running a time trend estimation separately for each point of the data and a small number of observations close to each point. The smoothed value (which is the estimate for intraday volatility $\hat{s}_{t,n}$) is then the predicted value for the particular point only, which means that a separate regression is performed for every point in the data. In addition, the observations are weighted so that observations close to the predicted point get larger weights than those further away.

Figure 3.2 presents the autocorrelation coefficients of the filtered absolute returns compared to the raw absolute returns, where the seasonality is filtered out by the LOWESS method. The length of the subset used in the estimation affects the smoothness of the estimated curve: the shorter the estimation subset, the more precisely the smoothed curve follows the original data. We present the results with three different values of parameter δ (0.0003, 0.0002, 0.0001), which controls the length of the estimation subset. The smaller the value of δ , the shorter the estimation subset.⁸ As can be seen, when the parameter value decreases there is less periodicity left after filtering. When the smoothness parameter gets the value 0.0001, there is no autocorrelation left in the filtered absolute returns at all.

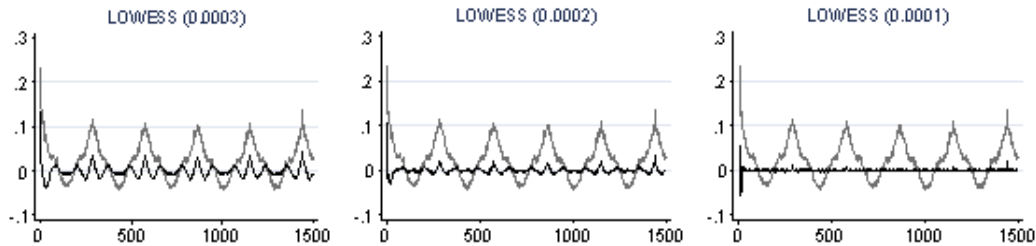


Figure 3.2 Autocorrelograms of raw and filtered absolute returns

The figures graph the five day autocorrelogram of the raw and filtered 5-minute absolute USD/EUR returns. Intraday periodicity was filtered by using the Locally Weighted Scatterplot Smoothing method. Three different values was considered for parameter δ (0.0003,0.0002,0.0001), which controls the length of the estimation subset. The smaller the value of δ , the shorter the estimation subset.

⁸For these parameter values the length of the subset is approximately six, four and two hours, respectively.

3.3 The Intraday Average Observations Model

The method that we consider of the average observations category is called the Intraday Average Observations Model (IAOM), introduced by Bauwens, Ben Omrane and Giot (2005). This recently proposed method has been used in the macro news literature by Bauwens et al (2005).

The method that we use differs slightly from the one originally proposed by the authors. First, they did not exclude any holidays from their data and second, their definition of weekend differs from ours.⁹ The intraday volatility estimate \hat{s}_{n_k} is computed by averaging the squared returns per each intraday interval (separately for each weekday) and then taking the square root

$$\hat{s}_{n_k} = \left(\frac{1}{M} \sum_{m=1}^M R_{m,k,n}^2 \right)^{1/2} \quad (3.4)$$

where $k = 1, \dots, 5$ denotes a weekday,¹⁰ $n = 1, \dots, N$ denotes the intraday interval ($N = 288$ for 5 minute intervals in the 24 hour market) and M denotes the number of weeks in the data set.

Also in the case of Intraday Average Observations Model, the filtering could be done separately in different subsets. Besides filtering the whole data at once, we filtered the returns separately for each year and each quarter. Figure 3.3 presents the autocorrelation coefficients of the filtered absolute returns compared to the raw absolute returns, where seasonality is filtered out by using the IAOM method. Shortening the subset affects the results as in the case of the other two methods: the shorter the subset the better the filter performs. However, the differences are not as large as in the case of the FFF and LOWESS methods.

⁹They exclude the intervals from Fridays at 22:05 through Saturday and Sunday and the first interval on Monday, while we use the definition of Andersen and Bollerslev (1998) and exclude the observations from Friday 21:05 until Sunday 21:00. We also tested for including holidays, and it did not have remarkable effect on the results. However, if the holidays and weekends are both included in the data, the missing observations cause a significant positive autocorrelation for the first lags in the filtered returns.

¹⁰For simplicity, Friday and Sunday observations are combined together and treated as one complete day $k = 5$. This should not have any implications to the results, since Friday and Sunday do not share any intervals. Friday observations always end at 21:00, while Sunday observations always begin at 21:05.

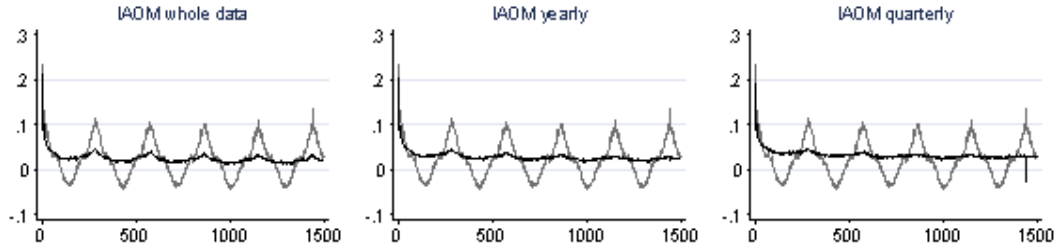


Figure 3.3 Autocorrelograms of raw and filtered absolute returns

The figures graph the five day autocorrelogram of the raw and filtered 5-minute absolute USD/EUR returns. Intraday periodicity was filtered by using Intraday Average Observations method. The IAOM model was estimated by using the whole data and re-estimating the model yearly and quarterly.

3.4 The Impact of News on Volatility

The intraday periodicity is caused by differences in trading times in the global foreign exchange markets (mainly US, European and Asian markets). Therefore, we consider the macro news from these three markets. The data contain all the scheduled macroeconomic news (for example GDP figures, interest rate announcements, confidence indices etc.) for Japan, Germany and the USA published in the World Economic Calendar (WECO) of Bloomberg.

Some of the macro figures are always announced at the same time, and therefore they might also cause some seasonality in the intraday volatility. Most influential of there kind of regular news are probably the US news announced at 13:30 GMT and 15:00 GMT. We wanted to study these news separately and therefore divided the US macro announcements into two categories: ‘USA regular news’ includes the news that are always announced at 13:30 GMT and 15:00 GMT and ‘USA news’ includes the rest of the US announcements.

In the original model of Andersen and Bollerslev (1997) the news variables $N_k(t, n)$ are included to the FFF model (3.3). This is also what we do when the FFF model is estimated once for the whole data set. On the other hand, when filtering is done in subsets, we need two steps for testing the news effects. The first step is to filter the returns, and the second step is to study the impact of news on volatility of the filtered returns. So, the two-step procedure is used in the subset FFF model, and in both LOWESS and IAOM methods, since the news variables cannot be included to the LOWESS and IAOM models in the same manner as in the FFF model.

To test the news effects in the second step, we use the model

$$2 \ln \frac{|\tilde{R}_{t,n} - \bar{R}|}{\hat{\sigma}_t / N^{1/2}} = c + \sum_{k=1}^4 \lambda_k N_k(t, n) + \varepsilon_{t,n} \quad (3.5)$$

where $\tilde{R}_{t,n}$ denotes the filtered returns and $N_k(t, n)$ the news variables ($k =$ Japan, Germany, USA, USA regular). c is the constant term and $\varepsilon_{t,n}$ is the error term of the model. If we were only interested in the impact that the macro figure has immediately after the announcement, the news variables $N_k(t, n)$ would be dummy variables taking the value one when the macro figure is announced and zero otherwise. However, it has been reported that the impact of news lasts from one to two hours (Andersen et al, 2003). Therefore we follow Andersen and Bollerslev (1998) and estimate the decay structure of the volatility response pattern of news (Figure 3.4) using a third order polynomial $\gamma(i)$

$$\gamma(i) = 0.054 (1 - (i/25)^3) - 0.009 (1 - (i/25)^2) i + 0.0007 (1 - (i/25)) i^2 \quad (3.6)$$

where $i = 1, 2, \dots, 25$ 5-minute intraday intervals. The polynomial captures the average decay structure (mean absolute returns after the news announcements) quite well and forces the impact to zero after two hours (when $i = 25$). Now, when the macro figure is announced, the news variable $N_k(t, n)$ takes the value $\gamma(i)$ for the first 25 intervals after the announcement and zero otherwise. The impact of news on volatility $M_k(i)$ can then be computed for each 25 intervals as $M_k(i) = \exp(\frac{\lambda_k \cdot \gamma(i)}{2}) - 1$.

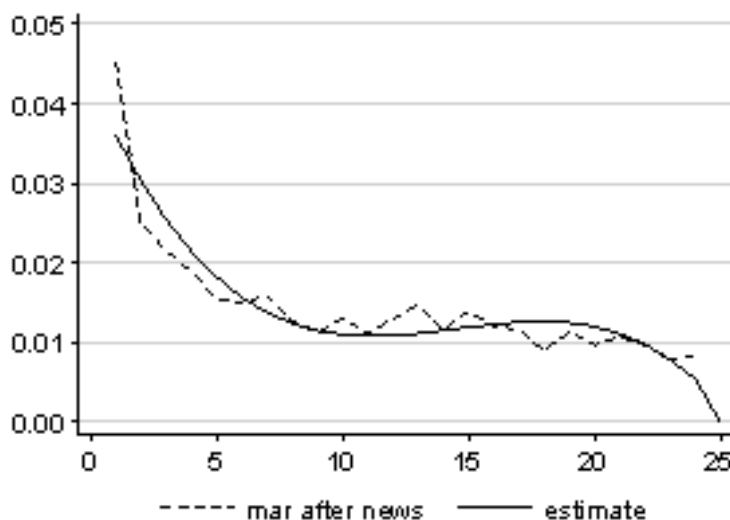


Figure 3.4 Decay structure of volatility response pattern after news

The figure presents the mean absolute returns (mar) after news announcements and the estimated news impact decay structure. Decay structure was estimated using the third order polynomial $\gamma(i)$.

Table 3.1 presents the coefficient values λ_k for the news variables, as well as the impact of news $M_k(1)$, computed for the first interval, ie five minutes after the announcement. The first row presents the results obtained with the returns which were not filtered at all. The following lines present the results obtained with the returns filtered with different methods. What is clear is that filtering

has a significant effect on the news effects. The results obtained with the non-filtered data are very different compared to the results obtained with any of the filtered data.

As can be seen, the regular US news seems to have much greater effects than any other news groups. While the news from Germany and the USA increase volatility significantly, it seems that the news from Japan does not have an effect on the volatility of the USD/EUR exchange rate. Also, the filtering method seems to affect the magnitude of the macro announcement coefficients. The estimated impact of macro news from Japan on volatility differs from decrease of 3% to increase of 10% depending on the used method and subset. The differences are very large also in other groups of news: the increase caused by the regular US news is estimated to be 52% at the lowest, and 140% at the highest. What is also worth noticing, is that there are quite clear patterns how the estimated news impacts depend on the used subset: in the case of the FFF models, the shorter the subset, the larger the news coefficients (except in the case of Japan). In contrast, for the LOWESS and IAOM models the news coefficients seem to decrease in size when the subset is shorter. However, the differences are not that large in the case of the IAOM than in the case of the LOWESS and FFF methods.

Figure 3.5 graphs the estimated intraday patterns $\hat{s}_{t,n}$ of different FFF and LOWESS methods, compared to the average volatility during a day. These figures might help us to understand the patterns seen in the news coefficient values. Since filtering is done by dividing the returns by $\hat{s}_{t,n}$, it is clear that the smaller is $\hat{s}_{t,n}$ the less we decrease the volatility. On the other hand, the less the volatility is decreased, the less news effects are cut down. Now, we can see from the estimates of $\hat{s}_{t,n}$ for the FFF model that the shorter the subset, the lower is the estimate $\hat{s}_{t,n}$ during the opening hours of the European markets and the US markets (from 7:00 GMT to 19:00 GMT). Therefore, the shorter is the subset, the less we cut down the US and European news effects while filtering seasonality. Yet, it can be understood that the shorter is the subset, the larger are the estimated news effects. The situation is completely opposite in the LOWESS method. The shorter is the subset, the higher is the estimate $\hat{s}_{t,n}$ during the opening hours of the biggest markets (Asia, Europe and USA). So now, the shorter is the subset, the more we cut down the volatility and also news effects. Therefore, it is understandable that while the subset is shortened, the estimated news effects are decreased.

Table 3.1 Impact of macroeconomic news on USD/EUR volatility

The impact of macroeconomic news on USD/EUR volatility was studied with returns which were filtered by different methods. FFF refers to Flexible Fourier Form models, LOWESS to Locally Weighted Scatterplot Smoothing method and IAOM to Intraday Average Observations Method. See details in the sections 3.1–3.3. λ_k refers to the coefficients of the news variables in the model (3.5) and $M_k(1)$ refers to the news impact on volatility 5 minutes after the news announcement. Newey-West t-values are in parentheses (288 lags).

	Japan		Germany		USA		USA regular	
	λ_k	$M_k(1)$	λ_k	$M_k(1)$	λ_k	$M_k(1)$	λ_k	$M_k(1)$
No filtering	-16.15 (-9.51)	-0.25	46.99 (30.95)	1.33	39.90 (12.76)	1.05	89.81 (59.88)	4.04
FFF whole data	2.17 (1.22)	0.04	6.44 (4.18)	0.12	20.23 (7.06)	0.44	33.54 (18.29)	0.83
FFF yearly	5.50 (3.55)	0.10	7.68 (5.93)	0.15	15.93 (6.08)	0.33	23.41 (17.78)	0.52
FFF quarterly	2.18 (1.40)	0.04	12.35 (9.29)	0.25	18.55 (6.84)	0.40	30.18 (22.42)	0.72
FFF weekly	0.09 (0.06)	0.00	16.21 (12.44)	0.34	19.42 (7.36)	0.42	34.83 (26.85)	0.87
LOWESS $\delta(0.0003)$	1.01 (0.71)	0.02	27.74 (24.62)	0.65	24.45 (10.76)	0.55	48.63 (45.63)	1.40
LOWESS $\delta(0.0002)$	-0.13 (-0.10)	0.00	23.69 (22.72)	0.53	21.00 (9.99)	0.46	40.20 (41.16)	1.06
LOWESS $\delta(0.0001)$	-1.43 (-1.18)	-0.03	18.94 (20.33)	0.41	16.10 (8.61)	0.34	31.69 (36.28)	0.77
IAOM whole data	-1.50 (-0.92)	-0.03	21.46 (15.01)	0.47	21.75 (7.96)	0.48	33.54 (24.97)	0.83
IAOM yearly	-0.30 (-0.19)	-0.01	21.06 (15.33)	0.46	21.49 (8.08)	0.47	31.88 (24.08)	0.78
IAOM quarterly	-0.41 (-0.26)	-0.01	20.53 (15.22)	0.45	20.54 (7.73)	0.45	30.91 (23.65)	0.74




Figure 3.5 Estimated intraday volatility pattern

Figure 3.5a graphs the estimated intraday volatility patterns of the FFF models compared to mean absolute five minute returns per intraday intervals (MAR). Figure 3.5b presents the same for the LOWESS models.

4 Properties of the filtered returns

In this section we study the statistical key figures of the original and filtered returns to see whether the filtering causes some unwanted effects (eg extra jumps) in the returns.

The key statistical figures for the different filtered returns are presented in Table 4.1. As expected, filtering does not affect the mean or standard deviation of the returns. On the other hand, filtering seems to have an effect on the third and fourth moments. While both LOWESS and IAOM method seem to significantly decrease both skewness and kurtosis, the FFF method has no such clear result. In some cases skewness and kurtosis are decreased and in some cases increased, when the returns are filtered with the FFF method. What is clear is that the kurtosis and skewness are larger in the returns filtered by FFF models than those filtered by LOWESS and IAOM methods.

Table 4.1 Key statistical figures

The key statistical figures for original and filtered returns. FFF refers to Flexible Fourier Form models, LOWESS to Locally Weighted Scatterplot Smoothing method and IAOM to Intraday Average Observations Method.

	mean	Standard Deviation	Skewness	Kurtosis	Minimum	Maximum
USD/EUR	5.0E-05	0.0432	0.781	65.94	-1.35	2.79
<hr/>						
FFF						
whole data	9.3E-05	0.0445	-0.230	76.00	-2.14	2.41
FFF						
yearly	1.2E-04	0.0464	-2.138	239.44	-2.13	2.49
FFF						
quarterly	8.6E-05	0.0439	-0.618	99.95	-2.51	2.38
FFF						
weekly	6.6E-05	0.0434	-0.154	40.92	-1.69	1.68
<hr/>						
LOWESS						
$\delta(0.0003)$	8.0E-05	0.0387	0.011	8.94	-0.95	0.74
LOWESS						
$\delta(0.0002)$	8.4E-05	0.0381	0.024	6.68	-0.67	0.56
LOWESS						
$\delta(0.0001)$	9.1E-05	0.0373	0.007	5.01	-0.48	0.41
<hr/>						
IAOM						
whole data	6.5E-05	0.0407	0.000	10.01	-0.64	0.73
IAOM						
yearly	8.1E-05	0.0400	0.002	6.06	-0.32	0.37
IAOM						
quarterly	8.3E-05	0.0385	0.000	3.89	-0.18	0.18




Figure 4.1 The raw and filtered USD/EUR returns

The figure in the upper left corner present the 5 minute returns on the USD/EUR exchange rate from 1 Jan. 1999 to 31 Dec. 2004. The figure in the upper right corner presents the returns filtered with the Flexible Fourier Form method. The figures in the lower left and right corner present the returns filtered with the Locally Scatterplot Smoothing Method and Intradaily Average Observation method, respectively.

The same can be also seen in Figure 4.1, which graphs the original returns and all the different filtered returns. It seems that while the LOWESS and IAOM methods seem to shrink the large jumps in the returns, the FFF model saves the jumps but also creates some new ones. Neither of results is what we want: the filter should not filter out ‘important’ jumps, nor should it create any new ones. However, compared to the returns filtered with the LOWESS and IAOM methods, the returns filtered with the FFF model most closely resemble the original returns in terms of key statistical figures.

5 Simulation Study

In this section we study more carefully whether the choice of filtering method has an effects on statistical inference concerning the impact of news on exchange rate volatility. We construct returns by using the properties of the USD/EUR data, simulate 2000 realizations with 288000 observations (1000

days), deseasonalize the simulated returns by the same three different filtering methods, and test the impact of news variables on the volatility of filtered returns.

5.1 Returns

The returns were constructed from the daily volatility component $\frac{\sigma_{t,n}}{\sqrt{288}}$, intraday volatility component $s_{t,n}$, the news component $\eta_{t,n}$ and the error term $\varepsilon_{t,n}$

$$R_{t,n}^S = \frac{\sigma_{t,n}}{\sqrt{288}} s_{t,n} \eta_{t,n} \varepsilon_{t,n} \quad (5.1)$$

The daily volatility component $\sigma_{t,n}$ was simulated using a GARCH(1,1) model. The estimated coefficient values from daily USD/EUR data were used as trend-setters for the simulated model.

To prevent any possibilities to favour any filtering methods in the simulation study, we do not want to use any of the studied models¹¹ to create the intraday seasonality component $s_{t,n}$. Therefore, we used a modified version of the dummy variable model proposed by Degenaro and Schrieves (1997). We estimated the following model using the USD/EUR data set

$$2 \ln \frac{|R_{t,n} - \bar{R}|}{\hat{\sigma}_t / N^{1/2}} = \alpha + \sum_{h=1}^{47} \beta_h D_h(t, n) + \varepsilon_{t,n} \quad (5.2)$$

where the $D_h(t, n)$ are 47 half-hour dummy variables.¹² To capture the possible time variation in intraday volatility, the model was estimated separately for each month in the data set. The fitted values of the model were then saved and used as intraday volatility component $s_{t,n}$.

The news component $\eta_{t,n}$ was constructed as follows: we first created a preliminary news variable (I_N), which takes the value of one in 2% of the observations and zero otherwise, by using equally distributed [1,0] random variables. We then created three dummy variables which indicate whether a particular market (US, European or Asian) is open or not. For example for the US markets this variable (I_{M_usa}) gets the value one if the US markets are open and zero otherwise. We also created a dummy variable I_{M_usar} which takes the value of one at 13.30 GMT and 15.00 GMT and zero otherwise. This was done to create the ‘regular US news’ variable.

The market-specific news variables were then created as $I_k = I_N \cdot I_{M_k}$, for $k = \text{Asia, Europe, USA, USA regular}$. We then had four news variables that take the value of one if the macro news are announced and zero otherwise. However, since the news impact lasts longer than five minutes, we used the

¹¹Flexible Fourier Form model is the only one of the three methods that could have been considered of using.

¹²There are 48 half-hour intervals in the 24 hour market. Since the model includes the constant, we need 47 dummy variables.

news impact pattern estimated from the USD/EUR data (see section 3.4). We used the estimated coefficient values from the USD/EUR data as trend-setters and set the news coefficients values such, that the Asian news increase volatility by 5%, the European news by 20%, the US news 40%, and the regular US by 80% 5 minutes after the announcement. The news impact patterns $\gamma_k(i)$ (where $k = asia, europe, usa, usar$ and $i = 1, \dots, 25$ five minute intraday intervals) are presented in Figure 5.1.

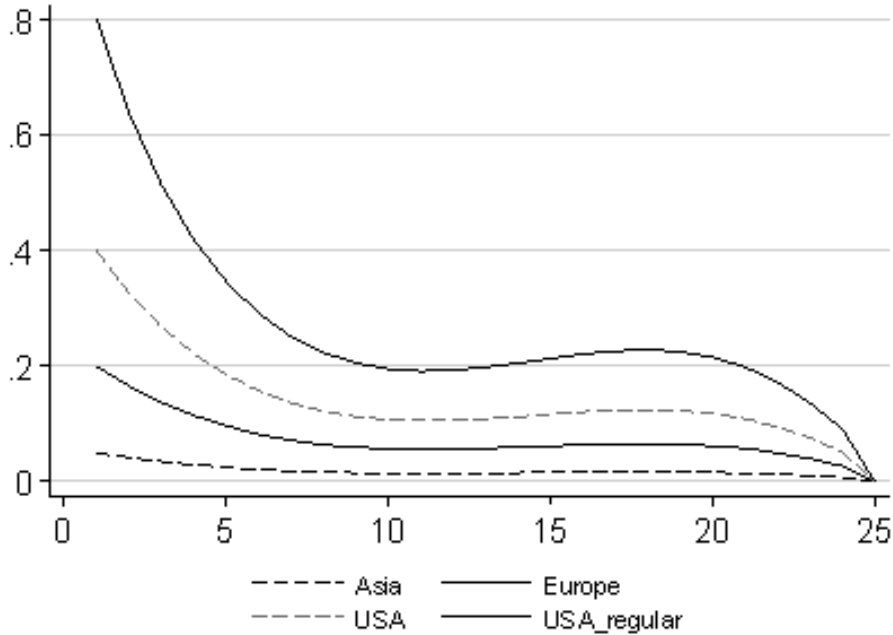


Figure 5.1 News impact decay structures $\gamma_k(i)$

Now, the variables $(N_{asia}, N_{europe}, N_{usa}, N_{usar})$ get the values of the polynomials $\gamma_k(i)$ during the first 25 intervals after the announcement (when the variable $I_{asia}, I_{europe}, I_{usa}$ or I_{usar} equals one) and zero otherwise. Finally, the news component $\eta_{t,n}$ in (5.1) is constructed as

$$\eta_{t,n} = 1 + N_{usa} + N_{europe} + N_{asia} + N_{usar} \quad (5.3)$$

To obtain returns that resemble the actual returns, we need to generate the shocks from a more leptokurtic distribution than the normal distribution. Therefore, we create the random shocks $\varepsilon_{t,n}$ as a mixture of two normally distributed random variables $\varepsilon_{1,tn}$ and $\varepsilon_{2,tn}$ such that $\varepsilon_{t,n} = \varepsilon_{1,tn} \sim N(0, 0.5)$ with probability 0.75 and $\varepsilon_{t,n} = \varepsilon_{2,tn} \sim N(0, 2.0)$ with probability 0.25.

Therefore, while the daily volatility component (which depends on $\varepsilon_{t,n}$) and the random shocks change in every round, the intraday volatility component and the news component remain the same.

5.2 Properties of the simulated returns

To see whether the simulated returns have the same kinds of properties as the USD/EUR returns, we computed the average volatility pattern and the autocorrelation function of the return volatility for one realization (Figure 5.2). As can be seen, the simulated series display an intraday seasonality similar to that of the USD/ERUR data set used before.

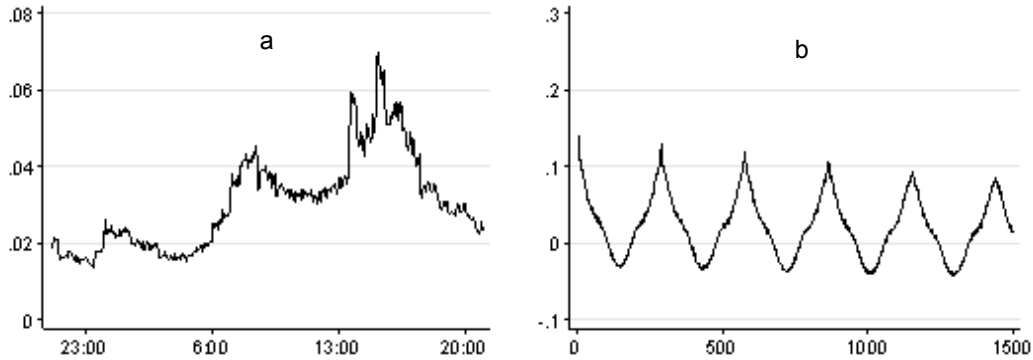


Figure 5.2 Intraday seasonality of volatility (simulated returns)

Figure 5.2a graphs the average intraday volatility pattern of the simulated absolute returns. Figure 5.2b graphs the five day autocorrelogram of the simulated absolute returns.

The simulated returns are filtered by the three methods: FFF, LOWESS and IAOM. The subset lengths in all of the methods are selected to be such that there is no periodicity left in volatility after filtering. For the IAOM method, no distinction between the weekdays was made. Since we did not create differences between weekdays in the intraday pattern, the distinction is not necessary.

To demonstrate how well the simulated returns resemble the USD/EUR returns, we computed the descriptive statistics for the simulated returns and the filtered simulated returns for one realization (Table 5.1). The key figures of the simulated returns are quite close to the ones of the USD/EUR returns. Mean and standard deviation of the simulated returns are very close to the ones of USD/EUR returns, but skewness and kurtosis of the simulated returns are smaller than those of the USD/EUR returns. When using the USD/EUR returns, filtering does not affect the mean and standard deviation, but rather the third and fourth moments. Similar findings can be made when the simulated returns are used.

Table 5.1 Key statistical figures for the simulated returns

The key statistical figures of the simulated returns and the simulated returns filtered with different methods. FFF refers to Flexible Fourier Form method, LOWESS to Locally weighted scatterplot smoothing method and IAOM to Intradaily Average Observations Model.

Method	Mean	St. Dev.	Skewness	Kurtosis
returns	-0.00006	0.056	0.016	21.34
FFF	-0.00011	0.053	0.040	17.10
LOWESS	-0.00012	0.045	0.003	7.03
IAOM	-0.00009	0.053	0.017	15.29

5.3 Simulation results

After deseasonalizing the simulated returns we studied the impact of news on volatility of filtered returns in the same manner as with the USD/EUR returns (equation 3.5). Table 5.2 presents the results of the simulation study. Besides using the returns filtered with the three different methods, we estimated the news impact on volatility by using returns that were not filtered at all. As can be seen, filtering out the seasonality is indeed crucial: if we do not filter the returns, the US and European news coefficients are biased upward and Asian news coefficients downward. We also estimated the news impact on volatility by using the returns which did not have intraday seasonality in volatility, ie excluding intraday seasonality component from the return process. As expected, the intraday volatility component is the one causing the problems, since without it the bias is very close to zero.

It seems that almost all of the filtering methods tend to produce downward biased estimates for US and European news, and upward biased estimates for Asian news. However, the magnitude of the bias depends on the filter. While the LOWESS is best in terms of filtering the periodicity in volatility, it also seems to produce a larger negative bias than the other two filters. This means that, while filtering the intraday seasonality, it also filters part of the news effects. While the IAOM seems to perform almost as well as the FFF in most cases, it performs much worse than the FFF model when the macro figures are announced regularly. The FFF model produces the smallest bias on average, and therefore we conclude that it performs the best in seasonality filtering.

Table 5.2 Simulation results

2000 simulated realizations of the return process were filtered by FFF, LOWESS and IAOM methods. After filtering the news impact on volatility of different filtered returns was tested. Table presents the key statistics of the estimated news coefficients. The first panel presents the results of the Asian news, the second panel the European news, the third panel the US news and the last panel the regular US news. The first two rows in each panel present the results when the returns were not filtered at all and when the intraday volatility component was excluded in the DGP when the returns were simulated. The next rows in each panel present the results when the returns were filtered with the FFF, LOWESS and IAOM methods, respectively. The mean bias, bias standard deviation and minimum and maximum of the estimated news variable coefficient values are presented in the four last columns.

ASIA, 1.05				
Method	Bias	St. Dev.	Min	Max
Returns – no filtering	-0.38	0.012	0.62	0.71
Returns – no intraday seasonality	0.00	0.019	0.98	1.11
FFF	0.04	0.016	1.03	1.15
LOWESS	-0.01	0.012	1.00	1.07
IAOM	0.01	0.022	0.99	1.14

EUROPE, 1.2				
Method	Bias	St. Dev.	Min	Max
Returns – no filtering	0.96	0.038	2.05	2.28
Returns – no intraday seasonality	-0.01	0.021	1.13	1.26
FFF	-0.02	0.018	1.11	1.23
LOWESS	-0.12	0.013	1.03	1.12
IAOM	0.06	0.026	1.16	1.33

USA, 1.4				
Method	Bias	St. Dev.	Min	Max
Returns – no filtering	-0.05	0.024	1.28	1.43
Returns – no intraday seasonality	-0.02	0.024	1.30	1.46
FFF	-0.09	0.019	1.24	1.38
LOWESS	-0.30	0.012	1.06	1.16
IAOM	-0.12	0.024	1.21	1.36

USA REGULAR, 1.8				
Method	Bias	St. Dev.	Min	Max
Returns – no filtering	2.95	0.092	4.48	5.06
Returns – no intraday seasonality	-0.04	0.034	1.66	1.88
FFF	-0.37	0.021	1.37	1.50
LOWESS	-0.55	0.017	1.20	1.31
IAOM	-0.51	0.031	1.18	1.39

6 Conclusions

In this paper we studied the capability of different methods to filter out the intraday seasonality of volatility and whether or not the choice of the filtering method affects the results concerning the impact of news on volatility. The results suggest that there are differences between the filters. The FFF model performs poorly if the model is estimated for the whole data set at once, but there is no periodicity left if the model is re-estimated every week. The success of the LOWESS method depends heavily on the chosen value of the smoothness parameter δ . When δ is set to be small enough, the filter is capable of getting rid of all the periodicity in the autocorrelation. The performance of the IAOM method also depends on the length of the estimation subset. The shorter the subset, the better the filter performs. On the other hand, the estimation subset length does not have as large impact on the IAOM method than on the FFF and LOWESS methods.

The choice of the filtering method affects the magnitude of the news coefficients. According to the simulation study, all the methods tend to produce downward biased estimates, which means that while filtering out the intraday seasonality, they also filter out part of the news effects. However, the size of the bias depends on the filter. The magnitude of the news impact and the announcement time regularity also affect the results: the larger news items are filtered more than the smaller ones and the news items that are announced regularly are filtered the most. While the LOWESS is capable of filtering out all the periodicity in volatility, it also seems to filter out more news effects than the other two filters. IAOM performs much worse than FFF in the case of news items that are always announced at the same time. The study supports the FFF model as the best seasonality filtering method.

References

- Andersen, T – Bollerslev, T (1997) **Heterogenous Information Arrivals and Return Volatility Dynamics: Unrecovering the Long-Run in High Frequency Returns.** The Journal of Finance 52, 975–1005.
- Andersen, T – Bollerslev, T (1998) **Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies.** The Journal of Finance 53, 219–265.
- Andersen T – Bollerslev, T – Diebold, F – Vega, C (2003) **Micro Effects and Macro Announcements: Real-Time Price Discovery in Foreign Exchange.** American Economic Review 93, 38–62.
- Bauwens, L – Ben Omrane, W – Giot, P (2005) **News Announcements, Market Activity and Volatility in the Euro/Dollar Foreign Exchange Market.** Journal of International Money and Finance 24, 1108–1125.
- Ben Omrane, W – Bodt, E (2005) **Using self-organizing maps to adjust intra-day seasonality.** Louvain School of Management Working paper 149/05.
- Cai, J – Cheung, Y – Lee, R – Melvin, M (2001) **Once in a generation yen volatility in 1998: fundamentals, intervention and order flow.** Journal of International Money and Finance 20, 327–347.
- Cleveland, W (1979) **Robust Locally Weighted Regression and Smoothing Scatterplots.** Journal of the American Statistical Association 74, 829–836.
- Dacorogna, M – Müller, 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.
- DeGennaro, R – Shrieves, R (1997) **Public information releases, private information arrival and volatility in the foreign exchange market.** Journal of Empirical Finance 4, 295–315.
- Dominguez, K – Panthaki, F (2006) **What defines ‘news’ in foreign exchange markets?** Journal of International Money and Finance 25, 168–198.
- Engle, R – Russell, J (1998) **Autoregressive conditional duration: a new approach for irregularly spaced transaction data.** Econometrica 66, 1127–1162.
- Granger, C – Sin, C-Y (1999) **Modelling the absolute returns of different stock indices: exploring the forecast ability of an alternative measure of risk.** University of California, San Diego, Department of economics.

Laakkonen, H (2007) **The Impact of macroeconomic news on exchange rate volatility.** Finnish Economic Papers 20, 23–40.

Martens, M – Chang, Y-C – Taylor, S (2002) **A comparison of seasonal adjustment methods when forecasting intraday volatility.** Journal of Financial Research 25, 283–299.

Melvin, M – Yin, X (2000) **Public information arrival, exchange rate volatility and quote frequency.** The Economic Journal 110, 644–661.

Veredas D – Rodriguez-Poo, J – Espasa, A (2002) **On the (intradaily) Seasonality and Dynamics of a Financial Point Process: A Semiparametric Approach.** CORE DP 2002/23.

**BANK OF FINLAND RESEARCH
DISCUSSION PAPERS**

ISSN 0785-3572, print; ISSN 1456-6184, online

- 1/2007 Timo Korkeamäki – Yrjö Koskinen – Tuomas Takalo **Phoenix rising: Legal reforms and changes in valuations in Finland during the economic crisis.** 2007. 39 p. ISBN 978-952-462-346-9, print; ISBN 978-952-462-347-6, online.
- 2/2007 Aaron Mehrotra **A note on the national contributions to euro area M3.** 2007. 25 p. ISBN 978-952-462-348-3, print; ISBN 978-952-462-349-0, online.
- 3/2007 Ilmo Pyyhtiä **Why is Europe lagging behind?** 2007. 41 p. ISBN 978-952-462-350-6, print; ISBN 978-952-462-351-3, online.
- 4/2007 Benedikt Goderis – Ian W Marsh – Judit Vall Castello – Wolf Wagner **Bank behaviour with access to credit risk transfer markets.** 2007. 28 p. ISBN 978-952-462-352-0, print; ISBN 978-952-462-353-7, online.
- 5/2007 Risto Herrala – Karlo Kauko **Household loan loss risk in Finland – estimations and simulations with micro data.** 2007. 44 p. ISBN 978-952-462-354-4, print; ISBN 978-952-462-355-1, online.
- 6/2007 Mikael Bask – Carina Selander **Robust Taylor rules in an open economy with heterogeneous expectations and least squares learning.** 2007. 54 p. ISBN 978-952-462-356-8, print; ISBN 978-952-462-357-5, online.
- 7/2007 David G Mayes – Maria J Nieto – Larry Wall **Multiple safety net regulators and agency problems in the EU: is Prompt Corrective Action a partial solution?** 2007. 39 p. ISBN 978-952-462-358-2, print; ISBN 978-952-462-359-9, online.
- 8/2007 Juha Kilponen – Kai Leitemo **Discretion and the transmission lags of monetary policy.** 2007. 24 p. ISBN 978-952-462-362-9, print; ISBN 978-952-462-363-6, online.
- 9/2007 Mika Kortelainen **Adjustment of the US current account deficit.** 2007. 35 p. ISBN 978-952-462-366-7, print; ISBN 978-952-462-367-4, online.
- 10/2007 Juha Kilponen – Torsten Santavirta **When do R&D subsidies boost innovation? Revisiting the inverted U-shape.** 2007. 30 p. ISBN 978-952-462-368-1, print; ISBN 978-952-462-369-8, online.
- 11/2007 Karlo Kauko **Managers and efficiency in banking.** 2007. 34 p. ISBN 978-952-462-370-4, print; ISBN 978-952-462-371-1, online.

- 12/2007 Helena Holopainen **Integration of financial supervision.** 2007. 30 p.
ISBN 978-952-462-372-8, print; ISBN 978-952-462-373-5, online.
- 13/2007 Esa Jokivuolle – Timo Vesala **Portfolio effects and efficiency of lending under Basel II.** 2007. 23 p. ISBN 978-952-462-374-2, print;
ISBN 978-952-462-375-9, online.
- 14/2007 Maritta Paloviita **Estimating a small DSGE model under rational and measured expectations: some comparisons.** 2007. 30 p.
ISBN 978-952-462-376-6, print; ISBN 978-952-462-377-3, online.
- 15/2007 Jarmo Pesola **Financial fragility, macroeconomic shocks and banks' loan losses: evidence from Europe.** 2007. 38 p. ISBN 978-952-462-378-0, print;
ISBN 978-952-462-379-7, online.
- 16/2007 Allen N Berger – Iftekhar Hasan – Mingming Zhou **Bank ownership and efficiency in China: what lies ahead in the world's largest nation?** 2007.
47 p. ISBN 978-952-462-380-3, print; ISBN 978-952-462-381-0, online.
- 17/2007 Jozsef Molnar **Pre-emptive horizontal mergers: theory and evidence.** 2007.
37 p. ISBN 978-952-462-382-7, print; ISBN 978-952-462-383-4, online.
- 18/2007 Federico Ravenna – Juha Seppälä **Monetary policy, expected inflation and inflation risk premia.** 2007. 33 p. ISBN 978-952-462-384-1, print;
ISBN 978-952-462-385-8, online.
- 19/2007 Mikael Bask **Long swings and chaos in the exchange rate in a DSGE model with a Taylor rule.** 2007. 28 p. ISBN 978-952-462-386-5, print;
ISBN 978-952-462-387-2, online.
- 20/2007 Mikael Bask **Measuring potential market risk.** 2007. 18 p.
ISBN 978-952-462-388-9, print; ISBN 978-952-462-389-6, online.
- 21/2007 Mikael Bask **Optimal monetary policy under heterogeneity in currency trade.** 2007. 28 p. ISBN 978-952-462-390-2, print; ISBN 978-952-462-391-9,
online.
- 22/2007 Mikael Bask **Instrument rules in monetary policy under heterogeneity in currency trade.** 2007. 29 p. ISBN 978-952-462-392-6, print;
ISBN 978-952-462-394-0, online
- 23/2007 Helinä Laakkonen **Exchange rate volatility, macro announcements and the choice of intraday seasonality filtering method.** 2007. 32 p.
ISBN 978-952-462-394-0, print; ISBN 978-952-462-395-7, online.

Suomen Pankki
Bank of Finland
P.O.Box 160
FI-00101 HELSINKI
Finland



.2343