ELSEVIER

Journal of Empirical Finance 4 (1997) 115–158

Journal of
EMPIRICAL
FINANCE

Intraday periodicity and volatility persistence in financial markets

Torben G. Andersen [a], Tim Bollerslev [b,c,*]

[a] Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208, USA

[b] Department of Economics, Rouss Hall, University of Virginia, Charlottesville, VA 22901, USA

[c] NBER, Cambridge, MA 02138, USA

Abstract

The pervasive intraday periodicity in the return volatility in foreign exchange and equity markets is shown to have a strong impact on the dynamic properties of high frequency returns. Only by taking account of this strong intraday periodicity is it possible to uncover the complex intraday volatility dynamics that exists both within and across different financial markets. The explicit periodic modeling procedure developed here provides such a framework and thus sets the stage for a formal integration of standard volatility models with market microstructure variables to allow for a more comprehensive empirical investigation of the fundamental determinants behind the volatility clustering phenomenon. © 1997 Elsevier Science B.V.

JEL classification: C14; C22; G14; G15
Keywords: Volatility; Intraday periodicity; Temporal aggregation; ARCH

1. Introduction

It is widely documented that return volatility varies systematically over the trading day and that this pattern is highly correlated with the intraday variation of trading volume and bid–ask spreads. Indeed, these strikingly regular patterns of market activity measures have provided the impetus for much theoretical work. On

* Corresponding author.

PII S0927-5398(97)00004-2

the other hand, the dynamics of the intraday return volatility process is mostly ignored in the empirical market microstructure literature. This is quite surprising given the notion that news arrivals and the resolution of their informational impact are intimately related to the dynamics of the return volatility process [1]. We conjecture that the intraday return dynamics is neglected primarily because standard time series models of volatility have proven inadequate when applied to high frequency returns data. In fact, previous results reported in the literature are often contradictory and generally defy theoretical predictions. Consequently, there is no well established paradigm for intraday volatility modeling, and at present its inclusion in market microstructure research is tenuous.

In this paper we demonstrate that the difficulties encountered by standard volatility models arise largely from the aforementioned systematic patterns in average volatility across the trading day. We further show how practical estimation and extraction of the intraday periodic component of return volatility is both feasible and indispensable for meaningful intraday dynamic analysis. Particular attention is paid to the differing impact of the periodic pattern on the dynamic return features at the various intraday frequencies. To illustrate the range of applicability of the developed procedures, the analysis is conducted in parallel for two different asset classes traded under widely different market structures, namely the over-the-counter foreign exchange interbank market and an organized exchange for futures equity index contracts. Moreover, to bring out the distinct character of the intraday returns process, the findings are contrasted to the corresponding features of interdaily returns series for the identical assets.

The empirical evidence on the properties of average intraday stock returns dates back to, at least, Wood et al. (1985) and Harris (1986a) who document the existence of a distinct U-shaped pattern in return volatility over the trading day i.e. volatility is high at the open and close of trading and low in the middle of the day. The existence of equally pronounced intraday patterns in foreign exchange markets has been demonstrated by Müller et al. (1990) and Baillie and Bollerslev (1991) [2].

Meanwhile, a separate time series oriented literature has modeled the dynamics of the intraday return volatility directly, building on the ARCH methodology of

[1] For example, theoretical work stress issues such as the process of price discovery, the optimal timing of trades designed to limit price impact, the differing price response to public versus private information, the clustering of discretionary liquidity trading and the associated increase in market depth when private information is short-lived and the particular market dynamics associated with periodic market openings and closures.

[2] Empirical work continues to refine and classify the regularities of high frequency returns in this dimension. Recent studies include Barclay et al. (1990) and Harvey and Huang (1991) on return variances over trading versus non-trading periods, Lockwood and Linn (1990) on overnight and intraday return volatility and Ederington and Lee (1993) on the impact of macroeconomic announcements on inter- and intraday return volatility.

Engle (1982). Most of these studies fall into one of three categories. Firstly, some authors investigate the interrelation between returns in geographically separated financial markets that trade sequentially, with a focus on the transmission of information as measured by the degree of spill-over in the mean returns and/or volatility from one market to the next [3]. A second strand of this literature is concerned with the lead–lag relations between two or more markets that trade simultaneously [4]. Finally, a third group of papers explores the role of information flow and other microstructure variables as determinants of intraday return volatility [5].

Direct comparison of these intraday volatility studies is complicated by the different sampling frequencies employed. Nonetheless, as noted by Ghose and Kroner (1994) and Guillaume (1994), the results regarding the implied degree of volatility persistence appear puzzling and in stark conflict with the aggregation results for ARCH models developed by Nelson (1990, 1992), Drost and Nijman (1993) and Drost and Werker (1996). One potential explanation is that these theoretical predictions about the relationship between parameter estimates at different sampling frequencies do not generally apply in the face of strong intraday periodicity, a fact that has gone largely unnoticed. The most comprehensive prior attempt at direct modeling of this intraday heteroskedastic pattern in returns is provided by a series of papers by the research group at Olsen and Associates on the foreign exchange market e.g. Müller et al. (1990, 1993) and Dacorogna et al. (1993). They apply time invariant polynomial approximations to the activity in the distinct geographical regions of the market over the 24-hour trading cycle [6]. Although this might be a reasonable assumption for the foreign exchange market, we propose an alternative and more general methodology that allows the shape of the periodic pattern to also depend on the current overall level of return volatility. This feature makes the procedure readily applicable to the analysis of high frequency financial data in general and turns out to be essential for our investigation of the stock market. While our approach accounts for the pronounced intraday patterns, we explicitly do not make any attempts to correct for the lower frequency interdaily patterns that also exist e.g. day-of-the-week and holiday effects which

[3] Examples of early contributions are Engle et al. (1990) for foreign exchange markets and Hamao et al. (1990) for various national equity index returns.

[4] See for example Baillie and Bollerslev (1991) and Chan et al. (1991).

[5] This literature is exemplified by Bollerslev and Domowitz (1993), Locke and Sayers (1993), Laux and Ng (1993), Foster and Viswanathan (1995) and Goodhart et al. (1993). This research is partially motivated by an attempt to identify the economic origins of the volatility clustering phenomenon as motivated by the mixture of distributions hypothesis; see for example Clark (1973), Tauchen and Pitts (1983), Harris (1986b, 1987), Gallant et al. (1991), Ross (1989) and Andersen (1994, 1996).

[6] One may note that the de-volatilization procedure proposed by Zhou (1992) implicitly adjusts for the intraday periodicity in the adaptive calculation of the volatility increments from tick-by-tick observations. Along similar lines, the notion of time deformation in modeling time varying volatility in financial markets has recently been advocated by Ghysels and Jasiak (1994).

are most certainly present in both of the data sets analyzed here. These inter-daily features are clearly less significant and not critical for the high frequency analysis pursued here. Yet, in analyses of longer run phenomena, accounting for these effects may be equally important and could in principle be incorporated along the same lines.

The remainder of the paper is organized as follows. Section 2 describes our data and summarizes the intraday average return patterns. Section 3 contains an analysis of the correlation structure of both raw and absolute 5-minute returns, as well as a comparison to the corresponding properties of the two daily time series. The impact of periodic heteroskedasticity on the 5-minute correlations is strong, while the evidence of standard conditional heteroskedasticity, although evident at the daily level, appears weak at many intraday frequencies. This motivates our simple model of intraday returns that renders formal assessments of the relation between the intra- and interdaily correlation patterns feasible. Section 4 investigates the properties of temporally aggregated intraday returns. Estimates of the degree of volatility persistence at the various sampling frequencies are contrasted to the theoretical aggregation results. Our estimation strategy for characterizing the intraday periodicity is presented in Section 5. A relatively simple model that allows for a direct interaction between the level of the daily volatility and the shape of the intradaily pattern provides a close fit to the average intradaily volatility patterns for both return series, with the interaction effect being less significant for the foreign exchange market. The corresponding time series properties of the filtered returns obtained by extracting the estimated volatility patterns from the raw series is also explored. Estimation results for these returns are much more in line with the theoretical predictions. Moreover, this analysis strongly suggests that several distinct component processes affect the volatility dynamics. This finding may help shed new light on the long-memory feature in low frequency return volatility documented by a number of recent studies. Section 6 contains concluding remarks. Details regarding the construction of the 5-minute foreign exchange and equity returns employed throughout and the flexible non-parametric procedure used in the estimation of the intraday periodicity are contained in the appendices.

2. Intraday return periodicity

Our primary data set consists of 5-minute returns for the Deutschemark–U.S.$ (DM–$) exchange rate from October 1, 1992 through September 30, 1993, comprising 74,880 observations, and the Standard and Poor's 500 (S&P 500) composite stock index futures contract from January 2, 1986, through December 31, 1989, consisting of a total of 79,280 observations. A more detailed description of the data sources and the calculation of the 5-minute returns is provided in Appendix A. In addition, we use two daily time series of 3,649 spot DM–$

exchange rates from March 14, 1979 through September 29, 1993 and 9,558 observations on the S&P 500 cash index from January 2, 1953 through December 31, 1990 [7]. All the empirical work is done in parallel, with tables and figures for the foreign exchange and equity data labelled 'a' and 'b', respectively.

2.1. The Deutschemark–U.S. dollar foreign exchange data

The sample mean of the 5-minute Deutschemark appreciation of 0.000175% is indistinguishable from zero at standard significance levels given the sample standard deviation of 0.047% [8]. However, the returns are clearly not normally distributed. For example, the sample skewness of 0.367 and the sample kurtosis of 21.5 are both highly statistically significant [9]. At the same time, the maximum and minimum 5-minute returns of 1.24% and −0.637% do not suggest the presence of sharp discontinuities in the series. A small negative first order autocorrelation coefficient of −0.04 provides some support for the hypothesis that foreign exchange dealers position their quotes asymmetrically relative to the perceived 'true market price' as a way to manage their inventory positions, thus causing the midpoint of the quoted prices to move around in a fashion similar to the bid-ask bounce often observed on organized exchanges [10]. A more detailed set of summary statistics are available in Andersen and Bollerslev (1994).

In order to evaluate the intraday periodicity of the returns, Fig. 1a plots the average sample mean for each 5-minute interval. The average returns are centered around zero but numerous violations of the constant 5% confidence band for the null of an i.i.d. series occur between 09.00 GMT and 18.00 GMT (interval range 108 to 216). Allowing for different return sample variances across the day produces a more realistic time-varying confidence band that is violated at seemingly random points in time and at a frequency consistent with the 5% band (13 violations over 288 intervals). Thus, there appears to be little evidence for any systematic DM–$ appreciation or depreciation through the regular trading day cycle [11].

---

[7] The initial observation on March 14, 1979 for the exchange rates corresponds to the beginning of the EMS. The Standard 90 was replaced by the broader S and P 500 composite index on January 2, 1953. Also, for reasons discussed below, the estimates for the S and P 500 data exclude the October 1987 crash period.

[8] Assuming the returns to be uncorrelated, the standard deviation for the mean equals $0.047/(74{,}880)^{1/2} = 0.000172\%$.

[9] The standard errors of these statistics in their corresponding asymptotic normal distributions are $(6/T)^{1/2}$ and $(24/T)^{1/2}$, or 0.009 and 0.018, respectively (see e.g. Jarque and Bera, 1987).

[10] The coefficient is small in an economic sense but given the large sample size it is highly statistically significant. Bollerslev and Domowitz (1993) also report a negative first order autocorrelation in 5-minute DM–$ returns over 3 months in 1989, but find the correlation for artificially constructed 5-minute pseudo transactions price returns to be positive.

[11] This is counter to Ito and Roley (1987) who found evidence for systematic dollar appreciation during the U.S. segment of the market but dollar depreciation during the European trading hours.

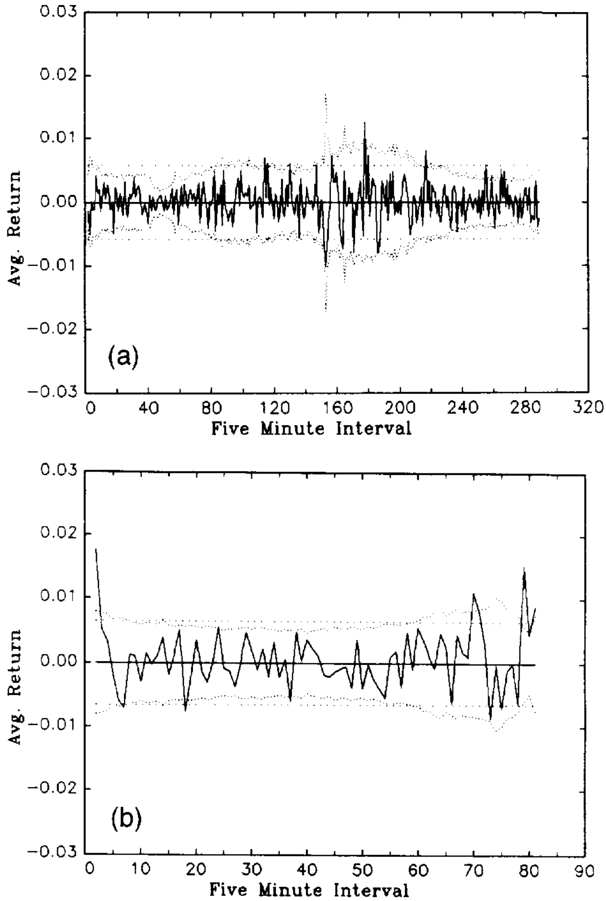

Fig. 1. Intraday average returns, (a) DM–$, (b) S&P 500.

In contrast, the exchange rate volatility fluctuates dramatically over the daily cycle. The average absolute returns over the 5-minute intervals are depicted in Fig. 2a. It reveals a pronounced difference in the volatility over the day, ranging from a low of around 0.01% at 04.00 GMT (interval 48) to a high of around 0.05% at 15.00 GMT (interval 180) [12]. This pattern is closely linked to the cycle of market activity in the various financial centers around the globe. The volatility starts out at a relatively high level followed by a slow decay up to around 03.00 GMT (interval 36). The strong drop between intervals 40 and 60 corresponds to the

[12] Since the average standard error for the absolute returns is 0.0022%, these differences are highly statistically significant.

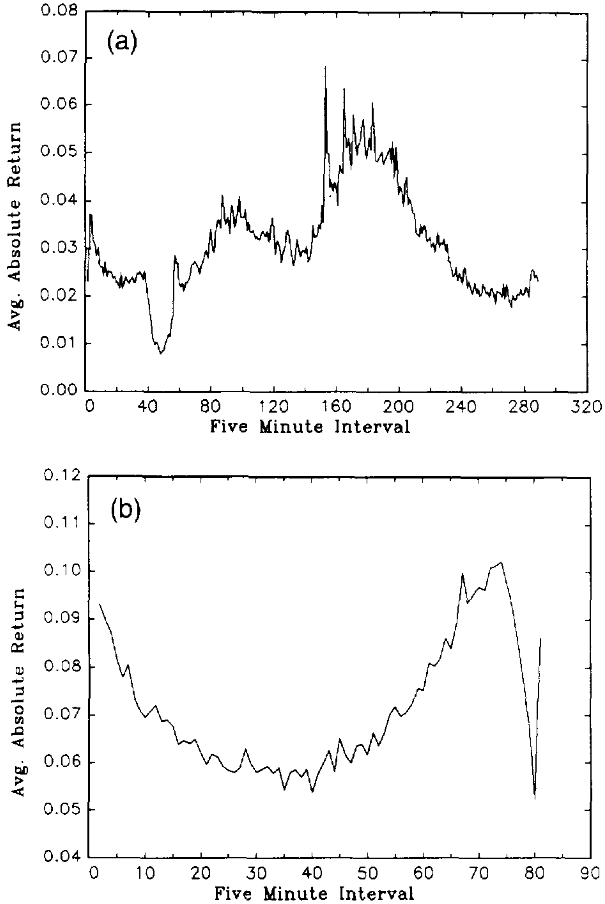

Fig. 2. Intraday average absolute returns, (a) DM–$, (b) S&P 500.

lunch hour in the Tokyo and Hong Kong markets. Activity then picks up during the afternoon session in the Far Eastern markets and is further fueled by the opening of the European markets around 07.00 GMT (interval 84). The market volatility then declines slowly until the European lunch hour at 11.30 GMT (interval 138), before it increases sharply during the overlap of afternoon trading in Europe and the opening of the U.S. markets around 13.00 GMT, or 7.00 a.m. New York (interval 156). After the European markets close volatility declines monotonically until trading associated with the Far Eastern markets starts to pick up again around 21.00 GMT (interval 252). The robustness of this intraday volatility pattern is confirmed by the sub-sample analysis and the sorting of days according to volatility levels reported in the more detailed analysis in Andersen and Bollerslev (1994), which is also consistent with earlier findings in Wasserfallen (1989), Müller et al. (1990), Baillie and Bollerslev (1991) and Dacorogna et al. (1993) [13]. Standard summary statistics further verify the overwhelming significance of this intraday volatility pattern. In particular, the first order autocorrelation coefficient for the absolute 5-minute returns of $\rho_1^A = 0.309$ exceeds the $1/\sqrt{T}$ asymptotic standard error by almost a factor of one hundred, while the Ljung–Box statistic for up to tenth order serial correlation in $|R_{t,n}|$ equals $Q^A(10) = 36{,}680$ [14].

2.2. The standard and poor's 500 stock index futures data

The basic features of the 5-minute S&P 500 are qualitatively similar to those of the 5-minute DM–$ returns. Perhaps, the most notable difference is that the standard deviation for the stock index futures return of 0.104% is more than double the value for the foreign exchange market. However, since the overnight returns for the S&P 500 are excluded, the average 5-minute standard deviation corresponds to active trading on the CME only, whereas the foreign exchange returns cover the entire 24-hour trading cycle and therefore include periods of relatively slow activity. Even so, when judged by auxiliary statistics such as the sample minimum and maximum of 2.22% and −2.76%, the equity market exhibits the more volatile returns. Another distinguishing feature is the virtual absence of autocorrelation in the futures returns. Although the first ten autocorrelation coefficients are highly significant, the coefficients are economically small and have unpredictable signs [15]. This lack of correlation contrasts sharply with results reported by most studies on the intraday S&P 500 cash market, where non-synchronous trading effects imply that stale prices may enter the calculation of the index (see e.g. Chan et al., 1991) [16].

The intraday periodic patterns over the eighty 5-minute intraday intervals are depicted in Fig. 1b and Fig. 2b. Apart from the positive returns over the initial 5-minute interval from 8.35 to 8.40 a.m. and towards the end of the trading day,

[13] We follow Dacorogna et al. (1993) in using GMT time scale throughout our analysis. Daylight savings time is observed in Europe and North America, but not in East Asia. From the sub-sample analysis in Andersen and Bollerslev (1994) this gives rise to a one hour difference in the peaks associated with the regular release of U.S. macroeconomic announcements at 08.30 a.m. corresponding to interval 162 for winter time and interval 150 for summer time. Ederington and Lee (1993) and Harvey and Huang (1991) also suggest that macroeconomic announcement effects have a distinct impact on the average volatility in early Friday morning trading in the U.S. segment of the market. We do not pursue this or any other day-of-the-week effects any further here, however.

[14] The first ten autocorrelations for $|R_{t,n}|$ are 0.309, 0.256, 0.238, 0.214, 0.212, 0.199, 0.204, 0.182, 0.185 and 0.182.

[15] The first ten autocorrelations of the returns are 0.009, −0.003, −0.009, 0.010, −0.004, 0.018, 0.009, 0.015, 0.013 and 0.008, respectively. The corresponding Ljung–Box statistic equals $Q(10) = 87.6$.

[16] The impact of non-synchronous trading has been explored extensively in the literature (see e.g. Scholes and Williams, 1977; Lo and MacKinlay, 1990).

the violations of the 5% confidence bands for the average returns are dispersed unpredictably over the trading day [17]. Nonetheless, as was the case for the foreign exchange market, the systematic return effects are dwarfed by the systematic movements in the return volatility, here documented in Fig. 2b. The average absolute returns attain the commonly observed intraday U-shape, starting out at 0.095% in the morning, followed by a smooth decline to a level of 0.055% around noon and a steady rise to 0.105% towards the end of trading in the cash market. The subsequent drop and rise over the last fifteen minutes corresponds to the post cash market trading on the CME [18]. The robustness of this intraday periodicity in the S&P 500 returns is again underscored by the more detailed analysis in Andersen and Bollerslev (1994) in which the full four year sample is divided into calendar years as well as four daily volatility categories. The only discernible difference across these sub-sample patterns is a tendency for the right part of the 'U' to occasionally rise above the left part, creating more of a 'J' shape. Interestingly, this tendency appears to be concentrated on high volatility days. The model proposed in Section 5 below explicitly accounts for this phenomenon.

Several recent studies have attempted to rationalize the pronounced U-shape pattern in intraday stock market volatility by strategic interaction of traders around market openings and closures (see e.g. Admati and Pfleiderer, 1988, 1989; Foster and Viswanathan, 1990; Son, 1991; Brock and Kleidon, 1992). Even though the foreign exchange market operates on a continuous basis, the volatility pattern for the DM–$ depicted in Fig. 2a may be viewed, tentatively, as a sum of two overlapping U-shapes corresponding to the Far East and European trading hours, along with an inverted U-shape for the U.S. segment of the market. Hence, in spite of obvious differences in market microstructures, the foreign exchange returns are calculated from quotes in a 24-hour over-the-counter market while the equity returns are obtained from transaction prices on an organized futures market with well defined daily closings, the pattern of intraday periodicity in the two markets share important common characteristics.

3. Characterization and modeling of the correlation structure in intraday returns

3.1. Intraday return correlations

While the intraday volatility patterns documented in the preceding section may be irrelevant for standard studies of the return dynamics based on price observa-

[17] This is related to the findings in Harris (1986a) who reports that the average positive returns in the equity markets tend to occur over the first 45 min of the trading day and the very last trade of the day. Notice also, that there is no indication of any abnormal positive returns after the cash market is closed.

[18] This U-shaped pattern in the volatility of S&P 500 futures prices following the closure of the cash market has also recently been documented by Chang et al. (1995).

tions at daily frequencies, conclusions drawn from the recent surge of empirical papers on return volatility and market microstructure variables at the intraday frequencies are likely subject to severe distortions due to the strong periodicity in returns. We therefore supplement the prior investigation of the unconditional volatility patterns with an explicit look at the dynamic features of our two return series.

Fig. 3a and b display the sample autocorrelations of the 5-minute returns for up to five days i.e. 1440 observations for the foreign exchange and 400 for the equity returns. All values are small and beyond the first few lags the series resemble realizations of white noise. Thus, we again detect little of interest in the mean

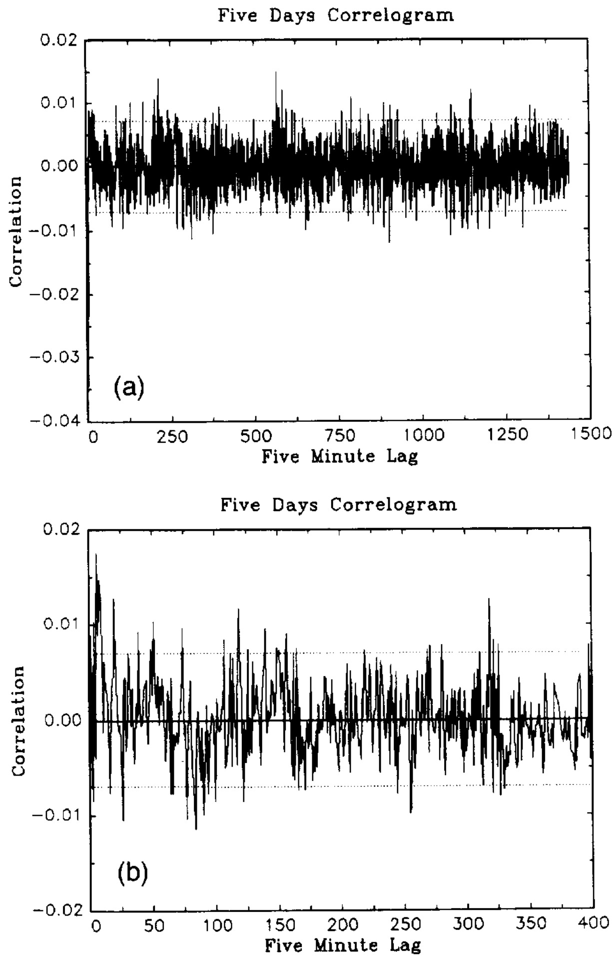

Fig. 3. Five days correlogram of intraday returns, (a) DM–$, (b) S&P 500.

returns. In contrast, the autocorrelation patterns for the absolute returns are strikingly regular. Consider the series for the DM–$ exchange rate in Fig. 4a. The strong intraday pattern induces a distorted U-shape in the sample correlogram [19]. Notice also how the size of the autocorrelations at the daily frequencies decay slowly over the first four days, only to increase slightly at the fifth, or weekly, frequency. This signals the presence of a minor day-of-the-week effect, which we ignore in the remainder. Fig. 4b for the S&P 500 futures returns is equally telling. The slowly declining U-shape occupies exactly 80 intervals, corresponding to the daily frequency.

3.2. Interpretation in terms of a suggestive intraday return model

The pronounced systematic fluctuations in the return correlogram provide an initial indication that direct ARCH modeling of the intraday return volatility would be hazardous. Standard ARCH models imply a geometric decay in the return autocorrelation structure and simply cannot accommodate strong regular cyclical patterns of the sort displayed in Fig. 4. Instead, it seems intuitively clear that the combination of recurring cycles at the daily frequency and a slow decay in the average autocorrelations may be explained by the joint presence of the pronounced intraday periodicity documented above coupled with the strong daily conditional heteroskedasticity [20]. The following stylized model provides a simple specification of the interaction between these two components,

$$R_t \equiv \sum_{n=1}^{N} R_{t,n} = \sigma_t \frac{1}{N^{1/2}} \sum_{n=1}^{N} s_n Z_{t,n}. \tag{1}$$

Here, R_t denotes the daily continuously compounded return calculated from the N uncorrelated intraday return components, $R_{t,n}$. The conditional volatility factor for day t is denoted by σ_t, while s_n refers to a deterministic intraday periodic component and $Z_{t,n}$ is an i.i.d. mean zero, unit variance error term assumed to be independent of the daily volatility process, $\{\sigma_t\}$. Both volatility components must be non-negative i.e. $\sigma_t > 0$ a.s. for all t and $s_n > 0$ for all n. The following terminology for the normalized, deterministic sample means and covariances for the periodic structure will prove convenient:

$$\frac{1}{N}\sum_{n=1}^{N} \equiv M(s) = 1, \qquad \frac{1}{N}\sum_{n=1}^{N} s_n s_{n-i} \equiv M(ss_i), \qquad \frac{1}{N}\sum_{n=1}^{N} S_n^2 \equiv M(s^2),$$

where $s_{n+jN} \equiv s_n$ for any integer j and $0 < n \le N$.

[19] A corresponding figure is presented by Dacorogna et al. (1993). However, in their analysis the correlations at the daily frequencies are sharply diminished due to a strong weekend effect. By excluding weekend returns we have effectively eliminated this distortion.

[20] The temporal variation in daily financial market volatility have been successfully modeled by ARCH type processes (see Bollerslev et al. (1992) for a survey of this extensive literature).

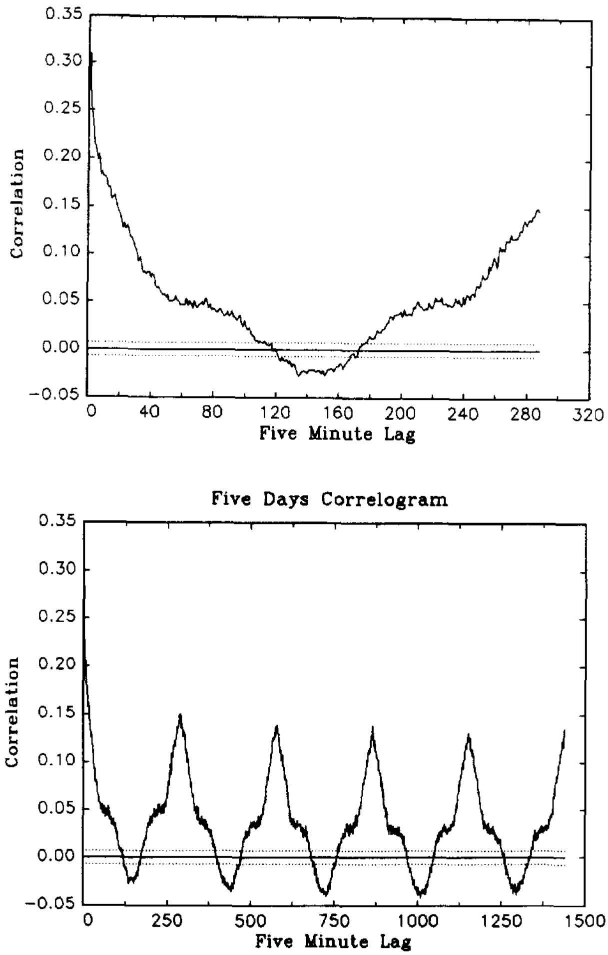

(a)

Fig. 4. Intraday absolute returns, (a) DM–\$, (b) S&P 500.

In the absence of intraday periodicity ($s_n = 1$ for all n) the daily returns may be represented in the form $R_t = \sigma_t N^{-1/2} \sum_{n=1,N} Z_{t,n}$, where the return component $N^{-1/2} \sum_{n=1,N} Z_{t,n}$ is i.i.d. with mean zero and unit variance. Thus, Eq. (1) extends the standard volatility model for daily returns to an intraday setting with independent return innovations and deterministic volatility patterns. Of course, this type of periodicity is annihilated when the returns are measured at the daily frequency. In particular, letting Z_t denote an i.i.d. random variable with $E(Z_t) = 0$ and $\mathrm{Var}(Z_t) = 1$, we have

$$R_t = M^{1/2}(s^2)\sigma_t Z_t, \tag{2}$$

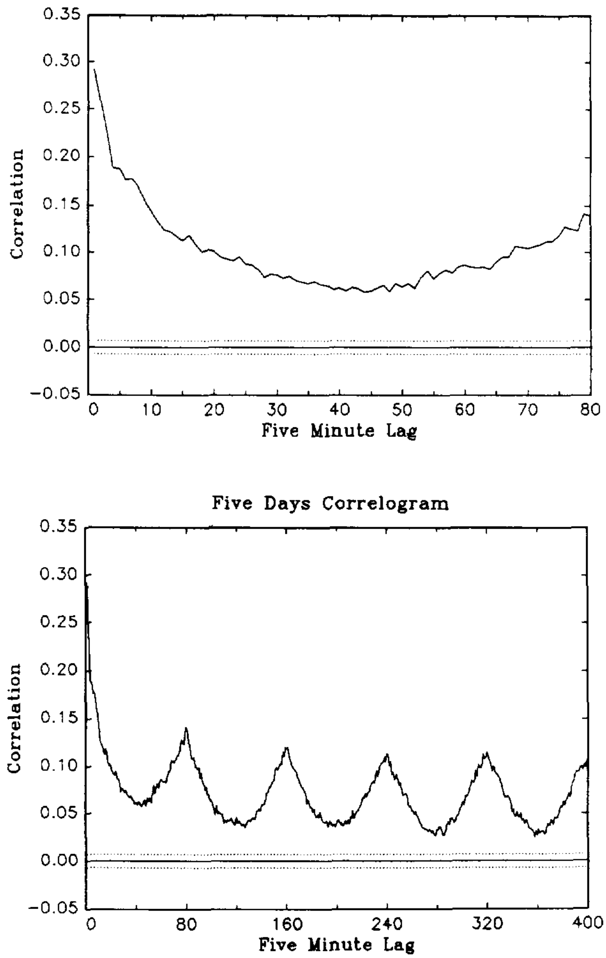

(b)

Fig. 4 (continued).

so that the expected absolute return equals $M^{1/2}(s^2)\sigma_t E|Z_t|$. Since $M^{1/2}(s^2) \geq 1$, the expected daily absolute return is an increasing function of the fluctuations in the intraday periodic pattern. However, this effect is limited to a scale factor. Letting $c \equiv (E|Z_t|)^{-2} - 1 > 0$, it follows that for $t \neq \tau$,

$$\mathrm{Corr}(|R_t|, |R_\tau|) = \frac{\mathrm{Cov}(\sigma_t, \sigma_\tau)}{\mathrm{Var}(\sigma_t) + cE(\sigma_t^2)}. \tag{3}$$

Hence, the presence of periodic components reduces the overall level of the interdaily return autocorrelations, without affecting the autocorrelation pattern [21]. In contrast, the periodicity may have a strong impact on the autocorrelation pattern for the absolute intraday returns. Straightforward calculations reveal

$$\mathrm{Corr}(|R_{t,n}|, |R_{\tau,m}|)$$
$$= \frac{M(ss_{n-m})\mathrm{Cov}(\sigma_t, \sigma_\tau) + \mathrm{Cov}(s_n, s_m)E^2(\sigma_t)}{M(s^2)\mathrm{Var}(\sigma_t) + c_N E(\sigma_t^2)M(s^2) + E^2(\sigma_t)\mathrm{Var}(s)}, \tag{4}$$

where $\mathrm{Var}(s) \equiv M(s^2) - M^2(s)$, $\mathrm{Cov}(s_n, s_m) \equiv M(ss_{n-m}) - M^2(s)$ and $c_N \equiv E^{-2}|Z_{t,n}| - 1$. Eq. (4) illustrates the interaction between the periodicity in absolute returns at the intradaily level and the conditional heteroskedasticity at the daily level. For adjacent trading days the impact of the positive correlation in the daily return volatility, captured by $\mathrm{Cov}(\sigma_t, \sigma_\tau)$, is strong and induces positive dependence in the absolute returns, but as the distance between t and τ grows this effect becomes less important which is consistent with the slow decay in the correlograms in the bottom panels of Fig. 4. At the same time, the correlograms are affected by the strong intraday periodicity. For example, consider the display for the absolute S&P 500 returns in Fig. 4b. The correlations attain their lowest values around lag forty, or half a trading day. This corresponds to the bottom of the U-shape for the average absolute returns depicted in Fig. 2b. Clearly, the population covariance, $\mathrm{Cov}(s_n, s_m)$, is minimized and significantly negative, at this frequency. Eq. (4) verifies that the negative correlation between the 5-minute absolute returns, realized about half-a-day apart, translates into a negative contribution to the corresponding correlogram at the 3–4 hour frequencies. Likewise, Fig. 2a indicates that there is strong negative correlation between the absolute foreign exchange returns in the intervals 80–225 (covering about half-a-day) and all the remaining 5-minute returns. Not surprisingly, the lower panel of Fig. 4a verifies that this again results in highly significant troughs in the correlogram around the 12 hour frequency (and its harmonics). Indeed, the impact is now sufficiently strong that the absolute return autocorrelations turn negative. This is truly remarkable given the very large positive autocorrelations found at the daily frequency and it is testimony to the profound impact of the periodic structure on the intraday return dynamics. In terms of the specification in Eq. (4), the size of the second, negative, term of the numerator exceeds the first, positive, term around the 12 hour frequency.

[21] Consistent with the findings of Granger and Ding (1996), informal investigations reveal that the dynamic dependencies are significantly more pronounced for the absolute as opposed to the squared sample returns. Consequently, our intraday modeling focuses on the patterns in $|R_t|$ and $|R_{t,n}|$, rather than R_t^2 and $R_{t,n}^2$.

3.3. Long-run implications and comparison to daily returns

To further assess the descriptive accuracy of the formulation in Eq. (1), we now investigate the long-run implications for the correlogram. It is convenient to focus on the daily frequency i.e. $n = m$,

$$\operatorname{Corr}(|R_{t,n}|, |R_{\tau,n}|) = \frac{\operatorname{Cov}(\sigma_t, \sigma_\tau) + (\operatorname{Var}(s)/M(s^2))E^2(\sigma_t)}{\operatorname{Var}(\sigma_t) + c_N E(\sigma_t^2) + (\operatorname{Var}(s)/M(s^2))E^2(\sigma_t)}. \tag{5}$$

Fig. 5a and b display the first forty autocorrelations for the absolute returns of the two daily time series on the DM–$ spot exchange rate and the S&P 500 cash

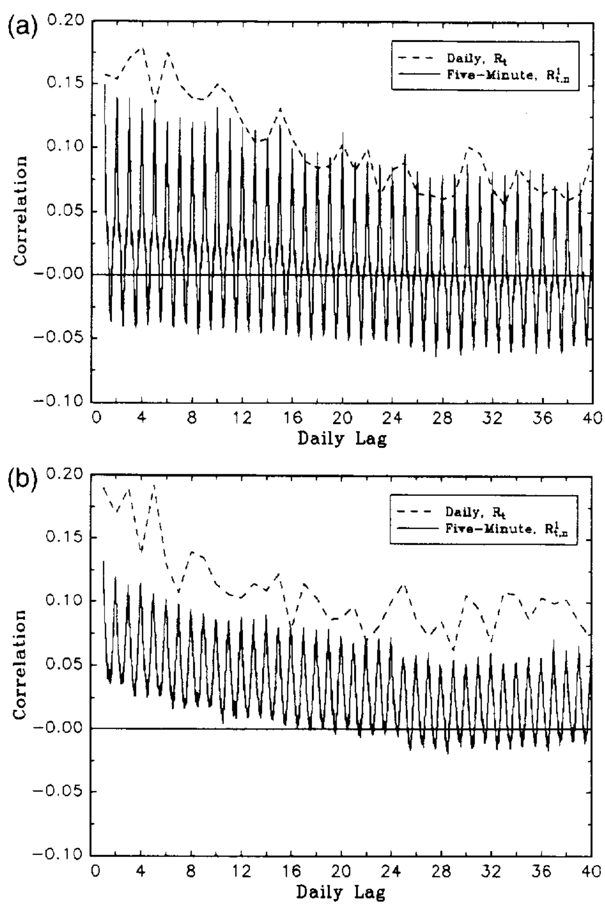

Fig. 5. Forty days correlogram of absolute returns, (a) DM–$, (b) S&P 500.

index, along with the corresponding 5-minute intraday autocorrelations out to lag 11,520 and 3,200, respectively [22]. Direct comparison between the empirically estimated $\mathrm{Corr}(|R_t|, |R_\tau|)$ in Eq. (3) and the expression for $\mathrm{Corr}(|R_{t,n}|, |R_{\tau,n}|)$ above are complicated by the different sample periods required for reliable inference. Nonetheless, it is clear that the decay rate in the local maxima of the intradaily absolute correlogram and the daily return autocorrelations should be qualitatively comparable, as the $\mathrm{Cov}(\sigma_t, \sigma_\tau)$ term governs both. Comparing the peaks in the intraday correlograms to the daily autocorrelations in Fig. 5a and b confirms this implication of our stylized model; the dominant rate of decay is strikingly similar for both markets [23].

The findings in this section demonstrate a strong correspondence between the qualitative implications of the model outlined in Eq. (1) and the stylized empirical facts. It suggests that our model, stressing the conditional heteroskedasticity at the daily level along with the strong deterministic periodicity at the intradaily level, may serve as a good starting point for high frequency volatility modeling and as such it constitutes the basis for the subsequent analysis in Section 5.

4. Implications for volatility modeling and high frequency return aggregation

Section 3 demonstrates that the distinct intraday periodicity has a strong impact on the autocorrelation patterns of the 5-minute returns. The question therefore arises whether more formal time series modeling of return volatility is similarly affected by the presence of periodic features, and if so, whether some observation intervals are preferable relative to others for the purpose of drawing inference concerning the dynamic features of interest. In order to address these issues this section presents an extensive analysis of the properties of the return series obtained at a range of different intradaily and interdaily frequencies.

4.1. Characterization of the intraday returns at the various frequencies

Summary statistics for the foreign exchange market are provided in Table 1a for all seventeen possible intraday returns with a 24-hour periodicity. The returns are continuously compounded i.e. the nth return on day t for the series at $(k \cdot 5)$-minute intervals is defined by $R_{t,n}^{k} \equiv \Sigma_{i=(n-1)k+1,nk} R_{t,i}$, $t = 1, 2, \ldots, 260$,

[22] The sample autocorrelations are generally negatively biased and become less precise as the lag length grows; see Percival (1993) who point out that the sum of all the sample autocorrelations by construction equals zero. We therefore limit our analysis to lag-lengths which may appear large, but nonetheless constitute a modest fraction of the total intraday sample.

[23] The slow rate of decay in the autocorrelation functions is also in accordance with the apparent long-memory feature of asset return volatility documented by a number of recent studies (see e.g. Baillie et al., 1996; Ding et al., 1993).

$n = 1, 2, \ldots, K$ where $K \equiv 288/k$ refers to the number of returns per day. Note that while the 5-minute return series consists of 74,880 observations, the hourly series contains only 6,240 observations and the 1/2-day return series has a mere 520 observations. These differences should be kept in mind when interpreting the evidence.

The standard deviations in Table 1a grow at a rate almost proportional to the square root of the sampling frequency. This is consistent with the 5-minute returns being approximately uncorrelated, although there is a small, but highly significant, negative first order autocorrelation coefficients at the higher frequencies. As mentioned, the weak negative correlation may be the result of spread positioning by dealers causing mean reversion in the quote midpoints; an effect similar to a bid–ask bounce in transactions data [24]. In line with this explanation, the ρ_1 coefficients generally turn insignificant at the 40-minute and lower frequencies. Further corroborating evidence along these lines is provided by the variance ratio statistics,

$$\mathrm{VR} = \frac{K \cdot \mathrm{Var}_T\left(R_{t,n}^k\right)}{\mathrm{Var}_T\left(\Sigma_{n=1}^{K} R_{t,n}^k\right)}, \tag{6}$$

where $\mathrm{Var}_T(R_{t,n}^k)$ and $\mathrm{Var}_T(\Sigma_{n=1,k} R_{t,n}^k)$ denote the sample variances for the intraday and daily returns, respectively. Expanding the daily variance estimate in the denominator demonstrates that a value of the VR-statistic below unity will result from positive autocorrelation between adjacent return components, while a statistic above one is indicative of predominantly negatively correlated intraday returns [25]. Finally, it is worth noting from Table 1a, that the kurtosis of the DM–$ returns increases almost monotonically with the sampling frequency.

The first order autocorrelations of the absolute returns, ρ_1^{A} are, not surprisingly, all highly significant for the shorter intervals. However, beyond the 2-hour sampling frequency the autocorrelations drop off very sharply and in fact turn negative at the 8 and 12 hourly frequencies ($k = 96, 144$). This is, of course, consistent with the negative region of the 5-minute absolute return correlogram in Fig. 4a. The VR^{A}-statistic reported in the final column of Table 1a is calculated by replacing $R_{t,n}^k$ with $|R_{t,n}^k|$ in the definition of VR in Eq. (6) [26]. The statistic starts out at 0.05 for the 5-minute returns and rises almost monotonically to 0.69 for the

[24] Note that the standard deviation of the 5-minute returns is less than the average quoted bid–ask spread. According to Bollerslev and Melvin (1994), more than half of the DM–$ quotes are posted with a spread of 0.10%, while the second most common and lowest regularly posted spread of 0.05% accounts for about a quarter of the quotations.

[25] Formal tests for serial correlation based on the VR-statistic may be calculated as outlined in Lo and MacKinlay (1989).

[26] Note that the denominator in this VR^{A}-statistic involves the variance of the sum of the absolute returns rather than the absolute value of the sum of the returns. The expected value of the VR-statistic would not equal unity under the latter definition.

Table 1

k	T/k	Mean	St. Dev.	Skew.	Kurtosis	ρ_1	$Q(10)$	VR	ρ_1^{A}	$Q^{\mathrm{A}}(10)$	VR^{A}
(a) Summary statistics for intraday DM–$ exchange rate											
1	74,880	0.018	0.047	0.368	21.5	−0.040	281	1.194	0.309	36,680	0.054
2	37,440	0.035	0.066	0.363	16.6	−0.070	263	1.162	0.313	17,563	0.071
3	24,960	0.053	0.079	0.200	13.6	−0.089	220	1.118	0.307	10,710	0.086
4	18,720	0.070	0.089	0.276	14.0	−0.082	154	1.084	0.287	6,296	0.099
6	12,480	0.105	0.107	0.534	12.6	−0.043	36.5	1.023	0.268	2,757	0.127
8	9,360	0.140	0.121	0.135	9.11	−0.023	22.0	0.994	0.272	1,736	0.149
9	8,320	0.158	0.126	0.345	10.1	0.002	18.9	0.948	0.251	1,212	0.161
12	6,240	0.210	0.148	0.326	11.0	−0.001	12.7	0.978	0.229	609	0.193
16	4,680	0.280	0.170	0.318	8.77	0.032	13.1	0.968	0.246	425	0.235
18	4,160	0.315	0.178	0.489	10.6	0.058	21.9	0.947	0.193	219	0.260
24	3,120	0.420	0.212	0.166	8.94	0.011	16.7	1.012	0.164	159	0.311
32	2,340	0.560	0.246	0.326	9.15	0.018	43.3	1.018	0.171	238	0.373
36	2,080	0.630	0.253	0.329	7.89	0.047	37.7	0.954	0.097	109	0.416
48	1,560	0.840	0.300	0.400	5.92	0.002	30.5	1.007	0.075	66.6	0.487
72	1,040	1.261	0.373	0.319	6.59	−0.019	20.6	1.042	0.007	67.6	0.679
96	780	1.681	0.423	0.389	5.19	−0.022	18.2	1.004	−0.025	53.2	0.653
144	520	2.521	0.520	0.192	4.31	−0.021	12.2	1.012	−0.033	28.2	0.692
(b) Summary statistics for intraday S&P 500 returns											
1	79,280	0.064	0.104	−0.597	29.3	0.009	87.6	0.774	0.292	32,425	0.099
2	39,640	0.128	0.150	−1.212	30.9	−0.009	81.5	0.801	0.285	12,641	0.128
4	19,820	0.255	0.212	−1.609	33.2	0.014	53.4	0.795	0.232	3,374	0.179
5	15,856	0.319	0.234	−1.755	33.3	0.032	60.6	0.780	0.243	2,323	0.205
8	9,910	0.511	0.299	−1.478	22.2	0.047	44.6	0.793	0.207	1,295	0.261
10	7,928	0.638	0.339	−1.417	21.2	0.039	41.7	0.819	0.211	973	0.300
16	4,955	1.021	0.437	−1.803	26.6	0.040	35.2	0.849	0.135	437	0.405
20	3,964	1.277	0.499	−1.869	27.4	0.016	24.4	0.884	0.114	268	0.463
40	1,982	2.553	0.728	−1.541	16.2	0.026	22.6	0.942	0.148	175	0.673

12 hourly returns. The results for the multiple day returns reported in Andersen and Bollerslev (1994) continue this near monotone ascent, reaching 1.94 for the biweekly sampling interval. The smooth increase suggests that a common component accounts for a substantial part of the positive higher order dependence in all of the return series. The corresponding ρ statistics of 0.123 and 0.118 for the weekly and biweekly sampling frequencies also testify to the importance of the interday heteroskedasticity [27].

[27] Hence, the VR^{A}-statistics convey a coherent message about the degree of conditional heteroskedasticity in the series. As a set of simple diagnostics, these statistics may therefore be more informative about the nature of the volatility process than the standard Ljung–Box statistics for tenth order serial correlation in the absolute returns, $Q^{\mathrm{A}}(10)$, which appear both erratic and highly dependent on the sample size.

The summary statistics for the S&P 500 index futures returns in Table 1b largely parallel those for the DM–$ returns. However, in contrast to the results for the exchange rates, the first order autocorrelations and the VR-statistics in Table 1b all indicate a slight positive intraday dependence. Moreover, the equity returns are negatively skewed and display very significant excess kurtosis [28]. Finally, the intraday return periodicity, here depicted in Fig. 2b, again have a strong effect on the correlations for the absolute intraday returns, although the decay in the ρ_1^A coefficients for the lower frequencies is less pronounced than for the exchange rates.

4.2. Specification of the volatility model and the associated persistence measures

Numerous recent studies have relied on more formal time series techniques in the analysis of high frequency return dynamics both within and across different markets. The most commonly employed formulation is the GARCH(1, 1) model proposed independently by Bollerslev (1986) and Taylor (1986). Thus, in order to evaluate the potential impact of the strong intraday periodicity in this context we

Notes to Table 1:

(a) The percentage returns are based on interpolated 5-minute logarithmic average bid-ask quotes for the Deutschemark–U.S. dollar spot exchange rate from October 1, 1992 through September 29, 1993. Quotes from Friday 21.00 Greenwich mean time (GMT) through Sunday 21.00 GMT have been excluded, resulting in a total of 74,880 return observations. The length of the different intraday return sampling intervals equals $5 \cdot k$ minutes. Each time series has a total of T/k non-overlapping return observations. The sample means have been multiplied by one hundred. The columns indicated by ρ_1 and ρ_1^A give the first order autocorrelations for the returns and the absolute returns. The Ljung and Box (1978) portmanteau test for up to tenth order serial correlation in the returns and the absolute returns are denoted by $Q(10)$ and $Q^A(10)$, respectively. The variance ratio's for the different sampling frequencies versus the daily return variance are denoted by VR. The corresponding variance ratio statistics for the absolute returns are given in the VR^A column.

(b) The returns are based on 79,280 interpolated 5-minute futures transactions prices for the Standard and Poor's 500 composite index. The sample period ranges from January 2, 1986 through December 31, 1989, excluding the period from October 15, 1987 through November 13, 1987. Overnight five minute returns have also been deleted, resulting in a total of 80 intraday return observations from 08.35 through 15.15 for each of the 991 days in the sample. The length of the different intraday return sampling intervals equals $5 \cdot k$ minutes. Each time series has a total of T/k non-overlapping return observations. The sample means have been multiplied by one hundred. The columns indicated by ρ_1 and ρ_1^A give the first order autocorrelations for the returns and the absolute returns. The Ljung and Box (1978) portmanteau test for up to tenth order serial correlation in the returns and the absolute returns are denoted by $Q(10)$ and $Q^A(10)$, respectively. The variance ratio's for the different sampling frequencies versus the daily return variance are denoted by VR. The corresponding variance ratio statistics for the absolute returns are given in the VR^A column.

[28] The negative skewness may be interpreted as evidence of the so-called 'leverage' and/or 'volatility feed-back' effects discussed by Black (1976), Christie (1982) and Nelson (1991), and Campbell and Hentschel (1992), respectively.

present MA(1)–GARCH(1, 1) estimation results for each of the intradaily sampling frequencies in Table 2a and b. Formally, the model is defined by

$$R_{t,n}^{k} = \mu_{(k)} + \theta_{(k)}\epsilon_{t,n-1}^{k} + \epsilon_{t,n}^{k},$$

and

$$\left(\sigma_{t,n}^{k}\right)^{2} = \omega_{(k)} + \alpha_{(k)}\left(\epsilon_{t,n-1}^{k}\right)^{2} + \beta_{(k)}\left(\sigma_{t,n-1}^{k}\right)^{2},$$

where $E_{t,n-1}(\epsilon_{t,n}^{k}) = 0$ and $E_{t,n-1}[(\epsilon_{t,n}^{k})^{2}] = (\sigma_{t,n}^{k})^{2}$ denotes the conditional return variance over the subsequent intraday period, with the subscript $(t, 0)$ defined to equal $(t-1, K)$. The reported parameter estimates for $\alpha_{(k)}$ and $\beta_{(k)}$ are obtained

Table 2

k	T/k	$\alpha_{(k)}$	$\beta_{(k)}$	$\alpha_{(k)}+\beta_{(k)}$	Half life	Mean lag	Median lag
(a) Persistence of MA(1)–GARCH(1, 1) models for intraday DM–$ exchange rate							
$R_{t,n}^{k} \equiv 100\cdot\sum_{i=(n-1)k+1,nk} R_{t,i} = \mu_{(k)} + \theta_{(k)}\epsilon_{t,n-1}^{k} + \epsilon_{t,n}^{k}$ $(\sigma_{t,n}^{k})^{2} = \omega_{(k)} + \alpha_{(k)}(\epsilon_{t,n-1}^{k})^{2} + \beta_{(k)}(\sigma_{t,n-1}^{k})^{2}$, $t = 1, 2, \cdots, 260$, $n = 1, 2, \cdots, 288/k$							
1	74,880	0.193 (0.011)	0.822 (0.009)	1.015	∞	∞	∞
2	37,440	0.229 (0.012)	0.774 (0.008)	1.003	∞	∞	∞
3	24,960	0.273 (0.018)	0.708 (0.014)	0.981	533	725	488
4	18,720	0.287 (0.019)	0.677 (0.016)	0.964	375	488	320
6	12,480	0.322 (0.035)	0.579 (0.033)	0.901	200	233	138
8	9,360	0.286 (0.028)	0.581 (0.037)	0.868	195	207	108
9	8,320	0.306 (0.035)	0.521 (0.042)	0.828	165	167	81
12	6,240	0.311 (0.047)	0.395 (0.069)	0.706	119	105	35
16	4,680	0.261 (0.039)	0.456 (0.074)	0.718	167	136	< 40
18	4,160	0.270 (0.061)	0.246 (0.124)	0.516	94	67	< 45
24	3,120	0.018 (0.015)	0.969 (0.026)	0.988	6,771	5,919	1,878
32	2,340	0.016 (0.008)	0.975 (0.013)	0.991	12,159	11,219	4,318
36	2,080	0.011 (0.004)	0.978 (0.005)	0.989	11,311	8,293	266
48	1,560	0.011 (0.004)	0.979 (0.005)	0.990	17,084	13,229	1,748
72	1,040	0.007 (0.005)	0.987 (0.004)	0.987	19,585	10,153	< 180
96	780	0.014 (0.008)	0.969 (0.007)	0.983	19,637	13,202	< 240
144	520	0.010 (0.010)	0.960 (0.007)	0.970	16,329	5,988	< 360
(b) Persistence of MA(1)–GARCH(1, 1) models for intraday S&P 500 returns							
$R_{t,n}^{k} \equiv 100\cdot\sum_{i=(n-1)k+1,nk} R_{t,i} = \mu_{(k)} + \theta_{(k)}\epsilon_{t,n-1}^{k} + \epsilon_{t,n}^{k}$ $(\sigma_{t,n}^{k})^{2} = \omega_{(k)} + \alpha_{(k)}(\epsilon_{t,n-1}^{k})^{2} + \beta_{(k)}(\sigma_{t,n-1}^{k})^{2}$, $t = 1, 2, \ldots, 991$, $n = 1, 2, \ldots, 80/k$							
1	79,280	0.137 (0.004)	0.838 (0.005)	0.975	137	168	105
2	39,640	0.180 (0.010)	0.765 (0.011)	0.945	121	138	79
4	19,820	0.223 (0.024)	0.664 (0.036)	0.887	116	118	57
5	15,856	0.230 (0.067)	0.630 (0.123)	0.861	116	112	49
8	9,910	0.053 (0.027)	0.935 (0.036)	0.988	2,213	2,602	1,559
10	7,928	0.048 (0.018)	0.940 (0.023)	0.988	2,947	3,437	2,043
16	4,955	0.148 (0.333)	0.764 (0.694)	0.912	606	575	240
20	3,964	0.060 (0.049)	0.890 (0.092)	0.951	1,376	1,124	246
40	1,982	0.108 (0.158)	0.798 (0.315)	0.906	1,397	1,128	228

by quasi-maximum likelihood methods assuming the innovations to be conditionally normally distributed. The corresponding robust standard errors for the estimates are provided in parentheses (see Bollerslev and Wooldridge, 1992). We note that, although it usually represents a reasonable approximation, the GARCH(1, 1) model is not necessarily the preferred specification for the return generating process in all, or even most, instances. However, estimating the same model across both asset classes and all return frequencies facilitates meaningful comparisons of the findings. Moreover, it corresponds to the class of models for which theoretical aggregation results are available. The MA(1) term is included to account for the economically minor, but occasionally highly statistically significant, first order autocorrelation in the returns.

Unfortunately, an unambiguous characterization of the estimated volatility dynamics and the associated persistence properties is not possible in this non-linear setting (see Bollerslev and Engle (1993), Bollerslev et al. (1994) and Gallant et al. (1993) for further discussion of these issues). Hence, we supplement the parameter estimates for $\alpha_{(k)}$ and $\beta_{(k)}$ in Table 2a and b with three additional summary measures for the implied degree of volatility persistence. In particular, if $\alpha_{(k)} + \beta_{(k)} < 1$, the j-step ahead prediction for the conditional variance may be written as

$$E_{t,n}\left[\left(\sigma_{t,n+j}^{k}\right)^{2}\right] = \sigma^{2} + \left(\alpha_{(k)} + \beta_{(k)}\right)^{j}\left[\left(\sigma_{t,n}^{k}\right)^{2} - \sigma^{2}\right]$$

where $\sigma^{2} \equiv \omega_{(k)}(1 - \alpha_{(k)} - \beta_{(k)})^{-1}$ equals the unconditional variance of the

Notes to Table 2:

(a) The returns are based on 288 interpolated five minute logarithmic average bid–ask quotes for the Deutschemark–U.S. dollar spot exchange rate from October 1, 1992 through September 29, 1993. Quotes from Friday 21.00 Greenwich mean time (GMT) through Sunday 21.00 GMT have been excluded, resulting in a total of 74,880 return observations. The length of the different intraday return sampling intervals equal $5 \cdot k$ minutes. The model estimates are based on T/k non-overlapping return observations. The $\alpha_{(k)}$ and $\beta_{(k)}$ columns give the Gaussian quasi-maximum likelihood estimates for the GARCH(1, 1) parameters. Robust standard errors are reported in parentheses. The half life of a shock to the conditional variance at frequency k is calculated as $-\log(2)/\log(\alpha_{(k)} + \beta_{(k)})$ and converted into minutes. The mean lag of a shock to the conditional variance is given by $\alpha_{(k)} + \beta_{(k)} > 1$. The median lag of a shock to the conditional variance is calculated by $\frac{1}{2} + [\log(1 - \beta_{(k)}) - \log)\alpha_{(k)}) - \log(2)]/\log(\alpha_{(k)} + \beta_{(k)})$ and reported in number of minutes. For $2\alpha_{(k)} < 1 - \beta_{(k)}$ the median lag is less than $\frac{1}{2}$. The median lag is also not defined for $\alpha_{(k)} + \beta_{(k)} > 1$.

(b) The returns are based on 79,280 interpolated five minute futures transactions prices for the Standard and Poor's 500 composite index. The sample period ranges from January 2, 1986 through December 31, 1989, excluding the period from October 15, 1987 through November 13, 1987. Overnight five minute returns have also been deleted, resulting in a total of 80 intraday return observations from 08.35 through 15.15 for each of the 991 days in the sample. The length of the different intraday return sampling intervals equal $5 \cdot k$ minutes. The model estimates are based on T/k non-overlapping return observations. The $\alpha_{(k)}$ and $\beta_{(k)}$ columns give the Gaussian quasi-maximum likelihood estimates for the GARCH(1, 1) parameters. See (a) for the definition of the half life, mean lag and median lag statistics.

return innovations. The 'half-life' of the volatility process is then defined as the number of time periods it takes for half of the expected reversion back towards σ^2 to occur i.e. $-\log(2)\cdot\log(\alpha_{(k)}+\beta_{(k)})^{-1}$. Alternatively, by defining the conditional heteroskedastic squared return innovations, $\nu_{t,n}^k \equiv (\epsilon_{t,n}^k)^2-(\sigma_{t,n}^k)^2$, the GARCH(1, 1) model may be expressed as an infinite MA model for $(\epsilon_{t,n}^k)^2$ with positive coefficients, θ_i^k,

$$\epsilon_{t,n}^k=\sigma^2+\alpha(k)\sum_{i=1}^{\infty}\left(\alpha_{(k)}+\beta_{(k)}\right)^{i-1}\nu_{t,n-i}^k+\nu_{t,n}^k\equiv\sigma^2+\sum_{i=0}^{\infty}\theta_i^k\nu_{t,n-i}^k.$$

This specification suggests the corresponding 'mean lag', $\alpha_{(k)}(1-\alpha_{(k)}-2\beta_{(k)}+\alpha_{(k)}\beta_{(k)}+\beta_{(k)})^{-1}$ and 'median lag', $\frac{1}{2}+[\log(1-\beta_{(k)})-\log(\alpha_{(k)})-\log(2)]\cdot\log(\alpha_{(k)}+\beta_{(k)})^{-1}$, as additional measures for characterizing the degree of volatility persistence and the duration of the dynamic adjustment process in squared returns across the different sampling frequencies [29]. Neither the mean nor the median lag is defined for $\alpha_{(k)}+\beta_{(k)}>1$. Also, the median lag is less than $1/2$ for $2\alpha_{(k)}+\beta_{(k)}<1$.

4.3. Interpretation of the GARCH results for different return frequencies

This section summarizes the evidence from fitting standard GARCH models to the return series at different frequencies. Particular emphasis is placed on the type of distortions that may be induced by the strong periodic intraday patterns which are ignored in these models. There are a couple of indirect ways to gauge the effect. First, there are theoretical predictions about the relation between the parameters at various frequencies. If these are most obviously violated at the particular frequencies where the intraday periodicity are expected to assert the maximal impact, this is therefore consistent with the periodic pattern being a dominant source of misspecification for these models. Second, to the extent that the periodic pattern is a strictly deterministic intraday phenomenon as suggested in Section 3, the distortions should be absent from models estimated at daily or multiple-day frequencies. Consequently, if the theoretical aggregation results work satisfactorily at the multiple-day frequencies but break down intradaily then this is further evidence of a significant impact of the periodic pattern on the dynamic properties of the intraday volatility process. We also relate our findings to the prior estimates reported from intraday volatility modeling. The comparison shows that our results are fully consistent with the diverse set of estimates reported in the literature once we control for the different return frequencies employed in the studies. Finally, the explicit incorporation of the cyclical pattern in Section 5 verifies that most of the distortions attributable to the intraday volatility cycle may

[29] The mean lag is given by $\Sigma_{i=0,\infty} i\theta_i^k$, whereas the median lag, m, is implicitly defined by $\Sigma_{i=0,m}\theta_i^k=1/2\cdot\Sigma_{i=0,\infty}\theta_i^k$ (see Harvey, 1981).

be eliminated. Hence, our findings apply readily to the majority of the prior high frequency studies in the literature, and, in particular, provide an indication of the magnitude of their potential biases due to the neglect of the intraday periodicity in the volatility process.

The MA(1)–GARCH(1, 1) results for the intraday foreign exchange rates are given in Table 2a. The implied persistence measures reveal an alarming degree of irregularity across the different sampling frequencies. For the longer intraday intervals the estimates, converted into minutes, point to half lives around 18,000, or about $12\frac{1}{2}$ trading days and mean lags of around 8–9 days. However, the corresponding measures collapse at the intermediate $\frac{1}{2}$–$1\frac{1}{2}$ hour frequencies ($k = 6$–18), becoming less than 4 *hours*, only to resurrect again at the lowest, 5–10 minute, intervals ($k = 1, 2$) where violations of the $\alpha_{(k)} + \beta_{(k)} < 1$ inequality cause the estimated processes to be covariance nonstationary.

These intraday results contrast sharply with the findings for the interdaily DM–\$ returns reported in Table 3a i.e. $R \equiv \Sigma_{\tau=(t-1)k+1,tk} R_\tau$, $t = 1, 2, \ldots, [3{,}649/k]$, $k = 1, 2, \ldots, 10$ where $[\cdot]$ denotes the integer value. Here, the persistence measures appear quite consistent over the different return intervals, with the half lives and mean lags fluctuating around 20 and 15 days, respectively [30]. As for the intraday returns, the median lag is always substantially lower than the mean lag and measured with some imprecision resulting in numerous violations of the inequality governing the lower bound of the statistic, particularly for the smaller sample sizes.

A formal framework for assessment of the parameter estimates obtained at the various sampling frequencies is available from the results on temporal aggregation in ARCH models provided by Nelson (1990, 1992), Drost and Nijman (1993) and Drost and Werker (1996). Specifically, assuming that the GARCH(1, 1) model serves as a reasonable approximation to the returns process at the daily frequency, it follows from Drost and Nijman (1993) that the estimates for the corresponding weak GARCH(1, 1) models at the lower interdaily frequencies should be related to the daily parameters via the simple formula $\alpha_{(k)} + \beta_{(k)} = (\alpha_{(1)} + \beta_{(1)})^k$. This implies that the estimated half lives, when converted to a common unit of measurement as in our tables, should be stable across the frequencies [31]. Our evidence in Table 3a is in line with this prediction and it is also consistent with

[30] The intraday measures in Table 2a are converted to minutes whereas the interdaily results in Table 3a are given in days. Furthermore, recall that the weekend returns have been excluded from the intraday series. This may induce a distortion in the return dynamics but, again, our informal analysis found this effect to be inconsequential.

[31] Note that any serial dependence in the mean will generally increase the order of the implied low frequency weak GARCH model beyond that of the high frequency GARCH(1, 1) model (see Drost and Nijman (1993) for further details). However, the estimate for the MA(1) term for the daily DM–\$ GARCH(1, 1) model is only −0.034 with an asymptotic standard error of 0.018.

Table 3

k	$[T/k]$	$\alpha_{(k)}$	$\beta_{(k)}$	$\alpha_{(k)}+\beta_{(k)}$	Half life	Mean lag	Median lag
(a) Persistence of MA(1)–GARCH(1, 1) models for daily DM–$ exchange rates $R_t^k \equiv 100\cdot\sum_{\tau=(t-1)k+1,tk} R_\tau = \mu_{(k)} + \theta_{(k)}\epsilon_{t-1}^k + \epsilon_t^k\ (\sigma_t^k)^2 = \omega_{(k)} + \alpha_{(k)}(\epsilon_{t-1}^k)^2 + \beta_{(k)}(\sigma_{t-1}^k)^2,\ t=1,2,\ldots,[T/k]$							
1	3,649	0.105 (0.015)	0.873 (0.015)	0.978	31.2	37.7	23.2
2	1,824	0.150 (0.024)	0.784 (0.026)	0.934	20.6	21.4	10.9
3	1,216	0.106 (0.021)	0.813 (0.037)	0.919	24.8	21.2	6.1
4	912	0.167 (0.036)	0.713 (0.042)	0.879	21.6	19.2	6.5
5	729	0.182 (0.049)	0.611 (0.081)	0.794	15.0	11.4	<2.5
6	608	0.191 (0.049)	0.646 (0.060)	0.838	23.5	20.9	5.6
7	521	0.129 (0.049)	0.674 (0.071)	0.803	22.2	14.2	<3.5
8	456	0.170 (0.051)	0.563 (0.176)	0.733	17.8	11.7	<4.0
9	405	0.133 (0.067)	0.641 (0.230)	0.774	24.3	14.7	<4.5
10	364	0.174 (0.079)	0.434 (0.186)	0.607	13.9	7.8	<5.0
(b) Persistence of MA(1)–GARCH(1, 1) models for daily S&P 500 returns							
1	9,558	0.089 (0.019)	0.906 (0.018)	0.995	147	202	137
2	4,779	0.087 (0.015)	0.902 (0.015)	0.990	135	175	114
3	3,186	0.108 (0.016)	0.870 (0.018)	0.979	98	119	74
4	2,389	0.093 (0.016)	0.889 (0.019)	0.983	158	194	121
5	1,911	0.101 (0.015)	0.900 (0.025)	1.000	∞	∞	∞
6	1,593	0.127 (0.037)	0.838 (0.044)	0.965	117	135	79
7	1,365	0.177 (0.085)	0.776 (0.076)	0.953	100	117	69
8	1,194	0.137 (0.050)	0.821 (0.044)	0.958	129	145	83
9	1,062	0.123 (0.035)	0.805 (0.052)	0.924	84	79	33
10	955	0.173 (0.066)	0.768 (0.030)	0.941	114	127	71

(a) The returns are based on 3,649 daily quotes for the Deutschemark–U.S. dollar spot exchange rate from March 14, 1979 through September 29, 1993. Weekend and holiday quotes have been excluded. The length of the return intervals equals k days, for a total of $[T/k]$ observations, where $[\cdot]$ denotes the integer value. See Table 2a for the definition of the half life, mean lag, and median lag. These measures are converted to trading days.
(b) The returns are based on 9,558 daily observations for the Standard and poor's 500 composite index from January 2, 1953 through December 31, 1990. The length of the return intervals equals k days, for a total of $[T/k]$ observations, where $[\cdot]$ denotes the integer value. See Table 2a for the definition of the half life, mean lag, and median lag. These measures are converted to trading days.

earlier evidence for other interdaily exchange rates reported in Baillie and Bollerslev (1989).

The observations above suggest that the results for the intraday exchange rates in Table 2a are indicative of serious model misspecification. For further analysis, we again use the estimates for the daily GARCH(1, 1) model ($\hat{\alpha}_{(288)} = 0.105$ and $\hat{\beta}_{(288)} = 0.873$) as a natural benchmark since these are unaffected by the intraday periodicity. The results of Drost and Nijman (1993) and Drost and Werker (1996) now imply that the intraday returns should follow weak GARCH(1, 1) processes with $\alpha_{(k)} + \beta_{(k)}$ converging to unity and $\alpha_{(k)}$ converging towards zero as the length of the sampling interval, k, decreases. In fact, Nelson (1990, 1992) establishes general conditions under which GARCH(1, 1) models, even if misspecified at all frequencies, will satisfy the above convergence results and produce consistent estimates for the true volatility process at the highest sampling frequencies. Unfortunately, these predictions do not allow for deterministic effects in the volatility process. Yet, given the estimated standard errors, the 12 hourly through 2 hourly returns ($k = 24–144$) are roughly in line with the qualitative predictions. Beyond this point the theoretical results are strongly contradicted, however. The most blatant violations are provided by the much lower volatility persistence, as measured by $\hat{\alpha}_{(k)} + \hat{\beta}_{(k)}$, for the models based on $\frac{1}{2}$–$1\frac{1}{2}$ hourly returns ($k = 6–18$). For the 5–15 minute returns ($k \leq 3$) the sum of the estimates for $\hat{\alpha}_{(k)}$ and $\hat{\beta}_{(k)}$ is again near unity, but the relative size of the coefficients does not conform to the theoretical predictions, as $\hat{\alpha}_{(k)}$ is too large.

Our intraday results in Table 2a are not unusual. They mirror the range of estimates previously obtained in the literature over corresponding return frequencies. In particular, Engle et al. (1990) and Hamao et al. (1990) who primarily rely on returns over six hours or longer find evidence of volatility persistence that is consistent with estimates from daily data. In contrast, Baillie and Bollerslev (1991) and Foster and Viswanathan (1995), on using hourly and half-hourly returns, find much lower volatility persistence [32]. However, the volatility persistence measures appear to rebound at the higher frequencies e.g. Bollerslev and Domowitz (1993) report 5-minute GARCH(1, 1) estimates for $\alpha_{(k)} + \beta_{(k)}$ close to one but, as in Table 2a, $\hat{\alpha}_{(k)}$ seems too large. For the very highest frequencies, Locke and Sayers (1993) find that 1-minute returns generally display little volatility persistence. Conversely, Goodhart et al. (1993) detect very strong persistence in quote-by-quote data, but also find a marked decline in the persistence once information events are taken explicitly into account, illustrating how specific news arrivals may overwhelm the underlying conditional heteroskedasticity at the extremely high frequencies.

Our findings provide strong, albeit indirect, evidence in support of the conjecture that a contributing factor to the systematic variation in volatility estimates across return frequencies is the interaction between the previously well documented interdaily conditional heteroskedasticity and the intraday periodicity. For the highest frequencies the change in the intraday pattern will generally appear smooth between adjacent returns, and thus have little impact on the overall estimated degree of volatility persistence. However, as argued more formally below, the existence of short-lived intraday volatility components (in addition to the intraday periodicity) will tend to increase the dependence of $(\sigma_{t,n}^{k})^2$ on the

[32] Interestingly, Laux and Ng (1993) deviate from these studies by finding high persistence in half-hourly data for the CME currency futures. However, the futures market only operates during the most active trading in the U.S. segment of the foreign exchange interbank market and this represents a period of relative stability for the intraday volatility pattern.

lagged squared innovation, $(\epsilon_{t,n-1}^{k})^2$, relative to the overall volatility level, $(\sigma_{t,n-1}^{k})^2$, hence explaining the relatively large estimates for $\alpha_{(k)}$ at the shortest return intervals. For the intermediate $\frac{1}{2}$–$1\frac{1}{2}$ hour return models the change in the average volatility between sampling intervals will typically appear much more abrupt, resulting in significantly lower persistence measures. Beyond the 2-hour intervals the periodic pattern is averaged over a substantial part of the 24-hour trading day, and the intraday exchange rate estimates are generally closer to the implications obtained from daily models.

The results for the S&P 500 equity returns tell a similar story. The interdaily estimates in Table 3b are again broadly consistent with the a priori predictions based on the daily GARCH(1, 1) model [33]. Although the volatility persistence is higher than for the foreign exchange returns, $\hat{\alpha}_{(k)} + \hat{\beta}_{(k)}$ again displays a general smooth decline and the explicit persistence measures are fairly stable across the different return horizons. The discrepancy between the half lives, mean lags and median lags implied by the intradaily and interdaily returns are even stronger than for the foreign exchange rate data, however [34]. Moreover, the pattern in the intraday estimates for $\alpha_{(k)} + \beta_{(k)}$ reported in Table 2b is again erratic, reaching lows at the $\frac{1}{2}$-day ($k = 40$) and 20–25 minute ($k = 4, 5$) return horizons, and highs at the 40–50 minute ($k = 8, 10$) and 5-minute ($k = 1$) horizons. We conclude that the daily GARCH models conform closely to the theoretical predictions, but the strong intraday periodic patterns in volatility render the intradaily estimates highly variable and generally hard to interpret.

5. The dynamics of filtered and standardized intraday returns

This section proposes a general framework for modeling of high frequency return volatility that explicitly incorporates the effect of the intraday periodicity. The preceding section suggests that this is a prerequisite for meaningful time series analysis. Our approach is motivated by the stylized model in Section 3. While the model almost certainly is overly simplistic, the previous analysis suggests that the representation does capture the dominant features of our foreign exchange and equity return series and thus may serve as a reasonable first approximation.

[33] In this case the estimate for the daily MA(1) term equals 0.186, which is highly significant when judged by the corresponding asymptotic standard error of 0.012. Consequently the GARCH(1, 1) models for the other frequencies are, at best, approximate representations of the data generating process (see Drost and Nijman (1993) for a formal analysis).

[34] The previous footnote about the deletion of weekend exchange rate returns are even more pertinent here as both weekend and overnight equity returns are excluded. However, our informal analysis again found this to be inconsequential (see Andersen and Bollerslev, 1994).

Specifically, consider the following decomposition for the intraday returns,

$$R_{t,n} = E(R_{t,n}) + \frac{\sigma_t s_{t,n} Z_{t,n}}{N^{1/2}}, \tag{7}$$

where $E(R_{t,n})$ denotes the unconditional mean, and N refers to the number of return intervals per day. Notice that this represents a generalization of the model in Section 3, in that the periodic component for the nth intraday interval, $s_{t,n}$, is allowed to depend on the characteristics of trading day, t [35]. Given the absence of any economic theory for stipulating a particular parametric form for the intraday periodic structure, a flexible nonparametric procedure seems natural. Although no one procedure is clearly superior, the smooth cyclical patterns documented in Fig. 2a and b naturally lend themselves to estimation by the Fourier flexible functional form introduced by Gallant, 1981, 1982 [36]. In a related context, Dacorogna et al. (1993) have proposed estimating the periodicity in the activity in the foreign exchange market as the sum of three polynomials corresponding to the distinct geographical locations of the market. Returns measured on their resulting theta-time scale correspond closely to our filtered returns defined and analyzed below. However, one advantage of the approach advocated here is that it allows the shape of the periodic pattern in the market to also depend on the overall level of the volatility; a feature which turns out to be important for the equity market. Also, the combination of trigonometric functions and polynomial terms are likely to result in better approximation properties when estimating regularly recurring patterns. Furthermore, our approach for estimating $s_{t,n}$ utilizes the full time series dimension of the returns data, as opposed to simply estimating the average pattern across the trading day. Full details of the approach are provided in Appendix B. Meanwhile, it is clear from the estimated average intraday periodic patterns depicted in Fig. 6a and b, that the fitted values, $\hat{s}_{t,n}$, provide a close approximation to the overall volatility patterns in both markets. Of course, the usefulness of the procedure will ultimately depend upon the degree to which it is successful in identifying the periodic components in a temporal dimension as well. If so, the approach may serve as the basis for a nonlinear filtering procedure that could eliminate the periodic components prior to the analysis of any intraday return volatility dynamics [37].

[35] This feature is particularly important for the equity returns for which the general U-shaped volatility pattern is transformed into more of a J-shape on the highest volatility days (see Andersen and Bollerslev, 1994).

[36] This technique has previously been applied to financial return series in a different context by Pagan and Schwert (1990).

[37] This same methodology may also be used directly for prediction of future volatility over different intraday time intervals. Such intraday volatility prediction may be particularly important in the pricing and/or continuous re-balancing of hedged intraday options positions. We shall not pursue this issue any further here, however.

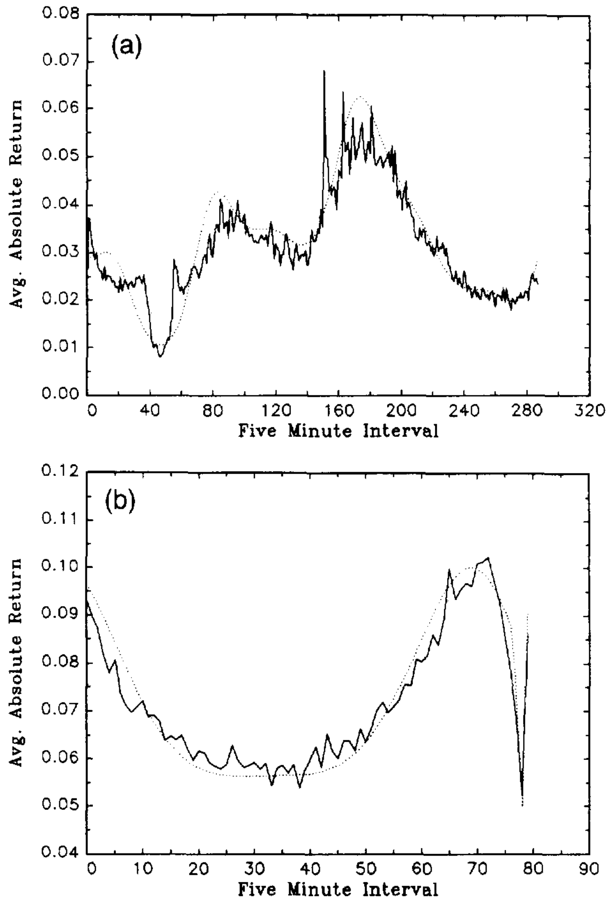

Fig. 6. Flexible Fourier functional form of intraday average absolute returns, (a) DM–$, (b) S&P 500.

5.1. Filtered foreign exchange returns

To further investigate these issues, define the filtered 5-minute return series; $\tilde{R}_{t,n} \equiv R_{t,n}/\hat{s}_{t,n}$ [38]. If the characterization of the 5-minute return series in Eq. (7) is perfect and the associated estimation error is negligible, then ignoring the impact of the weak first order return correlation, the filtered returns should conform more closely to the theoretical aggregation results for the GARCH(1, 1) model. We explicitly consider how well this hypothesis holds up, but we also keep in mind

[38] Alternatively, $\tilde{R}_{t,n}$ might also be mean adjusted. Since the mean return is practically zero, this is immaterial.

that the elimination of the main distorting effects of the intraday periodicity may bring out new features of the volatility process that were difficult to untangle prior to filtration.

First, we briefly summarize the main characteristics of the filtered series. While the mean and the standard deviation of these returns are virtually unchanged from Table 1a, both skewness and kurtosis are generally reduced by filtering the returns. For instance, the 5-minute skewness and kurtosis for $\tilde{R}_{t,n}$ equal 0.175 and 15.8, respectively. Interestingly, the evidence for negative return autocorrelations at the very highest frequencies becomes even more pronounced following the filtration, as measured by $\rho_1 = -0.090$ for the 5-minute returns. At the same time the first order absolute 5-minute return autocorrelations decline slightly to $\rho_1^A = 0.292$ [39]. The correlation structure for the absolute 5-minute filtered returns is further illustrated in Fig. 7a. The upper curves represent the correlogram for the raw returns, while the middle curves are for the filtered returns. The dramatic reduction in the periodic pattern is particularly striking for the longest lags. However, from the daily peaks in the 5 day correlogram it is clear that some periodicity remains, suggesting the presence of a stochastic periodic, or market specific, component in the intraday volatility [40]. Note also that the correlations for $|\tilde{R}_{t,n}|$ at the daily frequencies are always below the correlations for the raw absolute returns, $|R_{t,n}|$. This is consistent with the predictions from Eq. (5).

More direct evidence is provided by the estimated MA(1)–GARCH(1, 1) models reported in Table 4a. Compared to the results in Table 2a, the volatility parameters now display a much more coherent pattern across the return frequencies [41]. In accord with the theoretical aggregation results the estimates for $\alpha_{(k)} + \beta_{(k)}$ increase almost monotonically from the $\frac{1}{2}$ day to the 2 hour frequency. For the higher frequencies the theoretical predictions again begin to falter, however, although less starkly than before. The sum $\hat{\alpha}_{(k)} + \hat{\beta}_{(k)}$ remains high, but no longer increases monotonically and more importantly, $\hat{\alpha}_{(k)}$ starts to increase while $\hat{\beta}_{(k)}$ generally declines. These findings again should be interpreted in light of the absolute return autocorrelograms in Fig. 7a. The most striking feature is the initial rapid decay in the autocorrelations, followed by an extremely slow rate of decay thereafter. This pattern is not consistent with the exponential decay associated with a GARCH(1, 1) model for the 5-minute returns. Instead, these findings point to a slow hyperbolic rate of decay in the autocorrelation structure for the absolute returns, which is consistent with the presence of long-memory features in the

[39] See Andersen and Bollerslev for extensive documentation of the summary statistics for both filtered and standardized returns.

[40] The distinction between heat wave, or market specific, volatility as opposed to meteor shower, or global, volatility clustering was first discussed in Engle et al. (1990). This finding also points to the potential gains of more general dynamic periodic volatility type modeling (see e.g. Bollerslev and Ghysels, 1996).

[41] To conserve space, we do not report the half life or the mean and median lags.

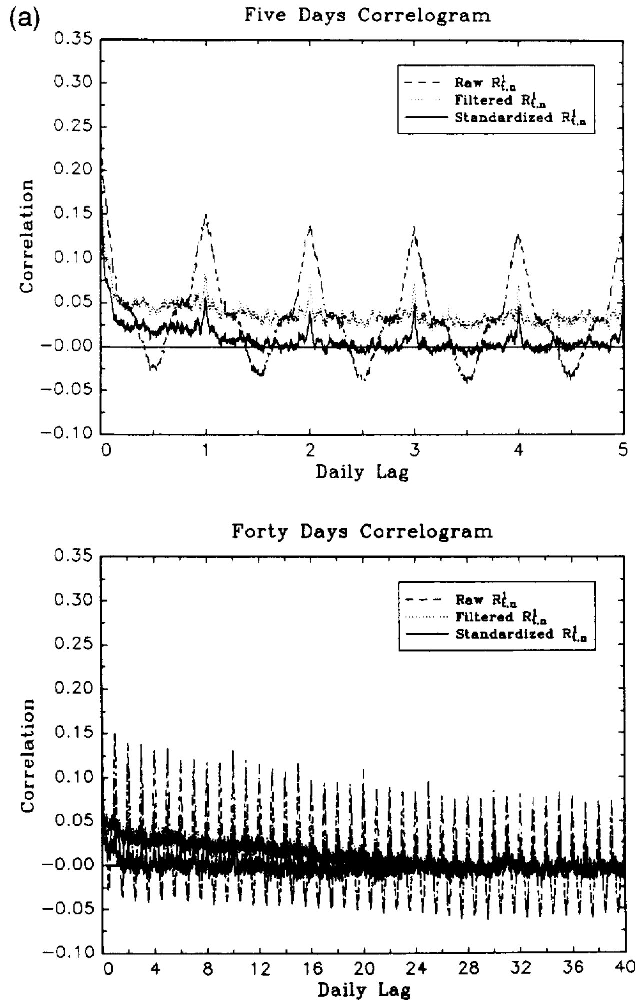

Fig. 7. Absolute returns, (a) DM–$, (b) S&P 500.

volatility process (see e.g. Baillie et al., 1996; Dacorogna et al., 1993; Ding et al., 1993). Note again the slightly higher peaks associated with the weekly frequencies.

Although the GARCH(1, 1) estimates for the high frequency filtered returns defy the theoretical predictions, the results are encouraging in terms of our ability to recover meaningful intraday volatility dynamics. In particular, by eliminating the deterministic periodicity we were able to uncover an interesting pattern in the absolute return correlogram which was largely invisible prior to the periodic filtering. A detailed investigation of the source of this phenomenon is well beyond

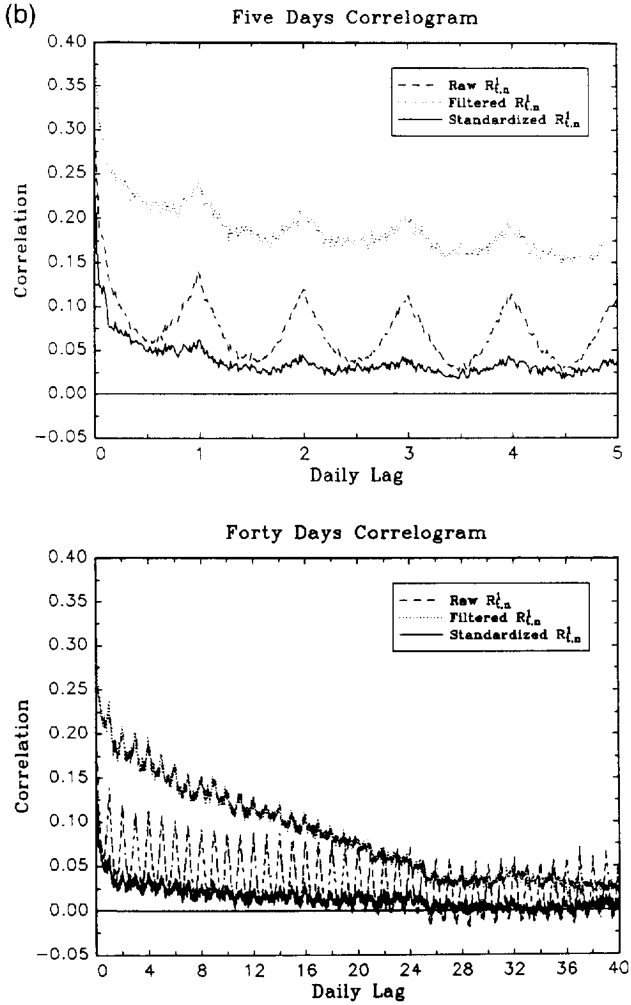

Fig. 7 (continued).

the scope of the present paper. However, we conjecture that the following factors have some impact on the observed correlation patterns. First, there may well be some positively cyclical correlated components left in the $|\tilde{R}_{t,n}|$ series, thus inducing spurious short-run dynamics in the return volatility. Second, and more importantly, the results also point to the potential importance of several distinct intraday volatility processes governed by e.g. economic announcements, the release of economic statistics, etc. each of which inherently may be of a less persistent nature than the volatility caused by changing trends in fundamental

Table 4

k	T/k	$\alpha_{(k)}$	$\beta_{(k)}$	$\alpha_{(k)}+\beta_{(k)}$
(a) MA(1)–GARCH(1, 1) models for filtered DM–$ exchange rate returns				
$\tilde{R}^k_{t,n} \equiv 100 \cdot \Sigma_{i=(n-1)k+1,nk} \tilde{R}_{t,i} = \mu_{(k)} + \theta_{(k)} \epsilon^k_{t,n-1} + \epsilon^k_{t,n}\ (\sigma^k_{t,n})^2 = \omega_{(k)} + \alpha_{(k)}(\epsilon^k_{t,n-1})^2 + \beta_{(k)}(\sigma^k_{t,n-1})^2,\ t=1, 2, \ldots, 260,\ n=1, 2, \ldots, 288/k$				
1	74,880	0.176 (0.013)	0.795 (0.016)	0.971
2	37,440	0.167 (0.012)	0.787 (0.015)	0.954
3	24,960	0.172 (0.018)	0.756 (0.025)	0.928
4	18,720	0.171 (0.019)	0.746 (0.029)	0.917
6	12,480	0.135 (0.052)	0.788 (0.096)	0.923
8	9,360	0.064 (0.051)	0.904 (0.082)	0.968
9	8,320	0.043 (0.011)	0.938 (0.016)	0.981
12	6,240	0.032 (0.011)	0.953 (0.016)	0.985
16	4,680	0.033 (0.010)	0.951 (0.015)	0.984
18	4,160	0.030 (0.009)	0.951 (0.013)	0.981
24	3,120	0.020 (0.007)	0.969 (0.009)	0.989
32	2,340	0.023 (0.006)	0.967 (0.007)	0.990
36	2,080	0.022 (0.007)	0.964 (0.007)	0.986
48	1,560	0.022 (0.008)	0.966 (0.007)	0.987
72	1,040	0.028 (0.013)	0.950 (0.009)	0.978
96	780	0.029 (0.014)	0.951 (0.009)	0.980
144	520	0.040 (0.032)	0.915 (0.016)	0.955
(b) MA(1)-GARCH(1,1) models for filtered S&P 500 returns				
$\tilde{R}^k_{t,n} \equiv 100 \cdot \Sigma_{i=(n-1)k+1,nk} \tilde{R}_{t,i} = \mu_{(k)} + \theta_{(k)} \epsilon^k_{t,n-1} + \epsilon^k_{t,n}\ (\sigma^k_{t,n})^2 = \omega_{(k)} + \alpha_{(k)}(\epsilon^k_{t,n-1})^2 + \beta_{(k)}(\sigma^k_{t,n-1})^2,\ t=1, 2, \ldots, 991,\ n=1, 2, \ldots, 80/k$				
1	79,280	0.096 (0.009)	0.892 (0.011)	0.988
2	39,640	0.086 (0.012)	0.905 (0.014)	0.991
4	19,820	0.088 (0.017)	0.904 (0.019)	0.992
5	15,856	0.071 (0.021)	0.923 (0.022)	0.994
8	9,910	0.058 (0.015)	0.937 (0.016)	0.994
10	7,928	0.058 (0.016)	0.936 (0.017)	0.994
16	4,955	0.109 (0.060)	0.869 (0.074)	0.978
20	3,964	0.084 (0.042)	0.893 (0.054)	0.977
40	1,982	0.099 (0.039)	0.873 (0.053)	0.972

(a) See Table 2a for construction of the raw return series. The method for obtaining the filtered returns, $\tilde{R}_{t,i}$, is described in the main text.
(b) See Table 2b for the construction of the raw return series. The method for obtaining the filtered returns, $\tilde{R}_{t,i}$, is described in the main text.

economic factors such as technology and productivity [42]. These distinct sources of volatility persistence could simultaneously influence the return series, resulting in a mixture distribution with different implications for the character of the short- and

[42] The distinct short-run volatility patterns induced by regularly scheduled macroeconomic announcements have been analyzed by Ederington and Lee (1993).

long-run dynamics. A promising first attempt at modeling this interaction between the volatility processes at different time resolutions within a unified framework have been suggested by Müller et al. (1995). In their so-called heterogeneous ARCH, or HARCH, model the volatility at the highest frequency is determined by the sum of numerous ARCH type processes defined over courser time intervals, where each of these components in turn may be linked to the actions of different types of traders with varying time horizons [43].

5.2. Standardized foreign exchange returns

The conjectures underlying a components type formulation of the volatility process are further reinforced by our analysis of the standardized 5-minute returns; $\hat{R}_{t,n} \equiv R_{t,n}/(\hat{\sigma}_t \hat{s}_{t,n})$. If our model provides a good approximation to the data generating process, then this series should display little ARCH effects at daily and lower frequencies, and the intraday ARCH effects should diminish. Consistent with this prediction, the absolute return autocorrelations at the lowest intraday frequencies have been reduced markedly. This is also manifest in the lower curves in Fig. 7a, which depict the correlograms for $|\hat{R}_{t,n}|$. Apart from small spikes associated with remaining stochastic periodicity at the daily frequency, the correlations for the absolute returns are generally close to zero beyond the two day lag. Thus, the daily GARCH(1, 1) volatility estimates appear to provide quite satisfactory estimates for the interday volatility dynamics [44]. At the same time, Fig. 7a is also indicative of important short-run dynamics that necessarily are unaccounted for by the daily GARCH(1, 1) volatility estimates. This again lends support to our conjecture of distinct short-run, or intraday, components in the fundamental return volatility generating process. The MA(1)–GARCH(1, 1) estimates for the standardized returns in Table 5a reinforce this interpretation by exhibiting a sharp decline in $\hat{\alpha}_{(k)} + \hat{\beta}_{(k)}$ as the return horizon increases from five minutes to one hour. In fact, beyond the one hour sampling frequency, the volatility clustering is sufficiently weak that the GARCH(1, 1) specification breaks down, and only ARCH(1) or homoskedastic MA(1) models are estimated.

5.3. Filtered equity returns

We now turn to the corresponding findings for the S&P 500 returns. In interpreting the results, it is important to recognize that the estimated intraday periodicity now involves interaction terms between the daily volatility level and

[43] Stationary conditions for this new class of time series models are developed in Darorogna et al. (1995).

[44] Of course, this apparent lack of any significant long-run correlations in the standardized returns may be due to the relatively short sample of only one year. With a longer span of data the GARCH(1, 1) model will most likely fail to capture all the low frequency dynamics (see e.g. Baillie et al., 1996).

Table 5

k	T/k	$\alpha_{(k)}$	$\beta_{(k)}$	$\alpha_{(k)}+\beta_{(k)}$
(a) MA(1)–GARCH(1, 1) models for standardized DM–\$ exchange rate				
$\hat{R}^k_{t,n} \equiv 100 \cdot \Sigma_{i=(n-1)k+1,nk} \hat{R}_{t,i} = \mu_{(k)} + \theta_{(k)} \epsilon^k_{t,n-1} + \epsilon^k_{t,n}$ $(\sigma^k_{t,n})^2 = \omega_{(k)} + \alpha_{(k)}(\epsilon^k_{t,n-1})^2 + \beta_{(k)}(\sigma^k_{t,n-1})^2$, $t = 1, 2, \ldots, 260$, $n = 1, 2, \ldots, 288/k$				
1	74,880	0.182 (0.014)	0.766 (0.021)	0.948
2	37,440	0.167 (0.015)	0.760 (0.029)	0.927
3	24,960	0.172 (0.019)	0.706 (0.037)	0.877
4	18,720	0.177 (0.016)	0.666 (0.034)	0.843
6	12,480	0.173 (0.023)	0.603 (0.047)	0.776
8	9,360	0.177 (0.028)	0.484 (0.105)	0.661
9	8,320	0.123 (0.030)	0.607 (0.143)	0.729
12	6,240	0.158 (0.028)	0.376 (0.071)	0.534
16	4,680	0.184 (0.034)	—	0.184
18	4,160	0.089 (0.026)	—	0.089
24	3,120	0.088 (0.034)	—	0.088
32	2,340	0.072 (0.031)	—	0.072
36	2,080	0.083 (0.049)	—	0.083
48	1,560	—	—	—
72	1,040	—	—	—
96	780	—	—	—
144	520	—	—	—
(b) MA(1)–GARCH(1, 1) models for standardized S&P 500 returns				
$\hat{R}^k_{t,n} \equiv 100 \cdot \Sigma_{i=(n-1)k+1,nk} \hat{R}_{t,i} = \mu_{(k)} + \theta_{(k)} \epsilon^k_{t,n-1} + \epsilon^k_{t,n}$ $(\sigma^k_{t,n})^2 = \omega_{(k)} + \alpha_{(k)}(\epsilon^k_{t,n-1})^2 + \beta_{(k)}(\sigma^k_{t,n-1})^2$, $t = 1, 2, \ldots, 991$, $n = 1, 2, \ldots, 80/k$				
1	79,280	0.095 (0.007)	0.877 (0.012)	0.973
2	39,640	0.107 (0.012)	0.841 (0.022)	0.949
4	19,820	0.127 (0.021)	0.764 (0.053)	0.890
5	15,856	0.124 (0.024)	0.765 (0.047)	0.889
8	9,910	0.117 (0.027)	0.727 (0.070)	0.843
10	7,928	0.126 (0.023)	0.681 (0.047)	0.807
16	4,955	0.215 (0.074)	0.512 (0.051)	0.727
20	3,964	0.159 (0.096)	0.551 (0.037)	0.710
40	1,982	0.250 (0.160)	—	0.250

(a) See Table 2a for construction of the raw return series. The method for standardizing the returns, $\tilde{R}_{t,i}$, is described in the main text.
(b) See Table 2b for construction of the raw return series. The standardized returns, $\tilde{R}_{t,i}$, are generated as described in the main text.

the Fourier functional form, so that not only the level but also the shape of the volatility pattern varies with σ_t. Thus, our stylized deterministic periodic model discussed in Section 3 is not strictly valid in this context i.e. generally $s_{t,n} \neq s_{\tau,n}$ for $t \neq \tau$. Counter to the results for the DM–\$ returns, this time-varying volatility component may weaken the autocorrelations for the raw absolute returns as, effectively, additional noise is injected into the returns process. Fig. 7b seem to indicate that this is indeed the case, as the correlogram for the filtered absolute returns, $|\tilde{R}_{t,n}|$, lies substantially above that for the raw series, $|R_{t,n}|$. Interestingly, this does not just occur at the 5-minute sampling frequency; the absolute return autocorrelation across the nine different intraday frequencies, as measured by ρ_1^A, $Q^A(10)$ and VR^A, are all markedly higher than the corresponding statistics for the raw returns in Table 1b [45]. This is also in line with the MA(1)–GARCH(1, 1) estimation results reported in Table 5b. The parameter estimates obtained as we move from the $\frac{1}{2}$ day to the 40-minute return horizon are again consistent with the theory. Thereafter, the sum $\hat{\alpha}_{(k)}$ and $\hat{\beta}_{(k)}$ decay slightly, but more importantly $\hat{\alpha}_{(k)}$ starts to increase. However, all the intraday estimates now consistently point towards a very high degree of volatility persistence and in all instances $\hat{\alpha}_{(k)} + \hat{\beta}_{(k)}$ are higher than the estimates for the raw return series in Table 2b. Note also, that in line with the findings for the DM–$ returns, the 5 day correlogram in Fig. 7b for the 5-minute filtered S&P 500 returns still retains a distinct periodic pattern, indicating the presence of even more complicated stochastic volatility components. Nonetheless, the simple filtering procedure again succeeds in eliminating a large proportion of the systematic intraday variation in the absolute returns and in so doing has unveiled a cleaner and starkly different picture of the volatility dynamics.

5.4. Standardized equity returns

In contrast to the results for the DM–$, the first order autocorrelations for the standardized absolute returns for the S&P 500, $|\hat{R}_{t,n}^k|$, remain highly significant for the lower intraday frequencies. Even at the $\frac{1}{2}$ day return frequency, $\rho_1^A = 0.094$ exceed the corresponding asymptotic standard error by more than a factor four. This is also confirmed by the much higher $\alpha_{(k)} + \beta_{(k)}$ estimates for the intraday GARCH(1, 1) models for $\hat{R}_{t,n}^k$ given in Table 5b. Similarly, the correlograms in Fig. 7b for the standardized returns indicate a much higher degree of volatility persistence in the 5-minute S&P 500 returns than was the case for DM–$ returns. In fact, the standardized absolute return correlogram stays mostly positive for the first 22 trading days, or about a month. This indication of more persistent volatility dynamics is likely attributable to the longer time span for the equity data. For example, Dacorogna et al. (1993) find that the absolute standardized return autocorrelations remain positive for one month when using 20-minute DM–$ data over a four year sample. Additionally, from Guillaume (1994) it is evident that our ability to detect significant long-horizon absolute return correlations is intimately linked to the length of the sampling period. Hence, although the interdaily GARCH(1, 1) model may capture a large portion of the day-to-day volatility clustering, the model's deficiency in dealing with long-memory behavior necessarily becomes more transparent when the time span of the data increases.

[45] For instance, ρ_1^A for the aggregated filtered returns, $|\tilde{R}_{t,n}^k|$, equal 0.371, 0.379, 0.341, 0.335, 0.305, 0.303, 0.292, 0.290 and 0.292 for $k = 1, 2, 4, 5, 8, 10, 16, 20$ and 40, respectively.

We conclude, that in spite of important institutional differences in the markets and the associated intradaily volatility patterns, there is strong indications that the volatility processes for the foreign exchange and the U.S. equity market share several important qualitative dynamic features. Moreover, these characteristics were largely invisible prior to our filtration of the intraday periodic structures in the high frequency return series. At the same time interesting differences between the average volatility level and volatility persistence in the two markets also emerge. These conclusions would be next to impossible to reach from the, at first sight, rather perplexing estimates obtained directly from the raw high frequency returns.

6. Concluding remarks

Our analysis of the intraday volatility patterns in the DM–$ foreign exchange and S&P 500 equity markets documents how traditional time series methods applied to raw high frequency returns may give rise to erroneous inference about the return volatility dynamics. Explicit allowance for the influence of the strong periodicity, as exemplified by our flexible Fourier form, is a necessary requirement for discovery of the salient intraday volatility features. Moreover, adjusting for the pronounced periodic structure appears critical in uncovering the complex link between the short- and long-run return components, which may help to explain the apparent conflict between the long-memory volatility characteristics observed in interday data and the rapid short-run decay associated with news arrivals in intraday data. More directly, however, our findings have immediate and important implications for a large range of issues in the rapidly growing literature using very high frequency financial data. Examples include investigations into the lead–lag relationship among returns and volatility both within and across different markets, the effect of cross listings of securities, the fundamental determinants behind the volatility clustering phenomenon, the development of real time trading and investment strategies and the evaluation of continuous option valuation and hedging decisions. Only future research will reveal the extent of the biases induced into these studies by the neglect of intraday periodic components.

Acknowledgements

We would like to thank Richard T. Bailie, the editor, an anonymous referee, Dominique Guillaume, Robert J. Hodrick, Charles Jones, Stephen J. Taylor, Kenneth F. Wallis, along with seminar participants at the Olsen and Associates Research Institute for Applied Economics, the workshop on 'Market Micro Structure' at the Aarhus School of Business, the Fall 1994 NBER Asset Pricing Meeting at the Wharton School, the HFDF-I Conference in Zürich, the 7th World

Congress of the Econometric Society in Tokyo, Duke University and the University of California at Santa Barbara for helpful comments.

Appendix A. Data description

A.1. The Deutschemark–U.S. dollar exchange rate data

The DM–$ exchange rate data consist of all the quotes that appeared on the interbank Reuters network during the October 1, 1992 through September 29, 1993 sample period. The data were collected and provided by Olsen and Associates. Each quote contains a bid and an ask price along with the time to the nearest even second. Approximately 0.36% of the 1,472,241 raw quotes were filtered out using the algorithm described in Dacorogna et al. (1993). During the most active trading hours, an average of five or more valid quotes arrive per minute; see Bollerslev and Domowitz (1993). The exchange rate figure for each 5-minute interval is determined as the interpolated average between the preceding and immediately following quotes weighted linearly by their inverse relative distance to the desired point in time. For instance, suppose that the bid–ask pair at 14.14.56 was 1.6055–1.6065, while the next quote at 14.15.02 was 1.6050–1.6055. The interpolated price at 14.15.00 would then be $\exp\{1/3 \cdot [\ln(1.6055) + \ln(1.6065)]/2 + 2/3 \cdot [\ln(1.6050) + \ln(1.6055)]/2\} = 1.6055$. The nth 5-minute return for day t, $R_{t,n}$, is then simply defined as the difference between the midpoint of the logarithmic bid and ask at these appropriately spaced time intervals. This definition of the returns has the advantage, that it is symmetric with respect to the denomination of the exchange rate. However, as noted by Müller et al. (1990), the numerical difference from the logarithm of the middle price is negligible. All 288 intervals during the 24-hour daily trading cycle are used. However, in order to avoid confounding the evidence in the correlation analysis conducted below by the decidedly slower trading patterns over weekends, all the returns from Friday 21.00 Greenwich mean time (GMT) through Sunday 21.00 GMT were excluded (see Bollerslev and Domowitz (1993) for a detailed analysis of the quote activity in the DM–$ interbank market and a justification for this 'weekend' definition). Similarly, to preserve the number of returns associated with one week we make no corrections for any worldwide or country specific holidays that occurred during the sample period. All in all, this leaves us with a sample of 260 days, for a total of 74,880 5-minute intraday return observations i.e. $R_{t,n}$, $n = 1, 2, \ldots, 288$, $t = 1, 2, \ldots, 260$.

A.2. The standard and poor's 500 stock index futures data

The intraday S&P 500 futures data are based on 'quote capture' information provided by the Chicago Mercantile Exchange (CME) from January 2, 1986 through December 31, 1989. The data specify the time, to the nearest 10 seconds and the exact price of the S&P 500 futures transaction whenever the price differs from the previously recorded price [46]. The calculation of the returns is based on the last recorded logarithmic prices for the nearby futures contract over consecutive five minute intervals. The price record covers the full trading day in the futures market from 8.30 a.m. (central standard time) to 3.15 p.m. Although, the New York Stock Exchange closes at 3.00 p.m., we retain the last three 5-minute returns from the futures market in the analysis reported on below. The first return for the trading day, i.e. from 8:30 to 8:35 a.m., constitutes another unusual time interval. This period incorporates adjustments to the information accumulated overnight, and consequently displays a much higher average return variability than any other 5-minute interval. In effect, this is not a 5-minute return, and we therefore delete it in the subsequent analysis. Alternatively, it would be possible to account for this special return interval using dummy variables. However, any such procedure is invariably ad hoc in nature. Furthermore, informal investigations reveal little sensitivity to the exact treatment of the overnight returns. We thus elect to work exclusively with the 5-minute returns. Following Chan et al. (1991), we also exclude the October 15 through November 13, 1987 time period around the stock market crash due to the frequent trading suspensions. Outside these four weeks trading suspensions were rare, but did occur. In these instances the missing prices were determined by linear interpolation, leading to identical returns over each of the intermediate intervals. This obviously smoothes the series over the missing data points which will mitigate the effect of sharp price changes subsequent to a trading suspension. Experimentation with exclusion of trading days with missing observations indicate that the findings pertaining to the degree of volatility persistence reported on here are virtually unaffected by this interpolation. All in all, these corrections result in a sample of 991 days, each consisting of 80 intraday 5-minute returns, for a total of 79,280 observations i.e. $R_{t,n}$, $n = 1, 2, \ldots, 80$, $t = 1, 2, \ldots, 991$.

Appendix B. Flexible Fourier form modeling of intraday periodic volatility components

From Eq. (7), define,

$$x_{t,n} \equiv 2\log[|R_{t,n} - E(R_{t,n})|] - \log\sigma_t^2 + \log N = \log s_{t,n}^2 + \log Z_{t,n}^2. \tag{A.1}$$

[46] We are grateful to G. Andrew Karolyi for providing us with this 5-minute price series. The same set of data has also been analyzed from a different perspective in Chan et al. (1991).

Our modeling approach is then based on a non-linear regression in the intraday time interval, n, and the daily volatility factor, σ_t,

$$x_{t,n} = f(\theta; \sigma_t, n) + u_{t,n}, \tag{A.2}$$

where the error, $u_{t,n} \equiv \log Z_{t,n}^2 - E(\log Z_{t,n}^2)$, is i.i.d. mean zero. In the actual implementation the non-linear regression function is approximated by the following parametric expression,

$$f(\theta; \sigma_t, n) = \sum_{j=0}^{J} \sigma_t^j \left[\mu_{0j} + \mu_{1j} \frac{n}{N_1} + \mu_{2j} \frac{n^2}{N_2} + \sum_{i=1}^{D} \lambda_{ij} I_{n=d_i} + \sum_{i=1}^{P} \left(\gamma_{pj} \cos \frac{pn2\pi}{N} + \delta_{pj} \sin \frac{pn2\pi}{N} \right) \right], \tag{A.3}$$

where $N_1 \equiv N^{-1} \Sigma_{i=1,N} i = (N+1)/2$ and $N_2 \equiv N^{-1} \Sigma_{i=1,N} i^2 = (N+1)(N+2)/6$ are normalizing constants. For $J = 0$ and $D = 0$, Eq. (A.3) reduces to the standard flexible Fourier functional form proposed by Gallant (1981, 1982). Allowing for $J \geq 1$ and thus a possible interaction effect between σ_t^j and the shape of the periodic pattern might be important in some markets, however. Each of the corresponding J flexible Fourier forms are parameterized by a quadratic component (terms with μ-coefficients) and a number of sinusoids (the γ- and δ-coefficients). Moreover, it may be advantageous to also include time specific dummies for applications in which some intraday intervals do not fit well within the overall regular periodic pattern (the λ-coefficients).

Practical estimation is most easily accomplished using a two-step procedure. Firstly, a generated $x_{t,n}$ series, $\hat{x}_{t,n}$, is obtained by replacing $E(R_{t,n})$ with the sample mean of the returns, $\bar{R}_{t,n}$ and σ_t with the estimates from a daily volatility model, say $\hat{\sigma}_t$. Substituting $\hat{\sigma}_t$ for σ_t and treating $\hat{x}_{t,n}$ as the dependent variable in the regression defined by Eqs. (A.2) and (A.3) allow the parameters to be estimated by ordinary least squares (OLS). Note that from Eq. (3), $\hat{\sigma}_t^2$ represents an estimate of $M(s^2)\sigma_t^2$, so that after substitution for σ_t in Eq. (A.2), the term $-\log M(s^2)$ is implicitly included in the constant term in Eq. (A.3), μ_{00}. Let $\hat{f}_{t,n} \equiv f(\hat{\theta}; \hat{\sigma}_t, n)$ denote the resulting estimate for the right hand side of Eq. (A.3) [47]. The normalization $T^{-1} \Sigma_{n=1,N} \Sigma_{t=1,[T/N]} s_{t,n} \equiv 1$, where $[T/N]$ denotes the number of trading days in the sample, then suggests the following estimator of the intraday periodic component for interval n on day t,

$$\hat{s}_{t,n} = \frac{T \cdot \exp\left(\hat{f}_{t,n}/2\right)}{\Sigma_{t=1}^{[T/N]} \Sigma_{n=1}^{N} \exp\left(\hat{f}_{t,n}/2\right)}. \tag{A.4}$$

Note that although the periodic modeling procedure is designed for fitting the

[47] Given consistent estimates for $\hat{\sigma}_t$, the resulting parameter estimates will generally be consistent. However, the use of generated regressors may result in a downward bias in the conventional OLS standard errors for the parameter estimates (see Pagan, 1984).

average volatility pattern across the N intraday intervals, the second-stage estimation of Eq. (A.3) is based on a time series regression that include all T intraday returns. Utilizing this additional information in the data rather than simply fitting the average intraday pattern, enhances the efficiency of the estimation.

The first step of our procedure involves the determination of the daily volatility factor estimates i.e. $\hat{\sigma}_t$. Given the relative success of the daily MA(1)–GARCH(1, 1) models in explaining the aggregation results for the interdaily frequencies in both markets, this appears to be a natural choice. Next, the number of interaction terms, J and the truncation lag for the Fourier expansion, P, must be determined. This is done primarily on the basis of parsimony i.e. for each of the return series we choose the model that best matches the basic shape of the periodic pattern with the fewest number of parameters. The resulting estimates for the DM–\$ returns with $J=0$ and $P=6$ are,

$$\hat{f}_{t,n}=\begin{bmatrix} \underset{(1.06)}{0.72} & -\underset{(-4.14)}{8.39}\dfrac{n}{N_1} & +\underset{(4.14)}{5.59}\dfrac{n^2}{N_2} & \\ -\underset{(-6.15)}{2.51}\cos\dfrac{2\pi n}{N} & -\underset{(-10.44)}{0.40}\sin\dfrac{2\pi n}{N} & -\underset{(-3.71)}{0.38}\cos\dfrac{2\pi 2n}{N} & +\underset{(2.70)}{0.06}\sin\dfrac{2\pi 2n}{N} \\ +\underset{(8.79)}{0.42}\cos\dfrac{2\pi 3n}{N} & -\underset{(-4.89)}{0.09}\sin\dfrac{2\pi 3n}{N} & -\underset{(-0.53)}{0.02}\cos\dfrac{2\pi 4n}{N} & +\underset{(20.48)}{0.35}\sin\dfrac{2\pi 4n}{N} \\ -\underset{(-5.38)}{0.12}\cos\dfrac{2\pi 5n}{N} & +\dfrac{0.22}{(13.35)}\sin\dfrac{2\pi 5n}{N} & -\underset{(-12.67)}{0.23}\cos\dfrac{2\pi 6n}{N} & +\underset{(0.45)}{0.01}\sin\dfrac{2\pi 6n}{N} \end{bmatrix}$$

where the numbers in parentheses indicate heteroskedastic robust t-statistics. It is evident from the corresponding fit in Fig. 6a, that this representation provides an excellent overall characterization of the average intradaily periodicity in the DM–\$ market. Consistent with Fig. 2a, the basic shape of the periodic pattern appears invariant to the daily volatility level i.e. $J=0$.

In contrast, our preferred model, the S&P 500, returns sets $J=1$ and $P=2$,

$$\hat{f}_{t,n}=\begin{bmatrix} -\underset{(-3.03)}{1.85} & -\underset{(1.62)}{3.07}\dfrac{n}{N_1} & -\underset{(-2.05)}{2.68}\dfrac{n^2}{N_2} & \\ -\underset{(-0.53)}{0.16}I_{n=d_1} & -\underset{(-1.83)}{0.62}I_{n=d_2} & +\underset{(2.99)}{1.11}I_{n=d_3} & \\ +\underset{(3.11)}{1.18}\cos\dfrac{2\pi n}{N} & -\underset{(-6.14)}{0.59}\sin\dfrac{2\pi n}{N} & +\underset{(2.94)}{0.28}\cos\dfrac{2\pi 2n}{N} & -\underset{(-2.31)}{0.14}\sin\dfrac{2\pi 2n}{N} \end{bmatrix}$$
$$+\hat{\sigma}_t\begin{bmatrix} \underset{(0.95)}{-0.54} & -\underset{(-0.98)}{1.73}\dfrac{n}{N_1} & +\underset{(1.29)}{1.57}\dfrac{n^2}{N_2} & \\ -\underset{(-0.39)}{0.11}I_{n=d_1} & -\underset{(-0.97)}{0.30}I_{n=d_2} & -\underset{(-2.02)}{0.69}I_{n=d_3} & \\ -\underset{(-1.06)}{0.37}\cos\dfrac{2\pi n}{N} & +\underset{(1.30)}{0.12}\sin\dfrac{2\pi n}{N} & -\underset{(-1.97)}{0.17}\cos\dfrac{2\pi 2n}{N} & -\underset{(-0.50)}{0.03}\sin\dfrac{2\pi 2n}{N} \end{bmatrix}$$

Although few of the coefficients in the expansion corresponding to $J=1$ are individually significant, leaving out the interaction effect results in a seemingly inferior overall fit. As seen in Fig. 2b, the volatility profile for the last fifteen minutes of trading (intervals 78, 79 and 80) shows an abrupt change from the overall smooth intraday pattern. Three dummy variables are included to minimize the distortions that may otherwise arise from this distinct period i.e. $d_1 = 78$, $d_2 = 79$ and $d_3 = 80$. The resulting fit depicted in Fig. 6b again testifies to the success of this relatively simple procedure for modeling the periodicity in intraday financial market volatility.

References

Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: Volume and price variability. Review of Financial Studies 1, 3–40.

Admati, A.R., Pfleiderer, P., 1989. Divide and conquer: A theory of intraday and day-of-the-week mean effects. Review of Financial Studies 2, 189–224.

Andersen, T.G., 1994. Stochastic autoregressive volatility: A framework for volatility modeling. Mathematical Finance 4, 75–102.

Andersen, T.G., 1996. Return volatility and trading volume: An information flow interpretation of stochastic volatility. Journal of Finance 51, 169–204.

Andersen, T.G., Bollerslev, T., 1994. Intraday seasonality and volatility persistence in foreign exchange and equity markets. Working Paper #186. Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University.

Baillie, R.T., Bollerslev, T., 1989. The message in daily exchange rates: A conditional variance tale. Journal of Business and Economic Statistics 7, 297–305.

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

Baillie, R.T., Bollerslev, T., Mikkelsen, H.O.Æ., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, forthcoming.

Barclay, M.J., Litzenberger, R.H., Warner, J.R., 1990. Private information, trading volume and stock-return variances. Review of Financial Studies 3, 233–253.

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

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

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance. Journal of Econometrics 52, 5–59.

Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. Journal of Finance 48, 1421–1443.

Bollerslev, T. and Engle, R.F., 1993. Common persistence in in conditional variances. Econometrica 61, 167–186.

Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models. In: Engle, R.F., McFadden, D. (Eds.), Handbook of Econometrics, vol.4. Elsevier Science B.V., Amsterdam.

Bollerslev, T., Ghysels, E., 1996. Periodic autoregressive conditional heteroskedasticity. Journal of Business and Economic Statistics 14, 139–151.

Bollerslev, T., Melvin, M., 1994. Bid–ask spreads and volatility in the foreign exchange market: An empirical analysis. Journal of International Economics 36, 355–372.

Bollerslev, T., Wooldridge, J.M., 1992. Quasi-maximum likelihood estimation of dynamic models with time varying covariances. Econometric Reviews 11, 143–172.

Brock, W.A., Kleidon, A.W., 1992. Periodic market closure and trading volume. Journal of Economic Dynamics and Control 16, 451–489.

Campbell, J.Y., Hentschel, L., 1992. No news is good news: An asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Chan, K., Chan, K.C., Karolyi, G.A., 1991. Intraday volatility in the stock index and stock index futures market. Review of Financial Studies 4, 657–684.

Chang, E.C., Jain, P.C., Locke, P.R., 1995. S&P 500 index futures volatility and price around the NYSE close. Journal of Business 68, 61–84.

Christie, A.A., 1982. The stochastic behavior of common stock variances: Value, leverage and interest rate effects. Journal of Financial Economics 10, 407–432.

Clark, P.K., 1973. A subordinated stochastic process model with finite variance for speculative prices. Econometrica 41, 135–155.

Darorogna, M.M., Müller, U.A., Embrechts, P., Samorodnitsky, G., 1995. Moment conditions for HARCH(k) models. Working Paper. Olsen and Associates, Research Institute for Applied Economics, Zürich, Switzerland.

Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V., 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.

Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.

Drost, F.C., Nijman, T.E., 1993. Temporal aggregation of GARCH processes. Econometrica 61, 909–927.

Drost, F.C., Werker, B.J.M., 1996. Closing the GARCH gap: Continuous time GARCH modeling. Journal of Econometrics 74, 31–57.

Ederington, L.H., Lee, J.H., 1993. How markets process information: News releases and volatility. Journal of Finance 48, 1161–1191.

Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50, 987–1007.

Engle, R.F., Ito, T., Lin, W.-L., 1990. Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market. Econometrica 58, 525–542.

Foster, F.D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance and trading costs in security markets. Review of Financial Studies 3, 593–624.

Foster, F.D., Viswanathan, S., 1995. Can speculative trading explain the volume–volatility relation?. Journal of Business and Economic Statistics 13, 379–396.

Gallant, A.R., 1981. On the bias in flexible functional forms and an essentially unbiased form: The Fourier flexible form. Journal of Econometrics 15, 211–245.

Gallant, A.R., 1982. Unbiased determination of production technologies. Journal of Econometrics 20, 285–323.

Gallant, A.R., Hsieh, D.A., Tauchen, G.E., 1991. On fitting a recalcitrant series: The pound/dollar exchange rate, 1974–83. In: Barnett, W.A., Powell, J., Tauchen, G.E. (Eds.), Nonparametric and Semiparametric Methods in Econometrics and Statistics. Proceedings of the Fifth International Symposium in Economic Theory and Econometrics. Cambridge University Press, pp. 199–240.

Gallant, A.R., Rossi, P.E., Tauchen, G.E., 1993. Nonlinear dynamic structures. Econometrica 61, 871–908.

Ghose, D., Kroner, K.F., 1994. Temporal aggregation of high frequency data. Unpublished manuscript. Department of Economics, University of Arizona.

Ghysels, E., Jasiak, J., 1994. Stochastic volatility and time deformation: An application to trading volume and leverage effects. Unpublished manuscript. C.R.D.E., Université de Montreal.

Goodhart, C.A.E., Hall, S.G., Henry, S.G.B., Pesaran, B., 1993. News effects in a high-Frequency model of the sterling–dollar exchange rate. Journal of Applied Econometrics 8, 1–13.

Granger, C.W.J., Ding, Z., 1996. Some properties of absolute returns: An alternative measure of risk. Annales d'Economie et de Statistique 40, 67–92.

Guillaume, D.M., 1994. Forecasting foreign exchange rate price changes conditional on volatility: A heterogeneous time horizon approach. Unpublished manuscript. Department of Economics, Catholic University of Leuven.

Hamao, Y., Masulis, R.W., Ng, V., 1990. Correlations in price changes and volatility across international stock markets. Review of Financial Studies 3, 281–307.

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

Harris, L., 1986b. Cross-security tests of the mixture of distributions hypothesis. Journal of Financial and Quantitative Analysis 21, 39–46.

Harris, L., 1987. Transaction data tests of the mixture of distributions hypothesis. Journal of Financial and Quantitative Analysis 22, 127–141.

Harvey, A.C., 1981. The Econometric Analysis of Time Series. Philip Allan, Oxford.

Harvey, C.R., Huang, R.D., 1991. Volatility in the foreign currency futures market. Review of Financial Studies 4, 543–569.

Ito, T., Roley, V.V., 1987. News from the U.S. and Japan: Which moves the yen/dollar exchange rate?. Journal of Monetary Economics 19, 255–277.

Jarque, C.M., Bera, A.K., 1987. Test for normality of observations and regression residuals. International and Statistical Review 55, 163–172.

Laux, P., Ng, L.K., 1993. The sources of GARCH: Empirical evidence from an intraday returns model incorporating systematic and unique risks. Journal of International Money and Finance 12, 543–560.

Ljung, G.M. and Box, G.E.P., 1978. On a measure of lack of fit in time-series models. Biometrika 65, 297–303.

Lo, A.W., MacKinlay, A.C., 1989. The size and power of the variance ratio test in finite samples: A Monte Carlo investigation. Journal of Econometrics 40, 203–238.

Lo, A.W., MacKinlay, A.C., 1990. An econometric analysis of nonsynchroneous trading. Journal of Econometrics 45, 181–212.

Locke, P.R., Sayers, C.L., 1993. Intra-day futures price volatility: Information effects and variance persistence. Journal of Applied Econometrics 8, 15–30.

Lockwood, L.J., Linn, S.C., 1990. An examination of stock market return volatility during overnight and intraday periods, 1964–1989. Journal of Finance 45, 591–601.

Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., Weizsäcker, J.E., 1995. Volatilities of different time resolutions: Analyzing the dynamics of market components. Working Paper. Olsen and Associates, Research Institute for Applied Economics, Zürich, Switzerland.

Müller, U.A., Darorogna, M.M., Davé, R.D., Pictet, O.V., Olsen, R.B., Ward, J.R., 1993. Fractals and intrinsic time: A challenge to econometricians. Working Paper. Olsen and Associates, Research Institute for Applied Economics, Zürich, Switzerland.

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

Nelson, D.B., 1990. ARCH models as diffusion approximations. Journal of Econometrics 45, 7–38.

Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: A new approach. Econometrica 59, 347–370.

Nelson, D.B., 1992. Filtering and forecasting with misspecified ARCH models I: Getting the right variance with the wrong model. Journal of Econometrics 52, 61–90.

Pagan, A.R., 1984. Econometric issues in the analysis of regressions with generated regressors. International Economic Review 25, 221–247.

Pagan, A.R., Schwert, G.W., 1990. Alternative models for conditional stock volatility. Journal of Econometrics 45, 267–290.

Percival, D.B., 1993. Three curious properties of the sample variance and autocovariance for stationary processes with unknown mean. American Statistician 47, 274–276.

Ross, S.A., 1989. Information and volatility: The no-arbitrage Martingale approach to timing and resolution irrelevancy. Journal of Finance 44, 1–18.

Scholes, M., Williams, J., 1977. Estimating betas from nonsynchronous data. Journal of Financial Economics 5, 309–327.

Son, G.S., 1991. Dealer inventory position and intraday patterns of price volatility, bid/ask spreads and trading volume. Working Paper. Department of Finance, University of Washington.

Tauchen, G.E., Pitts, M., 1983. The price variability–volume relationship on speculative markets. Econometrica 51, 485–505.

Taylor, S.J., 1986. Modeling Financial Time Series. John Wiley and Sons, Chichester.

Wasserfallen, W., 1989. Flexible exchange rates: A closer look. Journal of Monetary Economics 23, 511–521.

Wood, R.A., McInish, T.H., Ord, J.K., 1985. An investigation of transaction data for NYSE stocks. Journal of Finance 25, 723–739.

Zhou, B., 1992. Forecasting foreign exchange rates subject to de-volatilization. Working paper #3510. Sloan School of Management, MIT.