The Impact of News on Foreign Exchange Rates: Evidence from High Frequency Data*

Dirk Eddelbüttel† Thomas H. McCurdy‡

First Version: May 1996
This Version: May 1998

Abstract

This paper investigates the impact of the frequency of general and currency-specific news headlines on de-seasonalized intraday DEM-USD exchange rate changes. We find a significant relationship between volatility and the frequency of news. In particular, more news is associated with an increase in volatility. The result that spot exchange rates are more volatile during periods for which there is a lot of economic news accords with market participants' explanations for observed volatility clustering.

*We would like to thank Olsen & Associates, and especially Michel Dacorogna, for generously providing us with the data used in this study. We also thank Jean-Marie Dufour, Christian Dunis, Ramazan Gencay, James MacKinnon, Lonnie Magee, Baldev Raj, Francisco Ruge-Murcia as well as other participants at the Canadian Econometrics Study Group (Sept. 1996) and the Canadian Economics Association (May 1996) conferences. Financial support from the French Ministère de l'Éducation Nationale (Eddelbüttel) and the Social Sciences and Humanities Research Council of Canada (McCurdy) is gratefully acknowledged. McCurdy also acknowledges the hospitality and support of the Institute for Policy Analysis, University of Toronto.

†Bank of Montreal.

‡Rotman School of Management and Institute for Policy Analysis, University of Toronto.

1 Introduction

A prime example of the globalization of financial markets is the around-the-clock trading in the interbank foreign exchange (FX) market. The distributional characteristics of daily and lower frequency returns in this market have been the object of a very large literature. Now the increased availability of tick-by-tick data [1] has spurred interest in analysing *intraday* variation in returns and volatility. One important advantage of such high frequency data is that they allow a more detailed investigation of potential sources and characteristics of return volatility.

Guillaume et al. (1994) have surveyed some stylized facts associated with intraday patterns in FX markets. Some early evidence on the cross-sectional patterns in intraday FX data was provided by Müller et al. (1990). Early time-series applications, which built on the ARCH model of Engle (1982) to model the dynamics of intraday FX volatility, included: Engle, Ito, and Lin (1990) who used four points during the 24-hour trading day to study the transmission (spillover) of information across markets; and Baillie and Bollerslev (1991) who utilized hourly observations in a seasonal GARCH model with hourly dummy variables to study patterns of volatility across currencies and potential spillovers between markets.

Tick-by-tick data has allowed analyses of FX data at even higher frequencies. Goodhart and Figliuoli (1991) described the statistical characteristics of spot exchange rates with interval length of one minute, taken from several days of data from the Reuters FXFX screen. They documented the clear first-order autocorrelation in the data, especially after jumps in the level of the exchange rate, the changing level of activity throughout a 24-hour day as well as the effects of time aggregation on the data, but did not estimate an econometric model of the data. Moreover, while they attempted to link movements in the foreign exchange markets to news headlines from the Reuters AAMM screen, they failed to find such a relationship in the two days of Reuters AAMM data at their disposal.

Bollerslev and Domowitz (1993) examined the impact of quote arrival intensity and the size of the bid-ask spread on the volatility of five-minute returns from the DEM-USD foreign exchange market for the period April 9 to June 30, 1989. They found that only the bid-ask spread had a significant correlation with the conditional variance.

Goodhart, Hall, Henry, and Pesaran (1993) used eight weeks of data from the FXFX

[1] Perhaps the best known example is the Olsen & Associates' 'hfdf93' data set compiled from the interbank Reuters network for the period October 1, 1992 to September 29, 1993. These data include tick-by-tick bid and ask quotes for the Deutsche Mark–U.S. Dollar (DEM-USD), Japanese Yen–U.S. Dollar (JPY-USD), and JPY-DEM currencies, the ninety-day interbank deposit rates for these currencies, and the money-market news headlines (Reuters screen AAMM).

screen to estimate a GARCH-M model of the sterling-dollar exchange rate. They found that augmenting the conditional variance equation by adding two dummy variables, corresponding to two discrete news effects in their data, had a significant effect on their estimation results: the GARCH-M process became covariance stationary after having been close to an integrated GARCH process.

The year-long 'hfdf93' data sample has spurred further investigation of the effect of news on exchange rates. In a comprehensive analysis of DEM-USD volatility, Andersen and Bollerslev (1996) control for calendar effects, such as regional holidays and daylight savings time, as well as lunch breaks and data gaps. Their model incorporated the effects of regularly scheduled macroeconomic announcements[2], intraday activity patterns, and the persistent interdaily volatility structure. For example, they find that the announcements helped to explain day-of-the-week periodicities.

Baestens and Van den Bergh (1996) standardize the Reuters AAMM news headlines and estimate their impact on the dem/usd swap rate returns using ordinary least squares, an ARCH(2) model and neural net estimation. They find that some of the standardized news headlines systematically affect returns.

In a study of hourly JPY-USD exchange rate returns, DeGennaro and Shrieves (1995) add a weekend or vacation indicator variable, as well as the rate of quote arrival, to the conditional variance function. Having controlled for those two potential sources of seasonals in volatility, they find an increase in volatility *prior* to public announcements and a significant decrease in volatility during and following the hour of news arrival.

Our paper investigates the potential impact of the *frequency* of news on seasonally-adjusted intraday exchange rate changes.[3] In particular, we compute the news frequency from the money-market headlines of the Reuters screen AAMM. We consider different subsets

[2]Earlier analysis of the impact of macroeconomic announcements during the first hour of trading (usually at 8:30 EST) included Harvey and Huang (1991) and Hakkio and Pearce (1985). Ederington and Lee (1993) analyzed the impact of and speed of adjustment to nineteen *scheduled* monthly macroeconomic news releases on T-bond and Eurodollar interest rates and DEM foreign exchange futures. They found that while these announcements affected the direction of price adjustment for only one minute, the prices were considerably more volatile for fifteen minutes, and slightly more volatile for several hours.

[3]Two recent papers use similar data to analyse the effect of the frequency of news on exchange rates. Essential differences arise from different treatment of intraday seasonal effects. Low and Muthuswamy (1996) find that the number of contemporaneous news events is significantly related to the conditional volatility of *non-seasonally-adjusted* five-minute returns. Melvin and Yin (1996) use a longer sample (Dec. 1, 1993 to April 26, 1995) and deseasonalize by dividing hourly observations by their mean (over the sample) for that hour of the business week. That is, they investigate whether, for a particular hour of the business week, absolute returns are higher than average when news arrival is higher than average. They find that the previous hour's information variable marginally decreases the conditional variance of the JPY-USD returns but does not significantly affect the DEM-USD volatility.

of news (all versus currency-specific news) and also allow both the mean and the volatility of exchange rate changes to be a function of the relevant interest rate differential.[4] For these seasonally-adjusted intraday DEM-USD exchange rate changes, we find significant dependence of volatility on the frequency of news and on the interest rate differential. For example, more news is associated with an increase in volatility.

Given the pervasive impact that the intraday seasonal has on high frequency exchange rate dynamics (Andersen and Bollerslev (1994)), it is important to remove or account for the repetitive (unconditional or average) intraday seasonal pattern due to institutional features such as the level of market activity associated with the opening and closing of FX markets around the globe at different times of the day. To do so, we implement our analysis of conditional volatility on a time series of intraday exchange rate changes which has been deseasonalized by a particular time-scale transformation to ϑ-time (Dacorogna et al. 1993) which is based on the *average* market activity at each point in time. Since the Dacorogna et al. (1993) time scale removes the seasonal heteroskedasticity due to average market activity, we are able to focus on modelling the *conditional* heteroskedasticity due to economic news – including changes in the interest rate differential.

The ϑ-time transformation deseasonalizes volatility by expanding periods with high average market activity and contracting periods with low average market activity. As such, it controls for calendar effects such as those modelled explicitly in Andersen and Bollerslev (1996). Weekends will be contracted until the level of activity is similar to that during other periods. On the other hand, a more active market period will be expanded so that there is less clock-time in the ϑ-time interval. For example, a 10-minute ϑ-time interval may only contain clock-time from 8:30-8:34 a.m. In this case, we would count news events during that 4-minute interval. In other words, the time scale transformation for the news variable is appropriately matched to that for changes in the exchange rate returns. Therefore, news arrival will be deseasonalized to some extent as well. Note that we do not estimate an independent news activity time scale since this would create non-comparable time periods.

Several alternative time deformation approaches have been proposed to deal with intraday seasonals in foreign exchange data. For example, Ghysels, Gourieroux, and Jasiak (1995) proposed a stochastic volatility model which simultaneously accounts for both average and conditional market activity. Zhou (1996) used a statistical transformation. Andersen and Bollerslev (1994) utilized a flexible Fourier form that allows for a relationship between the

[4] At high intraday frequencies, any effect of changes in the interest rate differential could be interpreted as 'news effects' rather than yield effects.

level of daily volatility and the shape of the intraday seasonal pattern.

Of course, alternative methods of accounting for intraday seasonality will result in different amounts of remaining structure. Andersen and Bollerslev (1994), Ghose and Kroner (1996) and Guillaume, Pictet, and Dacorogna (1995) all concluded that accounting for intraday seasonals in volatility does not completely resolve apparent inconsistencies between estimates from models using lower frequency (interday) versus higher frequency (intraday) data, as stipulated by theoretical temporal aggregation properties of a correctly specified GARCH(1,1) process (Drost and Nijman 1993). As a result, Ghose and Kroner (1996) proposed a 'signal plus noise' model and also estimated the two-component GARCH model of Engle and Lee (1993) using various frequencies of FX returns ranging from 15-minutes to 8-hours. Analogous to fractionally-integrated models of volatility (for example, Baillie, Bollerslev, and Mikkelsen (1996)), this approach allowed a separation of components of volatility which may, for example, have different persistence properties.

If the standard GARCH(1,1) model is misspecified, one would not expect the theoretical temporal aggregation conditions to hold. We postulate that augmenting the volatility function by explicitly including information arrival variables, or other information which captures intraday stochastic structure, may improve the specification. As such, our approach is complementary to those using a general time-series approach such as a components or a fractionally-integrated specification of the conditional variance function. We discuss some of the alternative methods of deseasonalization and modelling of the intraday dynamics of FX return volatility further below.

The paper is organised as follows. Section 2 provides some notation. Section 3 describes the data set used in this study, briefly reviews some other recent analyses utilizing the same data, and presents some descriptive statistics. Section 4 sets out our model and presents the results. Finally, section 5 concludes the paper.

2 Notation

We follow the notation from the survey by Guillaume et al. (1994) and define the logarithmic *price* of a foreign currency as

$$x(t_j) \equiv \frac{\log p_{\text{bid}}(t_j) + \log p_{\text{ask}}(t_j)}{2},$$

in which t_j is the sequence of tick recording times which are unequally spaced. Many traditional statistical or econometric analyses require data which are equally spaced in time, denoted by t_i. To obtain prices at time t_i, one can either interpolate from adjacent ticks, that is, $t_{j-1} < t_i < t_j$, as advocated by Müller et al. (1990), or, alternatively, use the most recent price $x(t_j)$, as in Wasserfallen and Zimmermann (1985).

Similarly, we compute $i(t_i)$, an *interest rate* at time t_i, as

$$i(t_i) \equiv \frac{i_{\text{bid}}(t_i) + i_{\text{ask}}(t_i)}{2 \cdot 100}$$

where the division by 100 is needed to obtain net interest rates, that is values of, say, 0.0345 rather than 3.45.

The *change of log price* at time t_i, $r(t_i)$, is defined as

$$r(t_i) \equiv r(\Delta t; t_i) \equiv x(t_i) - x(t_i - \Delta t)$$

in which Δt is a fixed time interval such as 10 minutes or 1 hour. We will refer to the change in log price as the FX *return* even though the yield (interest rate) is not included. Uncovered interest rate parity relates the expected change in a bilateral spot exchange rate to the domestic-foreign interest rate differential. Although we have intraday interest rate data, a differential yield on the two currencies is unlikely to explain changes in the spot rate over 10 minute or hourly intervals. Nevertheless, as reported below, we test whether or not the relevant interest rates are important at such high frequencies. For risk averse agents, the expected change in the spot exchange rate may also be a function of a time-varying risk premium. Again, it would be very difficult to measure risk premiums at such high frequencies.

The 'news theory' of exchange rate determination relates innovations in spot exchange rates to 'news'. Therefore, it is possible that a 'news' variable may help explain the intraday conditional variance of changes in exchange rates. We use $a(t_i)$ to denote the number of headline news in the period preceding time t_i; details on how we computed these counts are given in section 3.3.

3 Data Issues

The 'hfdf93' data set from Olsen & Associates spans one year from October 1992 to September 1993 and is comprised of every bid and ask quote for the exchange rates between the US

Dollar, the Japanese Yen, and the German Mark. It is important to point out that these quotes are indicative, rather than firm; but actual transactions data are not available in these markets. Also supplied with the 'hfdf93' data set are the three-month interest rates for these three currencies as well as the money market headline news from the Reuters screen AAMM. This data set is extremely rich in detail and offers numerous challenges to econometricians. The mere volume is impressive: there are 1 472 241 records associated with the DEM-USD exchange rate.

As table 1 shows, each record comprises: the date in Greenwhich Mean Time (GMT) time up to the second, the bid and ask quotes, three codes — one each for the country and the city from which the quotes originated and one for the bank that made these quotes, and finally an indicator variable that signals whether the observation is to be taken as 'good' or 'bad'. The filter algorithm that is used to decide whether an observation is considered 'good' or 'bad' in described in Dacorogna et al. (1993); they find that 0.36% of the observations in the raw data to be in the latter category. The news headlines file contains 105 065 observations. Table 2 shows the first fifteen news headlines for the sample.

3.1 Seasonality

Currencies are traded 24 hours a day, 7 days a week, and are simultaneously traded in different regions of the world. For FX trading activity, one can divide the globe into three main geographical areas: East Asia with Tokyo as the most important place of trading, Europe with highest volume in London, and North America with New York as its major trading centre. Because these markets are in different time zones, when viewed on an intraday basis there are periods of little activity (the Japanese lunch break, 3:00-5:00 GMT, and when Europe and North America are inactive) and high activity (when both Europe and North America are active). Similarly, there very little activity during the weekend (21:00 GMT Fridays to 21:00 GMT Sundays).

This seasonality leaves a very strong pattern in the data as can be seen from figure 1 which corresponds to figure 4 in Guillaume et al. (1994). This figure shows intraday and intraweek histograms of the mean absolute values of the changes in logarithms of the foreign exchange prices, the mean relative spread between the bid and ask quotes and the average number of ticks per time interval.

Preceding the discussion on the computation of the news count variables in section 3.3 below, we also display histograms that show daily and weekly activity of these news variables; they are displayed in figures 2 to 4. They show the total number of news per sixty minute

Table 1: The Mark / Dollar exchange rate data from Olsen & Associates

1992-10-01	00:00:14	1.4116	1.4121	392	01	0058	1
1992-10-01	00:00:54	1.4108	1.4118	036	02	0130	1
1992-10-01	00:01:00	1.4110	1.4120	392	01	0452	1
1992-10-01	00:01:18	1.4115	1.4120	392	01	0041	1
1992-10-01	00:01:24	1.4107	1.4117	036	02	0130	1
1992-10-01	00:01:30	1.4115	1.4125	036	02	0089	1
1992-10-01	00:01:36	1.4113	1.4123	392	01	0041	1
1992-10-01	00:01:42	1.4110	1.4120	392	01	0053	1
1992-10-01	00:01:54	1.4118	1.4128	392	01	0041	1
1992-10-01	00:02:06	1.4113	1.4123	344	01	0055	1
1992-10-01	00:02:18	1.4115	1.4130	702	01	0126	1
1992-10-01	00:02:24	1.4105	1.4115	036	02	0130	1
1992-10-01	00:02:30	1.4115	1.4125	344	01	0089	1
1992-10-01	00:02:48	1.4115	1.4125	392	01	0041	1
1992-10-01	00:02:54	1.4116	1.4121	036	03	0070	1

Table 2: The market headline news from Olsen & Associates

1992-10-01	00:06:16	"AUSTRALIA BOND MARKET OPENS WEAKER AHEAD OF TENDER"
1992-10-01	00:08:42	"SENATE CALLS FOR MILITARY AID FOR BOSNIA"
1992-10-01	00:17:46	"U.S. DEBT FUTURES MIXED, CURRENCIES FIRM AT NIGHT"
1992-10-01	00:18:16	"NZ DEBT MARKETS STEADY PRE-TENDER ON HIGHER TWI"
1992-10-01	00:24:16	"DOLLAR OPENS ALMOST UNCHANGED IN RANGE-BOUND TOKYO"
1992-10-01	00:25:44	"BOJ INJECTS NET 1.15 TRILLION YEN INTO MARKET"
1992-10-01	01:01:50	"FINANCIAL TIMES CALLS FOR LAMONT'S RESIGNATION"
1992-10-01	01:14:22	"BOJ SEEKS OFFERS FOR 300 BILLION YEN BILLS"
1992-10-01	01:14:58	"BANK OF JAPAN OFFERS TO BUY 250 BLN YEN IN BONDS"
1992-10-01	01:31:18	"KEY 30-YR U.S. T-BOND SIGHTLY UP IN MORNING TOKYO"
1992-10-01	01:34:44	"POUND FALLS SHARPLY ON LAMONT RESIGNATION RUMOURS"
1992-10-01	01:55:14	"FOCUS--ITALY RACING AGAINST TIME TO APPROVE BUDGET"
1992-10-01	02:06:40	"GOLD OPENS SLIGHTLY FIRMER IN H.K. ON EARLY BUYING"
1992-10-01	02:07:18	"TOKYO OVERNIGHT RATE DROPS, LONGER RATES STEADY"
1992-10-01	02:09:02	"TAIWAN DOLLAR LITTLE CHANGED IN EARLY TRADE"
1992-10-01	02:13:30	"CHINA CONGRESS SEEN OPTING FOR CAPITALIST MARKETS"
1992-10-01	02:30:06	"AUSTRALIAN DOLLAR WEAKER AT NOON, SEEN SHAKY"
1992-10-01	02:33:34	"U.S. HOUSE, SENATE NEAR COMPROMISE ON DEFENCE BILL"

interval over the year, grouped into daily and weekly histograms. The differences among the three series for the DEM, USD and ALL that are graphed by these histograms is quite striking.

Starting with the 'DEM' newscount for the Deutsche Mark in figure 2, we recognise a distinct daily pattern that is clearly linked to market activity in Germany. There are very few news items before 07:00h (in GMT time) and after 17:00h. The highest number of news events occurs between 07:00h and 08:00h, activity then declines, takes a significant dip between 13:00h and 15:00h and then rebounds for two more hours. The weekly pattern is also very interesting as it shows a peak in the number of news items on Wednesdays, followed by slightly fewer items on Thursday. Monday and Friday show less activity than Tuesday, and there is almost none on weekends.

The 'USD' series for news related to the United States in figure 3 shows a very different picture. Around the clock, we see news arrival of over 600 headlines in each sixty minute interval (with the exception of the 05:00h to 06:00h interval where it drops to 468), and several distinct peaks at 03:00h, 07:00h to 09:00h, 13:00h to 18:00h and 20:00h to 22:00h. The weekly histogram confirms this analysis. There is clearly less of a 'daily' structure as for the DEM news. The biggest number of news events also occurs at the end of the week on Thursdays and Fridays, rather than in the middle. We can relate the relative clustering of news of Thursdays and Fridays to the higher number of official news announcements, as for example pointed out by Ederington and Lee (1993). However, the around-the-clock news arrival differentiates the data that we use here from the data sets use by either Ederington and Lee (1993) or Andersen and Bollerslev (1996) who focus on scheduled news announcements. These announcements occur mostly during the morning in Eastern time, reflected in the peak at the 14:00h to 15:00h GMT interval.

Finally, for all the news, we see from figure 4 that most news headlines are recorded during the European market hours between 07:00h and 17:00h. During the other hours of the day, the frequency drops to about one third of the peak value. As for the USD subset of news, the biggest number of news headlines are seen on Thursdays, the distribution over the week is fairly regular.

As mentioned above, several other papers provide detailed descriptive statistics for these data. For example, Ghysels, Gourieroux, and Jasiak (1995) analyze both real time, *i.e.* tick-by-tick, data and data sampled at fixed intervals of 20 minutes. They compare the distributions of $\Delta \log p_{\mathrm{bid}}$, $\Delta \log p_{\mathrm{ask}}$, and Δx or r; and, at least for the DEM-USD series, conclude that the general tendency of the market is well-approximated by the logarithmic

Figure 1: Graphical Evidence of Seasonality in FX Rates

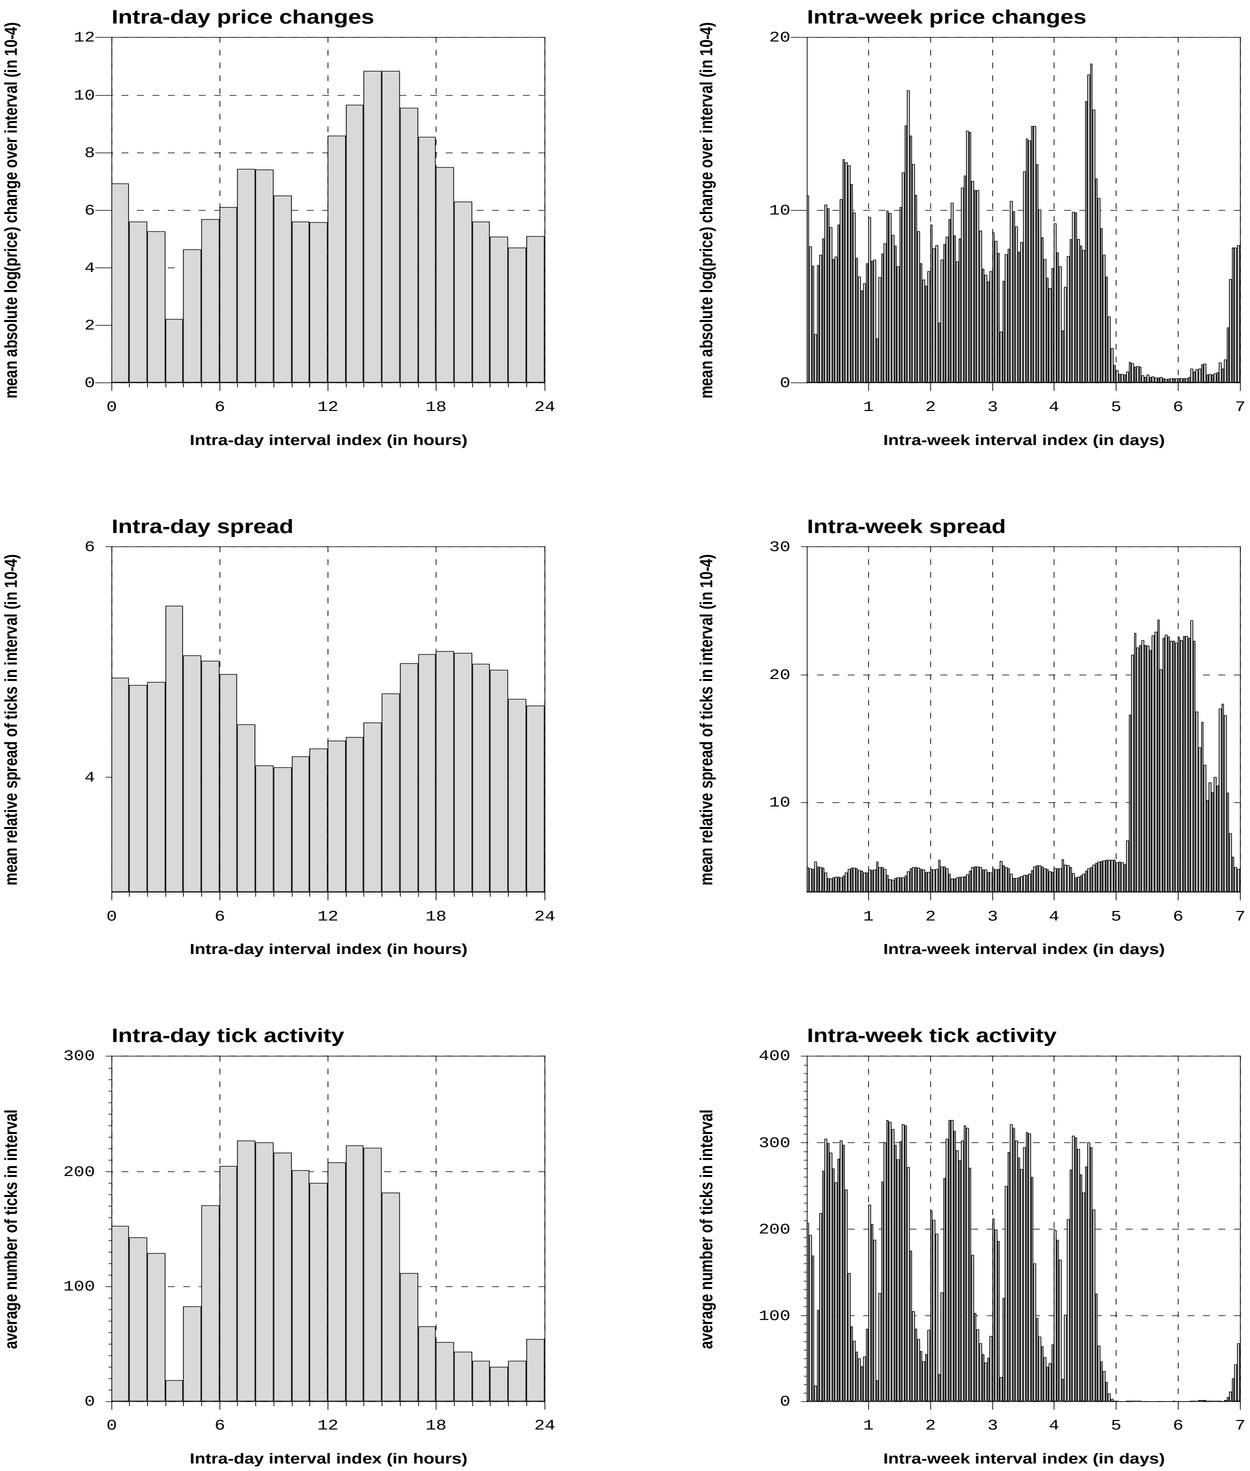

Note: We are grateful to Michel Dacorogna of Olsen & Associates for providing us with this figure which corresponds to figure 4 on page 15 of Guillaume et al. (1994). Note that the scaling factor (10-4) in the first four figures should be read as 10^{-4}.

Figure 2: Graphical Evidence of Seasonality in DEM News Arrival

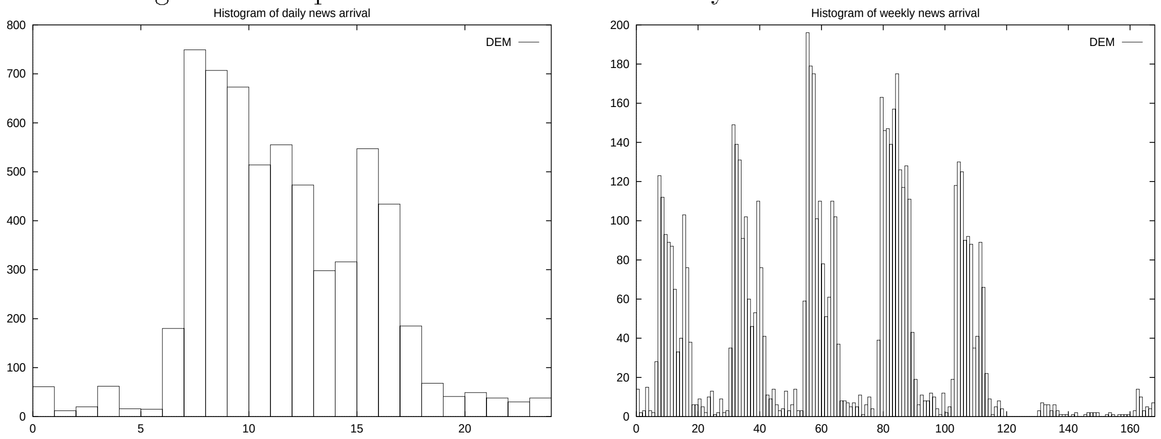

Figure 3: Graphical Evidence of Seasonality in USD News Arrival

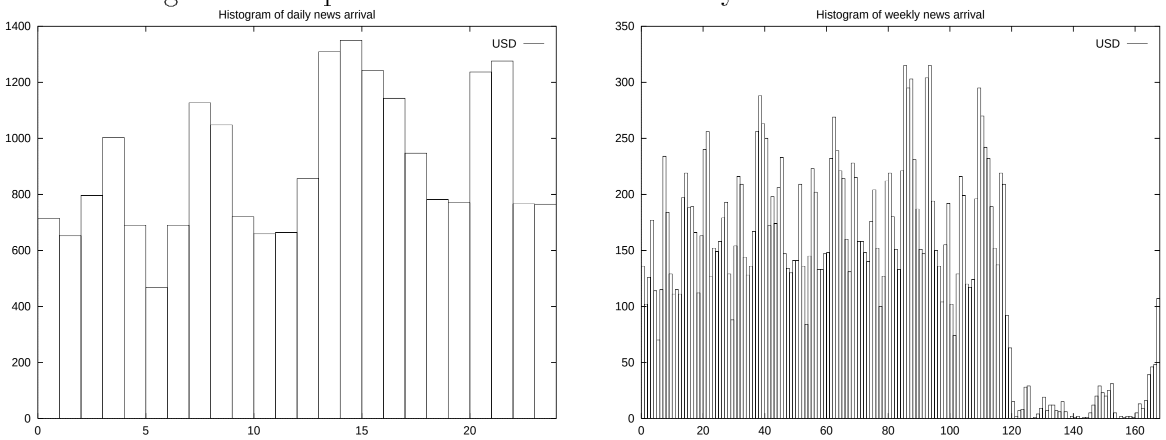

Figure 4: Graphical Evidence of Seasonality in ALL News Arrival

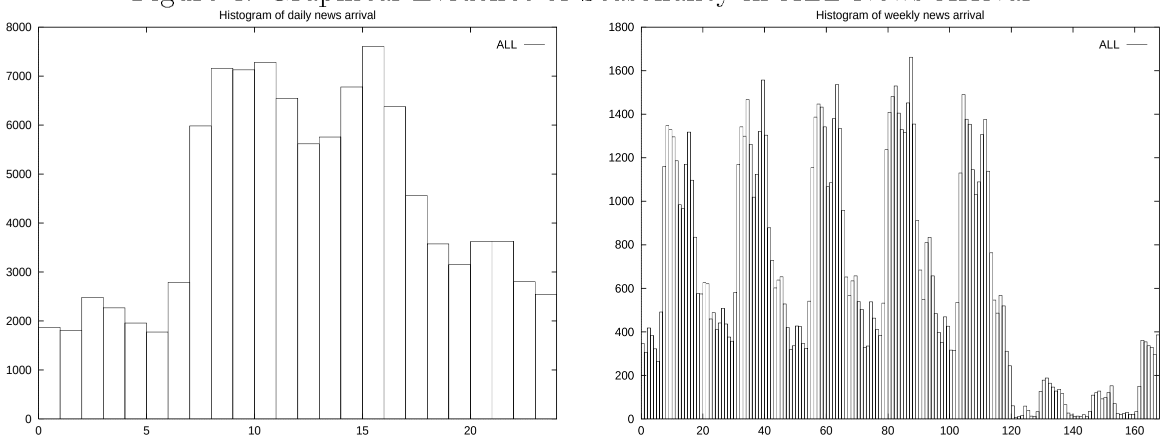

middle price increments or returns, $\Delta x(t_i) = r(t_i)$, associated with the 20 minute time interval. They also compare different measures of market activity as quote arrivals, bid-ask spreads, and absolute returns. They also note that trading volume, which was used as the proxy for market activity in the stock market studies by Clark (1973), Tauchen and Pitts (1983) and Ghysels and Jasiak (1995), is not readily available for foreign exchange markets.

Andersen and Bollerslev (1994) provide detailed summary statistics for 5-minute returns and absolute returns. Intraday heteroskedasticity patterns are illustrated using a plot of the average absolute returns for each of the 288 intervals associated with the 24-hour trading day. They conclude, confirming results reported by Wasserfallen (1989), Müller et al. (1990) and Baillie and Bollerslev (1991), that the pattern is 'closely linked to the daily cycle of market activity in the various financial centers around the globe'. In particular, the volatility starts out relatively high and declines slowly to around 3:00 GMT, drops sharply during lunch hour in the Tokyo and Hong Kong markets, recovers and then increases with the opening of European markets around 7:00 GMT, declines slowly until the European lunch hour, increases sharply during the overlap of European trading and the opening of North American markets (13:00 GMT), and declines after the close of European markets until Asian markets re-open.

3.2 Deseasonalization

As discussed in section 1 above, we are interested in analyzing the potential effect of economic fundamentals on the intraday dynamics of FX returns. Based on the above evidence of a systematic intraday volatility pattern, and on the results in Andersen and Bollerslev (1994) which demonstrate the pervasive effect that such seasonals have on the dynamic or time-series properties of intraday returns, it is important to account for the intraday seasonals.

To deseasonalize volatility, Dacorogna et al. (1993) proposed a time transformation to what they called a ϑ-time scale. This transformation was based on an empirical observation that was first reported by Müller et al. (1990) under the name 'empirical scaling law'. It relates the average of the unconditional volatility, measured as the absolute value of the return $r(t_i)$, over a time interval to the size of the time interval:

$$E|r(t_i)| = \left(\frac{\Delta t_i}{\Delta T}\right)^{1/E}$$

where the drift exponent $1/E$ is an estimated constant that Müller et al. (1990) find to be similar across different currencies and ΔT is a time constant that depends on the currency.

In order to compute the ϑ-time scale, Dacorogna et al. (1993) apply the inverse of this scaling law to the hourly average volatility, computed for every hour of the week, to obtain what they call 'activity statistics' as

$$a_{\text{stat},i} \equiv \frac{\Delta T}{\Delta t} \overline{|r(t_i)|}^E$$

where Δt equals one hour, and the index i refers to the hour of the week, that is $i \in [1, 2, \ldots, 168]$. Next, an activity function $a(t)$ is fitted to the $a_{\text{stat},i}$. Moreover, the activity function is divided into the three principal geographic regions East Asia, Europe and North America. Each of these three regions (indexed by k) is described by an activity variable $a_{0,k}$ that corresponds to the 'base level' during closing hours and an activity variable $a_{1,k(t)}$ which describes the activity when this market is open. Finally, ϑ-time is the time integral of global activity:

$$\vartheta(t) \equiv a_0(t - t_0) + \sum_{k=1}^{3} \int_{t_0}^{t} a_k(t')\mathrm{d}t'$$

The activity variable is normalized such that ϑ-time can be measured in the same units as physical time: an entire week in ϑ-time corresponds approximately to one week in physical time (normalization is done over four years to account for leap years).

In effect, the ϑ-time deseasonalizes volatility by expanding periods with high average volatility (high average market activity due to number of active markets and the market-specific seasonal pattern of intraday activity) and contracting periods with low average volatility (average market activity). This market-activity time deformation is reminiscent of, but more general than, the business-time scale proposed by Stock (1988).

Analogously, building on the subordinated stochastic process representation proposed by Clark (1973) which links volatility to the intensity of trading (the directing process), Ghysels, Gourieroux, and Jasiak (1995) propose a stochastic volatility model for FX returns which combines features of average *and* conditional market activity. In our application, we model the conditional volatility remaining after the intraday seasonal (unconditional) heteroskedasticity has been removed by the change in the time scale.

Other methods of accounting for or modelling the intraday seasonal have been proposed. For example, Andersen and Bollerslev (1994) use a flexible Fourier form which allows the shape of the seasonal pattern to depend on the level of the daily volatility. They suggest that this extension may be more important for equity returns than for those associated with

FX trading.

3.3 Data Extraction

While it might be possible, albeit rather expensive in terms of computer memory and execution time, to work with all $1,472,241$ observations on the DEM-USD exchange rate data, it also turns out that there are good reasons *not* to work with that many observations. Firstly, the signal-to-noise ratio is rather low using the tick-by-tick frequency. Secondly, this study analyses potential links between changes in exchange rates and the frequency of news (as well as the interest rate differential) of which we have fewer observations.

The ϑ-time scale observations on the DEM-USD exchange rate were provided by Olsen & Associates. We matched the time intervals between two successive observations in ϑ-time to the interest rate and news data. For every data point in ϑ-time on the exchange rate, we extracted the matching interest rates at that time point from the two files for German and American ninety-day deposit rates.

Similarly, the news counts were computed by counting the number of news items in the interval preceding the given reference time from the exchange rate data point in ϑ-time. We computed three different news counts: 'ALL', 'DEM' and 'USD'. The first, 'ALL', simply counts the number of news headlines from the Reuters AAMM screen in the interval, that is from the current datapoint in the ϑ-time resolution to the preceding one. News counts denoted by the three-letter ISO symbol for a given currency, that is DEM or USD, indicate how many headlines contained text relevant to the respective currency.

As Andersen and Bollerslev (1996, p. 17) write in context of their modelisation with macroeconomic announcements: *'Modeling volatility in the wake of news arrivals poses a problem in the regression context. We expect the induced volatility to depend on the degree of surprise, or innovation, in the announcement. However, we do not have reliable information on this dimension for most of the events.* Our study of news encompasses macroeconomic announcements by government agencies and includes every possible headline from the Reuters screen. So in order to avoid qualitative judgement about 'good' or 'bad' news, we only note whether or not there is a news headline. Or in other words, we restrict the analysis to the counting of news headlines in a given period of time. While counting the news items may well be an imperfect treatment of the information available to the market participants, it serves as an approximation. In order to count news headlines for each currency, we define a set of matching keywords. Using this set of keywords, we examine all of the 105 065 news headlines one by one and examine whether any of the matching keywords are included. But we retain a headline that matches one of the keywords only if none of the excluding keywords is present. This second set of excluding keywords is needed as we would otherwise record spurious matches. Consider the fourth news headline in our data set that is displayed in table 2 and which contains the word 'market'. Using 'mark' as a matching keyword for the German currency, we would record a match here. But because 'market' is an excluding keyword, we do not retain this headline. Similarly, when we search for 'dollar' we have to ensure that no news item that refers to the Australian or Canadian currency is mistakenly retained as a USD news item. Specifically, for the US Dollar, we use the terms 'U.S.', 'DOLLAR' or 'FED' as matching keywords and the terms 'CANAD', 'TAIWAN', 'AUSTRAL' or 'NZ DOLL' as excluding keywords. So if a given news headline contains the string 'DOLLAR' while it does not contain any of the excluding terms, we count it. For the German Mark, search terms are 'BUBA', 'BUNDESBANK', 'MARK' and 'BONN', while terms 'MARKET', 'MARKED' or 'MARKKA' form the exclusion set. Using these sets, we first construct two new news files that refer exclusively to USD or German news, respectively. We then match these headlines to the ϑ-time scale data points and count the number of news events between data points.

3.4 Statistical Properties of the Data

Table 3 shows the first four moments for the differences in log prices, or returns, associated with the DEM-USD spot exchange rate. These moments are reported for both the ϑ-time scale and the physical-time scale at a time resolution of sixty minutes. The data in ϑ-time are used as the dependent variable in section 4 below. The transformation from physical time into the ϑ-time scale does not change the mean of the logarithmic price changes. However, the skewness and excess kurtosis both decrease for the ϑ-time scale relative to the physical-time scale. For example, the coefficient of excess kurtosis is reduced substantially from 11.525 to 4.052.

In table 4, we display Ljung-Box-Pierce portmanteau statistics for changes in the foreign currency prices computed both in physical and in ϑ-time. Since non-random structure in higher-order moments will affect the asymptotic distribution of these tests, the critical values are only meant to be indicative. However, Monte Carlo experiments by Bollerslev and Mikkelsen (1996) provide evidence that these tests are quite well-behaved for standardized residuals; see below for a further discussion of this. Table 4 shows that there is a high degree of persistence associated with the two alternative measures of volatility, *i.e* absolute values and squares of log price changes, and in some cases this persistence actually *increases* with

Table 3: Moment Statistics for Data in Physical and Theta-Time Scales

	$100 \cdot r$	$100 \cdot r$	$a_{\text{ALL},t-1}$	$i_{\text{diff},t-1}$
	physical time	theta time		
observations	8759	8735	8735	8735
mean	0.002	0.002	0.120	0.0455
variance	0.016	0.016	0.008	0.0000
skewness	0.452	0.261	3.816	-0.3661
kurtosis	11.525	4.052	27.978	-0.9305

Note: As indicated in section 2, r denotes the logarithmic middle price increments.

Table 4: Portmanteau Statistics in Physical and Theta-Time Scales

	$100 \cdot r$	$\|100 \cdot r\|$	$(100 \cdot r)^2$
		physical-time scale	
Q_2	10.98	2042.96	488.15
Q_{24}	76.37	6365.28	756.35
Q_{168}	290.11	14165.23	1671.31
		physical-time scale: no weekends	
Q_2	6.14	555.56	261.57
Q_{24}	58.93	1206.89	380.99
Q_{168}	267.52	4630.57	1317.36
		theta-time scale	
Q_2	22.51	383.44	310.72
Q_{24}	60.58	1453.24	776.01
Q_{168}	272.73	5099.59	2650.99

Note: The portmanteau statistic Q_p for p lags is, under the null of white noise data, asymptotically distributed as a $\chi^2(p)$. At the 1% level, the unadjusted critical values for 2, 24 and 168 lags are 9.22, 43.00 and 213.57 respectively.

the Olsen & Associates deseasonalization. This increase in persistence conforms with the Andersen and Bollerslev (1994) conclusion concerning the effect of the intraday seasonal on the dynamic properties of the return volatility.

3.5 Unit Root Tests

We test for stationarity of the relevant time series using the standard z' tests of Phillips and Perron (1988). Optimal lag lengths are determined using the AIC2 criterion from Pantula, Gonzales-Farias, and Fuller (1994). These stationarity tests are conducted using three alternative test equations. In the Davidson and MacKinnon (1993) notation, we use specifications with just a constant, denoted by a subscript c, with a constant and a time trend, denoted by a subscript ct and with a constant, a time trend and a squared time trend, denoted by a subscript ctt. Table 5 shows the test statistics, and their P-values, computed for log prices, x, and change in log prices, r, for the DEM-USD spot exchange rate from both the physical-time and the transformed ϑ-time scale. We also report the test statistics for the relevant interest rate series which have been extracted for matching time intervals.

Briefly, we find that the log prices series x might be described by a non-stationary process as we fail to reject the null hypothesis of a unit root for all six test statistics. However, for our dependent variable r, that is the first differences of the log prices, we can very decisively reject the null hypothesis of a unit root. For the interest rate series, we clearly reject the null hypothesis of a unit root for the Mark and the US Dollar as well as for the yield-difference series for the test regression that comprise a trend or a trend and a squared trend. We verified the results of the Phillips-Perron test with the KPSS test of Kwiatkowski et al. (1992) which tests the null hypothesis of stationarity. We very decisively reject this null hypothesis for the log price series x, and can only reject it for the returns series r at the relatively high sizes of 5% and 10%. Finally, there is remarkably little difference between the test statistics for the untransformed 'physical' time and the transformed 'theta' time with both tests.

4 Model Specification and Results

In order to examine the impact of news headlines and interest rates on the first and second conditional moments of the changes in the deseasonalized DEM-USD exchange rate, we use the GARCH(1,1) model of Bollerslev (1986). We augment the basic model by including a set of regressors $\boldsymbol{V}$ in the conditional variance equation. In order to account for serial dependence in the first moment, we also include an MA(1) term in the conditional mean

Table 5: Unit Root Tests

Series	z'_c		z'_{ct}		z'_{ctt}	
	physical-time scale					
x	-10.97	(0.107)	-12.90	(0.270)	-14.37	(0.444)
r	-6494.57	(0.000)	-6448.19	(0.000)	-6440.84	(0.000)
i_{DEM}	-2.18	(0.310)	-99.30	(0.000)	-124.02	(0.000)
i_{USD}	-52.38	(0.000)	-100.41	(0.000)	-114.97	(0.000)
i_{diff}	-5.16	(0.419)	-135.93	(0.000)	-230.77	(0.000)
	theta-time scale					
x	-10.38	(0.12)	-11.89	(0.321)	-13.18	(0.513)
r	-9277.48	(0.000)	-9276.22	(0.000)	-9275.87	(0.000)
i_{DEM}	-2.47	(0.281)	-112.40	(0.000)	-137.33	(0.000)
i_{USD}	-56.45	(0.000)	-106.53	(0.000)	-122.28	(0.000)
i_{diff}	-5.70	(0.371)	-159.59	(0.000)	-262.74	(0.000)

Note: The P-values that are reported in parentheses were calculated using the `urcdist` program by MacKinnon (1996).

equation. Formally, the model is written as:

$$
\begin{aligned}
100 \cdot r_t &= \boldsymbol{X}_{t-1}\boldsymbol{\beta} + \phi\,\varepsilon_{t-1} + \varepsilon_t; \\
\varepsilon_t|\Psi_{t-1} &\sim N(0, h_t); \\
h_t &= \alpha_0 + \alpha_i \varepsilon_{t-1}^2 + \gamma_j h_{t-1} + \boldsymbol{V}_{t-1}\boldsymbol{\delta}.
\end{aligned}
\tag{1}
$$

where $\boldsymbol{V}$ might comprise the same regressors as $\boldsymbol{X}$, with the exception of the column of constants which is already included with the regressor α_0.

As the dependent variable, we use the change in log price of the DEM-USD spot exchange rate. As explanatory variables, we use a constant, the difference between the German and American interest rates and three news-count variables. The first news count is the frequency of news where we count all news headlines. This count variable is denoted by a subscript 'ALL'. Specific news counts for the Mark and the Dollar, computed as described in section 3.3, are denoted by 'DEM' and 'USD', respectively. We scale the news count variables by a factor of 0.01 as is indicated in table 6.

Table 6 presents Quasi-Maximum Likelihood Estimates for four different versions of our augmented GARCH(1,1) model. Standard errors from the estimated Hessian are reported in parentheses and robust standard errors (Bollerslev and Wooldridge 1992) are shown in square brackets. The first version is the standard GARCH(1,1) model augmented with the interest rate differential in the mean (Model 1). Next, we estimated the base model augmented with the 'ALL' news-count variable in the conditional variance function (Model 2). Model 3 adds the interest rate differential to the variance function as well. Finally, in addition to the interest rate differential, we include currency-specific news-count variables in the variance (Model 4).

Table 7 summarizes some analyses of the standardized residuals from the models reported in Table 6. In particular, we report descriptive statistics for the first four moments of the standardized residuals. We also report Ljung-Box-Pierce portmanteau statistics for autocorrelation up to 2 lags (2 hours), 24 lags (1 day), and 168 lags (1 week). These portmanteau statistics are computed for the standardized residuals in order to detect remaining predictable structure in the conditional mean, and also computed for the absolute value of the standardized residuals and the square of the standardized residuals to determine if there is neglected conditional heteroskedasticity. As these portmanteau test statistics are not robust to ARCH effects, we must not interpret the associated P-values literally. However, Bollerslev and Mikkelsen (1996) provide some simulation evidence that these tests are fairly accurate for the standardized residuals. For absolute standardized residuals and squared

Table 6: GARCH ML Estimation Results

	Model 1	Model 2	Model 3	Model 4
	Mean Parameters			
c	-0.01316	-0.01311	-0.01186	-0.01283
	(0.00758)	(0.00739)	(0.00704)	(0.00715)
	[0.00760]	[0.00740]	[0.00730]	[0.00756]
ϕ	-0.05668	-0.06118	-0.06303	-0.06382
	(0.01125)	(0.01172)	(0.01176)	(0.01191)
	[0.01272]	[0.01261]	[0.01254]	[0.01258]
$i_{\text{diff},t-1}$	0.31975	0.32632	0.29921	0.31444
	(0.16710)	(0.16240)	(0.15730)	(0.15440)
	[0.16889]	[0.16626]	[0.16387]	[0.16905]
	GARCH(1,1) Parameters			
α_0	0.00026	0.00005	-0.00138	-0.00072
	(0.00004)	(0.00011)	(0.00022)	(0.00027)
	[0.00008]	[0.00025]	[0.00034]	[0.00047]
α_1	0.03414	0.07556	0.08103	0.09700
	(0.00382)	(0.00635)	(0.00699)	(0.00775)
	[0.00484]	[0.01109]	[0.01164]	[0.01249]
γ_1	0.94904	0.86566	0.83669	0.79834
	(0.00512)	(0.01054)	(0.01272)	(0.01444)
	[0.00823]	[0.02029]	[0.02150]	[0.02269]
	Additional Variance Parameters			
$i_{\text{diff},t-1}$			0.03206	0.03554
			(0.00509)	(0.00631)
			[0.00865]	[0.00979]
$0.01 \cdot a_{\text{ALL},t-1}$		0.00751	0.01045	
		(0.00094)	(0.00110)	
		[0.00193]	[0.00195]	
$0.01 \cdot a_{\text{DEM},t-1}$				0.00521
				(0.00808)
				[0.01285]
$0.01 \cdot a_{\text{USD},t-1}$				0.01756
				(0.00416)
				[0.00683]
	Loglikelihood			
	6003.62	6031.80	6059.43	6028.75

Note: Standard errors from the estimated Hessian are in parentheses and robust standard errors (Bollerslev and Wooldridge 1992) are in square brackets.

standardized residuals, they recommend a simple adjustment to the degrees of freedom, that is, subtracting the number of estimated ARCH parameters from the number of lags.

Several conclusions can be drawn from the results reported in tables 6 and 7. Firstly, and as expected when modelling returns on financial assets, we note that the intercept is insignificant at conventional test sizes for the conditional mean of all four models reported in table 6. However, the estimate for the MA(1) coefficient ϕ, where we find a small negative value, is statistically significant in all four models. This is consistent with earlier results by Baillie and Bollerslev (1991) who use exchange rate bid prices, also with a frequency of sixty minutes. Recall that we use the midpoint of the log bid and log ask quotes. Therefore, as suggested by the market micro-structure literature, some of our negative serial dependence might be due to spread positioning by dealers. That is, as summarized by Andersen and Bollerslev (1994), dealers may position their quotes asymmetrically about the perceived 'true market price' in order to manage their inventory positions. This strategy can generate mean reversion in the midpoint of bid and ask quotes which is analogous to the 'bid-ask bounce' observed on organized exchanges.

The interest rate differential implied by economic theory is only marginally significant; the robust t-statistic is 1.89 for the mean equation. The fact that the interest rate differential has low explanatory power for the conditional mean was predicted at the end of section 2 above based on the premise that it is unlikely that a differential yield on ninety-day deposits in the two currencies would explain changes in the spot exchange rate over such short intervals as one hour.

However, adding the time $t-1$ ALL news-count variable to the conditional variance function significantly improves the loglikelihood by 28.18 (from 6003.62 to 6031.80). Furthermore, adding the interest rate differential to the conditional variance function (model 3) increases the loglikelihood still further to 6059.43. Changes in interest rates are likely to be related to news concerning monetary policy. In sum, both the frequency of ALL news and the interest rate differential contribute to explaining the conditional variance of the bilateral DEM-USD spot exchange rate. In particular, our results show that the conditional volatility of the spot exchange rate responds positively to the frequency of news. A comparison of the loglikelihoods associated with models 3 and 4 indicates that the frequency of ALL news contributes explanatory power in addition to that provided by the frequency of local news which we measured using the subsets of news that we identified with certain keywords associated with the DEM or the USD.

We also estimated model 3 at higher frequencies. For example, for 20-minute returns the

Table 7: GARCH ML Estimation Residual-Based Diagnostics

	Model 1	Model 2	Model 3	Model 4
	Descriptive statistics for stand. residuals			
mean	0.000	−0.001	−0.001	0.001
variance	1.004	1.003	1.003	1.003
skewness	0.111	0.149	0.137	0.110
excess kurtosis	3.412	3.223	2.958	3.041
	Q_p statistics for stand. residuals			
Q_2	3.38	2.58	2.37	1.99
Q_{24}	31.48	28.05	28.56	29.59
Q_{168}	230.60	229.25	232.64	237.31
	Q_p^a statistics for absolute stand. residuals			
Q_2^a	61.30	16.90	12.50	7.71
Q_{24}^a	84.23	46.13	47.93	48.64
Q_{168}^a	328.68	372.28	395.83	443.96
	Q_p^2 statistics for squared stand. residuals			
Q_2^2	38.82	10.66	8.03	4.20
Q_{24}^2	65.06	36.04	31.64	28.60
Q_{168}^2	211.53	252.82	254.25	255.64

Note: The portmanteau statistic Q_p for p lags is, under the null of white noise data, asymptotically distributed as a $\chi^2(p)$. At the 1% level, the unadjusted critical values for 2, 24 and 168 lags are 9.22, 43.00 and 213.57 respectively.

scaled 'ALL' news-count variable in the conditional variance had a coefficient of .00441 with standard and robust t-statistics of .00036 and .00082 respectively.

The result that spot exchange rates are more volatile during periods for which there is a lot of economic news accords with casual observations in the industry. Furthermore, news arrival has often been the informal explanation used for the observed volatility clustering of financial data.

Note that including the news variables in the GARCH(1,1) variance function results in a covariance-stationary, as opposed to near-integrated, conditional volatility process: the sum of α_1 and γ_1 falls from around 0.98 to 0.92. This result has the appealing feature that volatility forecasts will exhibit mean reversion. That is, an innovation to exchange rate changes will not affect forecasts of conditional variance for all future periods – the impulse response weights will eventually die-out.[5]

The portmanteau Q statistic for standardised residuals shows that there is virtually no remaining cumulative persistence in the conditional mean over the bi-hourly (2 lags) and daily horizons (24 lags). All models have test statistics that are lower than the 1% critical value. Models which include the news count variables have slightly lower portmanteau statistics than model 1. However, persistence remains over a longer period of one week (168 lags) with test statistics that exceed the 1% critical value.

With respect to remaining conditional heteroskedasticity, we also compute the portmanteau statistics associated with the absolute and squared values of the standardized residuals. For up to 2 lags (2 hours) and 24 lags (1 day), the Q statistics decrease dramatically for the models with news variables included. For example, for 24 lags of the squared residuals, the Q^2_{24} statistic decreases from 65.06 to 31.64 which is lower than the critical 5% value of the χ^2 distribution with 24 degrees of freedom. We can conclude that the frequency of news helps explain a part of the conditional volatility of the DEM-USD spot exchange rate.

Note that at 168 lags (1 week), there is still significant persistence in volatility which is not captured by any of the GARCH(1,1) specifications. Whether this remaining structure or long memory would be better captured by extending the number of lags in the GARCH specification, by conditioning on additional information, or by modelling the volatility as a fractionally-integrated process as in Baillie, Bollerslev, and Mikkelsen (1996) or as a components model as in Ghose and Kroner (1996), remains to be seen.

[5]Note that volatility persistence is difficult to characterize in nonlinear models (see for example Nelson (1990) and Bollerslev and Engle (1993)). In addition to the parameter estimates for α_1 and γ_1, Andersen and Bollerslev (1994) provide three other measures for persistence: the mean lag, the median lag, and the 'half-life' of the volatility process.

5 Conclusion

Using the 'hfdf93' data set from Olsen & Associates, we constructed activity count variables from the Reuters news headlines. Simple keyword matches permitted us to count news that pertain to either the German Mark or the US Dollar, and we also counted global news. Using a resolution of one hour in the ϑ-time scale, we fitted a standard GARCH(1,1) model and included the news counts as well as the relevant interest rate or yield differential. As expected, our results showed that the interest rates (yields) do not have a lot of explanatory power at this high frequency. However, both the frequency of news and the interest rate differential were seen to be significant in the conditional variance equation. In particular, at a frequency of one hour on the ϑ-time scale, more global news increases the conditional volatility of the DEM-USD spot exchange rate.

In sum, our paper contributes to the rapidly growing body of empirical research which uses intraday financial data by providing time-series estimates of the impacts of news headlines on the intraday volatility of the DEM-USD spot exchange rate.

Adding the news variables to the conditional variance function resulted in a covariance-stationary, as opposed to near-integrated, GARCH(1,1) volatility function. This result has the appealing feature that an innovation to exchange rate changes will no longer affect forecasts of conditional variance for all future periods. In this regard, our paper has confirmed, and extended, the findings reported by Goodhart, Hall, Henry, and Pesaran (1993) who used a smaller dataset and examined only two discrete news events.

Much work remains to be done. Although we have attempted to avoid the impact of repetitive intraday seasonals on the time-series dynamics in order to focus on modelling conditional heteroskedasticity due to economic news and changes in yield, other methods of accounting for the intraday seasonals could be tried as well. There is some evidence of additional structure, particularly with respect to remaining persistence in volatility at longer lags, which has not been captured by our simple specification. Bollerslev and Domowitz (1993) showed that bid-ask spreads had a positive effect on return volatility in the DEM-USD interbank market. We plan to investigate the impact of bid-ask spread dynamics in our model as well. Also, we plan to analyse other functions of the news headlines, including dynamic effects as in Engle, Ito, and Lin (1990), and to conduct further investigations of our specification including its persistence and temporal aggregation properties. We hope to address all of these issues in future research.

References

Andersen, Torben G. and Tim Bollerslev (1994). Intraday seasonality and volatility persistence in foreign exchange and equity markets. Working Paper No. 186, Department of Finance, Kellogg Graduate School of Management, Northwestern University.

Andersen, Torben G. and Tim Bollerslev (1996). DM-Dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. Working Paper No. 217, Department of Finance, Kellogg Graduate School of Management, Northwestern University.

Baestens, Dirk J. E. and Willem M. Van den Bergh (1996). Money market headline news flashes, effective news and the dem/usd swap rate: An intraday analysis in operational time. In *Forecasting Financial Markets: New Advances for Exchange Rates, Interest Rates and Asset Management*, London. Third International Conference sponsored by Chemical Bank and Imperial College.

Baillie, Richard T. and Tim Bollerslev (1991). Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies 58(3)*, 565–585.

Baillie, Richard T., Tim Bollerslev, and Hans Ole Æ. Mikkelsen (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 74(1)*, 3–30.

Bollerslev, Tim (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics 31(3)*, 307–327.

Bollerslev, Tim and Ian Domowitz (1993). Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance 48(4)*, 1421–1443.

Bollerslev, Tim and Robert F. Engle (1993). Commom persistence in conditional variance. *Econometrica 61(1)*, 167–186.

Bollerslev, Tim and Hans Ole Æ. Mikkelsen (1996). Modeling and pricing long-memory in stock market volatility. *Journal of Econometrics 73(1)*, 151–184.

Bollerslev, Tim and Jeffrey M. Wooldridge (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. *Econometric Reviews 11(2)*, 143–172.

Clark, Peter K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica 41(1)*, 135–155.

Dacorogna, Michel M., Ulrich A. Müller, Robert J. Nagler, Richard B. Olsen, and Olivier V. Pictet (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.

Davidson, Russell and James G. MacKinnon (1993). *Estimation and Inference in Econometrics*. New York: Oxford University Press.

DeGennaro, Ramon P. and Ronald E. Shrieves (1995). Public information releases, private information arrival, and volatility in the foreign exchange market. In *First International Conference on High Frequency Data in Finance*, Volume 1, Zürich. Olsen & Associates.

Drost, Feike C. and Theo E. Nijman (1993). Temporal aggregation of garch processes. *Econometrica 61(4)*, 909–927.

Ederington, Louis H. and Jae Ha Lee (1993). How markets process information: News releases and volatility. *Journal of Finance 48(4)*, 1161–1191.

Engle, Robert F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica 50(4)*, 987–1007.

Engle, Robert F., Takatoshi Ito, and Wen-Ling Lin (1990). Meteor showers or heat waves? Heteroskedastic intra day volatility in the foreign exchange market. *Econometrica 58(3)*, 525–542.

Engle, Robert F. and Gary G. J. Lee (1993). A permanent and transitory component model of stock return volatility. Unpublished manuscript, Department of Economics, University of California at San Diego.

Ghose, Devajyoti and Kenneth F. Kroner (1996). Components of volatility in foreign exchange markets: An empirical analysis of high frequency data. Unpublished manuscript, Department of Economics, University of Arizona.

Ghysels, Eric, Christian Gourieroux, and Joanna Jasiak (1995). Trading patterns, time deformation and stochastic volatility in foreign exchange markets. Scientific Series 95s-42, CIRANO, Montreal.

Ghysels, Eric and Joanna Jasiak (1995). Trading patterns, time deformation and stochastic volatility in foreign exchange markets. In *First International Conference on High Frequency Data in Finance*, Volume 1, Zürich. Olsen & Associates.

Goodhart, Charles A. E. and L. Figliuoli (1991). Every minute counts in financial markets. *Journal of International Money and Finance 10(1)*, 23–52.

Goodhart, Charles A. E., Steven G. Hall, S. G. Brian Henry, and Bahram Pesaran (1993). News effects in a high-frequency model of the sterling-dollar exchange rate. *Journal of Applied Econometrics 8(1)*, 1–13.

Guillaume, Dominique M., Michel M. Dacorogna, Rakhal R. Davé, Ulrich M. Müller, Richard B. Olsen, and Olivier V. Pictet (1994). From the bird's eye to the microscope: A survey of new stylized facts of the intra-day foreign exchange markets. Discussion Paper DMG.1994-04-06, Olsen & Associates, Zürich.

Guillaume, Dominique M., Olivier V. Pictet, and Michel M. Dacorogna (1995). On the intra-day performance of GARCH processes. Discussion Paper DMG.1994-07-31, Olsen & Associates, Zürich.

Hakkio, Craig S. and Douglas K. Pearce (1985). The reaction of exchange rates to economic news. *Economic Inquiry 23(4)*, 621–636.

Harvey, Campbell R. and Roger D. Huang (1991). Volatility in the foreign currency futures market. *Review of Financial Studies 4(3)*, 543–569.

Kwiatkowski, Denis, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics 54(1-3)*, 159–178.

Low, Aaron H. W. and Jayaram Muthuswamy (1996). Information flows in high frequency exchange rates. In Christian Dunis (Ed.), *Forecasting Financial Markets: Exchange Rates, Interest Rates, and Asset Management*, Series in Financial Economics and Quantitative Analysis, New York, pp. 3–32. John Wiley & Sons.

MacKinnon, James G. (1996). Numerical distribution functions for unit root and cointegration tests. *Journal of Applied Econometrics 11*, 601–618.

Melvin, Michael and Xixi Yin (1996). Public information arrival, exchange rate volatility, and quote frequency. Unpublished manuscript, Department of Economics, Arizona State University.

Müller, Ulrich A., Michel M. Dacorogna, Richard B. Olsen, Oliver V. Pictet, Matthias Schwarz, and Claude Morgenegg (1990). Statistical study of foreign exchange rates, empirical evidence of a price scaling law, and intraday analysis. *Journal of Banking and Finance 14(6)*, 1189–1208.

Nelson, Daniel B. (1990). Stationarity and persistence in the GARCH(1,1) model. *Econometric Theory 6(3)*, 318–334.

Pantula, Sastry G., Graciela Gonzales-Farias, and Wayne A. Fuller (1994). A comparison of unit-root test criteria. *Journal of Business and Economic Statistics 12(4)*, 449–459.

Phillips, Peter C. B. and Pierre Perron (1988). Testing for a unit root in time series regression. *Biometrika 75(2)*, 335–346.

Stock, James H. (1988). Estimating continous-time processes subject to time deformation: An application to postwar U.S. gnp. *Journal of the American Statistical Association 83(401)*, 77–85.

Tauchen, George E. and Mark Pitts (1983). The price variability-volume relationship on speculative markets. *Econometrica 51(2)*, 485–505.

Wasserfallen, Walter (1989). Flexible exchange rates: A closer look. *Journal of Monetary Economics 23(3)*, 511–521.

Wasserfallen, Walter and Heinz Zimmermann (1985). The behavior of intra-daily exchange rates. *Journal of Banking and Finance 9(1)*, 55–72.

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