NBER WORKING PAPER SERIES

DM-DOLLAR VOLATILITY: INTRADAY ACTIVITY PATTERNS, MACROECONOMIC ANNOUNCEMENTS, AND LONGER RUN DEPENDENCIES

Torben G. Andersen
Tim Bollerslev

Working Paper 5783

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1996

We are grateful for financial support from the Institute for Quantitative Research in Finance (the Q-Group). Special thanks are also due to Olsen and Associates for making the foreign exchange quotes available. Furthermore, we have received valuable comments from seminar participants at Northwestern University, the University of Houston, the NBER Universities Research Conference on the Determination of Exchange Rates, and the Econometric Society Summer Meetings in Iowa City. We are in particular indebted to John Heaton, Bob Hodrick, Takatoshi Ito, Ken Kavajecz, Richard Lyons, and Bruce Mizrach for comments. Needless to say, we remain fully responsible for the content. This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper 5783
October 1996

DM-DOLLAR VOLATILITY: INTRADAY ACTIVITY PATTERNS, MACROECONOMIC ANNOUNCEMENTS, AND LONGER RUN DEPENDENCIES

ABSTRACT

This paper characterizes the volatility in the DM-dollar foreign exchange market using an annual sample of five-minute returns. Our modeling approach explicitly captures the pronounced intraday activity patterns, the strong macroeconomic announcement effects, and the volatility persistence, or ARCH effects, familiar from lower frequency returns. The different features are separately quantified and shown, in conjunction, to account for a substantial fraction of the realized return variability, both at the intradaily and daily levels. Moreover, we demonstrate how the high frequency returns, when properly modeled, constitute an extremely valuable and vastly underutilized resource for better understanding the volatility dynamics at the daily or lower frequencies.

Torben G. Andersen
Department of Finance
J.L. Kellogg Graduate School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
t-andersen@nwu.edu

Tim Bollerslev
Department of Economics
University of Virginia
Charlottesville, VA 22901
and NBER
bollerslev@virginia.edu

This paper provides a detailed characterization of the volatility in the DM-dollar foreign exchange market based on a one year sample of five-minute returns extracted from continuously recorded quotes on the Reuters interbank network. Our model captures the significant intraday activity patterns and the strong announcement effects associated with the regularly scheduled releases of macroeconomic statistics, along with the persistent interdaily volatility dependencies familiar from studies of returns at daily and lower frequencies. This decomposition into three distinct factors consisting of largely predictable calendar effects, announcements effects, and volatility clustering, or ARCH, effects, is fundamental to our approach, and allows for the explanation of a substantial fraction of the realized return variability, both at the extreme high frequency and at the daily level. Interestingly, the relative predictive power of the various components change dramatically with the observation frequency.

The calendar effects are naturally split into intraday, weekly, and more irregular Holiday and time change patterns. The intraday effects combine an overall smooth volatility pattern over the 24-hour cycle, reflecting the trading and economic news emanating from the financial centers around the globe, and more abrupt changes associated with the Tokyo market opening and lunch periods. The weekly features are driven strictly by the weekends. The markets virtually shut down over Saturday and Sunday, and volatility just prior to and following the weekend is subdued. Finally, regional Holidays and the introduction of Daylight Savings Time cause highly significant changes in the volatility pattern.

The announcement effects are quite distinct. They resemble one-time price adjustment processes that induce dramatic, but short-lived, bursts of volatility. We investigate the impact of all major scheduled macroeconomic releases, including both U.S. and German announcements. The investigation focuses on the significance of each individual type of announcement and the associated dynamic volatility response pattern. This is relevant for proper modeling of the volatility process, but also interesting in its own right as a measure of the significance that the market attaches to each type of announcement.

Movements in volatility at the daily level permeates intraday returns. This suggests that standard daily volatility estimates have explanatory power for intraday volatility, but also, more importantly, that intraday returns possess valuable information regarding the evolution of volatility at lower frequencies.

Our results extend the literature in several ways. In particular, the use of continuous spot market quotations for the full 24-hour trading cycle allows us to investigate the effects of all the major U.S. and German announcements, the vast majority of which have not previously been analyzed in the literature, along with various specific one-time events, e.g., the widening of the ERM band during our sample period. However, counter to the analysis in Ederington and Lee (1993) who rely on U.S. futures prices, we conclude that once the systematic intraday variability is properly accounted for, the announcement effects, although statistically significant, are of only secondary importance in explaining overall volatility. The major announcements clearly exercise a dominant role during a few predictable five-minute intervals immediately following the news release, but their explanatory power is less than that of the intraday pattern at the high frequencies, and much less than that of standard volatility forecasts at the daily level.

While we confirm previous findings regarding the qualitative features of the intraday volatility pattern, established in, e.g., Bollerslev and Domowitz (1993), Dacorogna et al. (1993), and Andersen and Bollerslev (1996a), we also enrich the characterization in several dimensions. First, only the latter study accounts for volatility movements at the daily level, and, second, none of these studies include announcement, regional Holiday, and Daylight Savings Time effects. Third, the identification of a Japanese market opening effect with a highly significant influence over half an hour is, as far as we know, new. Fourth, we explicitly control for the impact of the effective market closures around the Tokyo lunch time and the implications of gaps in the data transmission on the inference.

Our analysis also shows how intraday returns may provide valuable information for measurement of market volatility at daily and lower frequencies. As such our methodology sets the stage for more accurate analyses of longer-run volatility and price dependencies in liquid financial markets through the estimation of the latent volatility process by the cumulative sum of the intraday absolute returns rather than the absolute value of the corresponding long-run returns. While information beyond past daily returns has been used in estimating day-by-day volatility before, our results compare favorably to those from earlier studies.[1] This novel aspect of our analysis is further supported by additional evidence from other financial markets and longer calendar time spans reported in Andersen and Bollerslev (1996c).

The general framework is furthermore relevant for a number of topics that have been studied mainly through daily data, but should benefit from the additional information contained in high frequency returns. These include issues regarding information transmission and volatility spill-over, both across markets and intertemporally between the geographical regions of the identical market, see, e.g., Engle, Ito and Lin (1990), Hamao, Masulis and Ng (1992) and Hogan and Melvin (1994), as well as the relation between information flow, return volatility and variables such as bid-ask spreads or transactions volume, see Goodhart and O'Hara (1996) for a survey of this literature. Finally, continuous (real time) decision making regarding option pricing, implementation of hedges, and portfolio allocation require instantaneous

[1] The literature is large. For example, Garman and Klass (1980), Parkinson (1980), Beckers (1983), Ball and Torous (1984), Rogers and Satchell (1991), and Kunitomo (1992) explore the information content in daily observations on high and low prices, while Latane and Rendleman (1976), Day and Lewis (1992), Canina and Figlewski (1993), Jorion (1995), and Xu and Taylor (1995) study the implied volatility extracted from option prices, and Clark (1973), Epps and Epps (1976), Tauchen and Pitts (1983), Lamoureux and Lastrapes (1990a), Gallant et al. (1992), Andersen (1994, 1996), and Jones et al. (1994) investigate trading volume.

evaluation of volatility dynamics. Our approach provides a natural starting point for such endeavors.

The paper is structured as follows. Section I reports on the data sources and the construction of the five-minute return series. Section II provides a preliminary data analysis, motivating our formal modeling approach. The robust regression procedure is developed in section III; the details of this discussion may be skipped by the reader primarily interested in the qualitative findings. Section IV presents the empirical results regarding the calendar and announcement effects. It includes a quantitative assessment of the instantaneous impact of each factor as well as their overall cumulative effect. Section V provides evidence on the explanatory power of the documented volatility factors at both the intradaily and daily levels. Section VI concludes. Further details regarding the model are provided in appendices.

I. Data Sources and Construction

Our primary data set consists of five-minute returns for the Deutschemark-U.S.Dollar (DM-$) spot exchange rate from October 1, 1992, through September 30, 1993.[2] In addition, we utilize a longer daily time series of 3,649 spot DM-$ exchange rates from March 14, 1979, through September 29, 1993. The five-minute returns were constructed from the DM-$ exchange rate quotes that appeared on the interbank Reuters network over the sample period. Each quote contains a bid and an ask price along with the time to the nearest even second. At the end of each five-minute interval, we use the immediately preceding and following quote to construct the relevant bid and ask prices. The quotes are weighted by their inverse relative distance to the endpoint, and the log-price, $\log(P_{t,n})$, is then defined as the midpoint of the logarithmic bid and ask. The n'th return within day t, $R_{t,n}$, is now the change in log-prices during the corresponding period. All $N=288$ intervals during the 24-hour cycle are used. However, to avoid confounding the evidence by the decidedly slower trading patterns over the weekends, all returns from Friday 21:00 Greenwich Mean Time (GMT) through Sunday 21:00 GMT were excluded; see Bollerslev and Domowitz (1993) for an analysis of the interbank quote activity that justifies this "weekend" definition. To maintain a fixed number of returns over the span of one week, we did not remove any observations due to worldwide or country-specific Holidays, although we control explicitly for their impact in the analysis below. This leaves us with a sample of $T=260$ weekdays for a total of 74,880 five-minute return observations; i.e., $R_{t,n}$, $n=1,2,\dots,N$, $t=1,2,\dots,T$. The data set also includes all of the news headlines that appeared on the Reuters money news-alert screens. During the sample period

[2] Going to a finer sampling interval results in the bid-ask bounce effect becoming dominant, as evidenced by the increasingly significant negative sample autocorrelations reported in Guillaume et al. (1995). These findings are also consistent with the standard deviation of our 5-minute return series being slightly less than the average quoted spread; see Bollerslev and Melvin (1994).

from October 1, 1992 through September 30, 1993, a total of 105,065 such headlines appeared. These are time stamped to the second and constitute the basis for our analysis of announcement effects.

II. Preliminary Data Analysis

This section provides an initial investigation of our high frequency foreign exchange return series that serves to motivate our modeling approach. It falls naturally in three parts, corresponding to each of the general factors that we identify as important determinants of the volatility process.

A. Daily ARCH Effects

Market microstructure theories concerning the relation between information flow, return volatility and trading activity often ignore the lower frequency movements in volatility that are associated with the conditional heteroskedasticity of daily returns. This is probably related to the fact that, until recently, empirical studies have been unable to document that intraday return volatility displays characteristics that are consistent with those observed at the lower frequencies. At the face of it, this is utterly puzzling. How can the intraday return volatility process be void of ARCH features when the identical data, aggregated to the daily level, provide overwhelming evidence of conditional heteroskedasticity? An answer is provided by Andersen and Bollerslev (1996a) who demonstrate that the strong intraday volatility pattern interferes with, and garbles, the time series structure of intraday volatility. Only by explicitly modeling the intraday pattern is it possible to recover meaningful volatility dynamics. Nonetheless, the question remains as to whether the ARCH features are of secondary importance at the highest frequencies. A useful assessment is provided by the ability of standard models, based on daily returns, to forecast the variability of high frequency returns over the following day. If the volatility clustering at the daily level has little predictive value for subsequent intraday volatility then it may well be advisable to ignore ARCH effects at the daily level when studying general intraday return dynamics.

For concreteness, we explore the relation between one-step-ahead volatility forecasts generated by a MA(1)-GARCH(1,1) model of daily returns and alternative measures of return variability based on intraday data.[3] The GARCH model is estimated from daily data over the longer sample period. The associated estimates of the conditional standard deviation for each day of our high frequency sample are depicted in Figure 1. Volatility starts out at a high level, and consistently declines over the initial one

[3] While the GARCH(1,1) model is not necessarily the preferred model, it does represent a simple and popular model that provides a reasonable approximation to the second order dependency in the series, see, e.g., Baillie and Bollerslev (1989).

and a half month, followed by a more stable level over the remainder of the sample. However, even the latter period is characterized by sudden bursts of volatility that die out only gradually. Finally, there is an apparent surge in volatility at the end of the sample. This overall development is broadly consistent with the dramatic events surrounding the European Monetary System (EMS) over the sample period.[4]

To develop intuition about the properties of alternative (ex post) volatility measures, it is useful to contemplate an explicit model of intraday returns. Suppose that the exchange rate is determined by

$$d\log(P_\tau) \;=\; \mu_\tau \cdot d\tau \;+\; \sigma_\tau \cdot Dw_\tau\,,$$

where $\tau \geq 0$, W_τ is a standard Brownian motion with unit variance per day, and the instantaneous mean, μ_τ, and volatility, σ_τ, may be governed by separate stochastic processes. Much of modern asset pricing theory is cast in terms of such continuous time diffusions. In the notation for the discretely sampled intradaily returns defined above $R_{t,n} \equiv \log(P_{t+n/N}) - \log(P_{t+(n-1)/N})$, where $t=1,2,\ldots T$, and $n=1,2,\ldots,N$. In empirical applications to high frequency data, it is often assumed that the mean return is constant,

$$E(R_{t,n}) \;=\; \mu_{t+n/N} \;\approx\; \mu,$$

while allowance is made for time variation in the corresponding volatility process,

$$E(|R_{t,n}|) \;=\; \sigma_{t+n/N}\,.$$

One common approach used in the evaluation of daily (or lower) frequency volatility estimates, say $\hat{\sigma}_t$, relies on direct comparison with the corresponding realized absolute returns, ($t = 1,2,\ldots,T$),

$$|R_t| \;\equiv\; |\log(P_t) - \log(P_{t-1})|\,.$$

Studies using this approach include Cumby, Figlewski and Hasbrouck (1993), Figlewski (1995), Jorion

[4] Early September 1992, the Finnish Markkaa gave up its peg to the main European currencies, and later that month the British Pound and Italian Lira left the EMS, which limited exchange rates, through the European Rate Mechanism (ERM), to fluctuate by only 2.25% versus each other. This created intense speculation that other currencies would leave the EMS, and the volatility in October 1992 reflect the repercussions of these events in the DM-dollar market. The more dramatic episodes include the abolition of the Swedish Krona peg on 11/19, the 6% devaluation of the Peseta and the Escudo on 11/23, and the abolition of the Norwegian Krone peg on 12/10. By Christmas, this round of turmoil had been weathered, but uncertainty arose again during a speculative attack on the Irish Punt in late January. The Punt was devalued by 10% on 01/30, and the market remained unsettled for most of February. Market sentiment focused on the willingness of the Bundesbank to support the weaker currencies by loosening its monetary policy. In fact, EMS-tensions were reduced by a German interest rate cut on 02/04. Later, the Peseta and Escudo devalued again, on 05/13. This decision may have been associated with the upcoming vote, in Denmark, regarding the country's participation in the Maastricht treaty and the EMS, but the popular verdict, on 05/18, came out in favor of the treaty. The final bout of ERM-related volatility occurred during the latter three weeks of July, but came to a dramatic halt with the announcement, on 08/01, of a widening of the ERM-band from 2.25% to 15% for all currencies except the Dutch Guilder vis-a-vis the DM.

(1995), and West and Cho (1995), among others. However, realized absolute or squared daily returns are imprecise gauges of the underlying volatility. For example, the price may fluctuate rather wildly, but nonetheless end up close to the opening price, thus falsely signaling a low volatility state. A richer measure for the latent volatility might instead be based on the sum of the intradaily absolute returns, i.e.,

$$\sum_{n=1}^{N} |R_{t,n}| \ .$$

This measure will be referred to as the *cumulative absolute returns* in the discussion below.

To illustrate the potential efficiency gains associated with the latter measure, consider the extreme case where volatility remains constant within each day; i.e., $\tilde{\sigma}_\tau \equiv \sigma_{[\tau]}$ for $\tau \geq 0$, and $[\cdot]$ denotes the integer value operator. Given the distributional assumptions regarding P_τ, it follows that

$$E(|R_t|) = \tilde{\sigma}_{t-1} \cdot (2/\pi)^{1/2} \ ,$$

and

$$E\left(\sum_{n=1}^{N} |R_{t,n}|\right) = N^{1/2} \cdot \tilde{\sigma}_{t-1} \cdot (2/\pi)^{1/2} \ ,$$

which suggest the following two ex-post measures of the daily volatility

$$\hat{\sigma}_{t,1} = (\pi/2)^{1/2} \cdot |R_t| \ ,$$

and

$$\hat{\sigma}_{t,N} = N^{-1/2} \cdot (\pi/2)^{1/2} \cdot \sum_{n=1}^{N} |R_{t,n}| \ .$$

While both estimators are unbiased, the latter is vastly superior. In the present simplified setting

$$\begin{aligned} \mathrm{Var}(\hat{\sigma}_{t,N}) &= N^{-1} \cdot (\pi/2) \cdot \mathrm{Var}\left(\sum_{n=1}^{N} |R_{t,n}|\right) \\ &= N^{-1} \cdot (\pi/2) \cdot \left[N \cdot \mathrm{Var}(|R_{t,n}|) + 2 \cdot \sum_{i=1}^{N-1} (N-i) \cdot \mathrm{Cov}(|R_{t,n}|, |R_{t,n-i}|) \right] \\ &= (\pi/2) \cdot \mathrm{Var}(|R_{t,n}|) = (\pi/2) \cdot N^{-1} \cdot \mathrm{Var}(|R_t|) \end{aligned}$$

$$= N^{-1} \cdot \mathrm{Var}(\hat{\sigma}_{t,1}) \ .$$

Thus, with N=288 five-minute intraday returns, the standard deviation is reduced by a factor of close to seventeen. While the intraday volatility dynamics are much more complex than assumed above, the calculation is suggestive of the greatly improved ex-post measurement of the latent volatility process afforded by the cumulative absolute returns. As long as the discretely sampled intradaily returns are uncorrelated and the absolute returns not perfectly correlated, this intraday measure is superior; for a theoretical exposition on related issues, see Nelson (1992), and Nelson and Foster (1995, 1996).

The practical implication of the above is illustrated in Figure 2, which displays the two alternative daily volatility measures along with the GARCH forecasts from Figure 1.[5] The low correlation between the GARCH forecasts and the realized absolute daily returns is evident in Figure 2.A. The daily absolute returns are scattered almost arbitrarily around the forecasted values. In some sense this is inevitable. Daily returns are inherently noisy, and innovations are typically large relative to their expected values. Table I underscores this point. The sample correlation between the one-step-ahead GARCH volatility forecasts and two different ex-post measures of the absolute and the squared daily returns are disturbingly low, attaining a maximum of 0.107.[6] Clearly, a regression of the ex-post volatility measure on the GARCH forecasts has negligible explanatory power, with an explained variability of, at best, around $(0.107)^2 \approx 1.1\%$. Given this evidence and the inadequacies of standard ARCH models when applied directly to intraday returns, it is perhaps understandable that many studies ignore such volatility estimates.

The fallacy of this approach is, however, evident from Figure 2.B. The cumulative absolute returns are intimately related to the GARCH volatility predictions. The first two columns of Table I reinforce this conclusion. Over the annual sample, the correlation between the forecasts and the cumulate absolute returns is as high as 0.672. In other words, about $(0.672)^2 \approx 45.2\%$ of the variation in the sum of absolute intraday returns is predicted by the daily forecasts generated by a simple GARCH model.[7]

[5] Note that the GARCH volatility estimates rely solely on the *preceding* squared daily returns and the parameter estimates obtained over the longer sample. Since these parameter estimates are largely unaffected by the realization of returns over the final year of the sample, the volatility estimates are effectively one-step-ahead volatility forecasts based on prior daily returns only.

[6] The second to last column calculates returns in accordance with our definition of the trading day, from 21 GMT one day to 21 GMT the following day. The last column uses exchange rates from 12 GMT instead. The latter definition corresponds to the convention for the longer DM-$ series that underlies our MA(1)-GARCH(1,1) estimates. These timing conventions are inconsequential for the correlation measures. The results are also unaffected by the exclusion of Holidays. For instance, the correlation between $|R_t|$ and the GARCH volatility estimate for non-Holidays equals 0.082, compared to 0.086 for the full sample.

[7] This R^2 goes beyond 50% if simple adjustments are made for Holidays with predictably low return volatility.

Furthermore, this variation is at the daily level. It is impossible to explain this phenomenon by any intraday variation in volatility which is annihilated when aggregated over the entire trading day. Consequently, ignoring the information conveyed by the volatility forecasts implies that a large component of *predictable* return variability, entirely unrelated to intraday patterns or events, is excluded from the high frequency analysis. Clearly, a misleading picture may emerge if there is no control for this source of common variation across the intraday returns.

B. Calendar Effects

It is well documented that high frequency returns display pronounced intraday volatility patterns as well as other systematic calendar features such as day-of-the-week and Holiday effects. In fact, we have, at the outset, excluded weekend observations due to the effective market closure over this period. This section describes the strictly calendar related characteristics of our five-minute DM-$ return series.

While there is very little evidence of predictability in the conditional mean, the series displays pronounced intraday volatility and activity patterns.[8] Figure 3 depicts the average absolute return for each five-minute interval across all 260 weekdays in our sample. The initial observation corresponds to the interval ending at 21:00 Greenwich Mean Time (GMT), while the last observation represents the interval 20:50-20:55 GMT. Thus, our week originates Monday morning in the Pacific segment where trading is dominated by banks located in Wellington and Sydney. Trading volume and return volatility is rather subdued at this hour. There is a significant jump in (average) volatility at 0:00 GMT, or 9am Tokyo time, corresponding to the simultaneous opening of trading in a number of financial markets, including the Tokyo foreign exchange interbank market and markets for U.S. debt securities. At this point, the market must both interpret innovations to U.S. bond yields that have occurred since the close of U.S. trading, and absorb any customer orders that have accumulated overnight at authorized currency-dealing banks in Japan.[9] While Yen-$ dealings comprise the largest portion of the Asian foreign exchange market, the quotes in the Yen-$ and DM-$ markets are intimately linked through a triangular arbitrage relationship. Thus, it is perhaps not surprising that the impact of the opening of the Tokyo market resembles the market opening effects documented for equity markets by, e.g., Wood et al. (1985),

[8] There is evidence of weak negative first order autocorrelation, most likely induced by spread positioning of dealers attempting to correct inventory imbalances by posting quotes that attract customers on one side of the market only, see, e.g., Müller et al. (1990), Bollerslev and Domowitz (1993), and Zhou (1996).

[9] Prior to December 1994, The Committee of Tokyo Foreign Exchange Market Customs prohibited all authorized foreign exchange trading in Japan prior to 9am, between 12:00-1:30pm, and after 3:30pm local time, see Ito, Lyons and Melvin (1996).

and Harris (1986). Similar effects are identifiable during and following the lunch hour, 3:00-4:30 GMT, where the Tokyo segment shuts down, and the overall market typically approaches a stand-still. Ignoring the Lunch effect, we may loosely identify a u-shape pattern in volatility over the Asian segment, with the latter part leading into the European segment at 6:00 GMT. Volatility is notably higher during European trading which remains active until about 15:00 GMT. This is to be expected, as more economic events of relevance for the DM-$ rate may hit the market during this part of the trading day. Interestingly, we may again identify the rough outlines of a u-shape in the volatility pattern over this regional segment.[10] Notice, however, that the latter part of the u-shape in either case may reflect an overlap in market activity: first the Asian market coexists with the European, and later, between 12:00 and 15:00 GMT, the two most active centers trade simultaneously as it is afternoon in London and morning in New York. Finally, after the close of the London market, volatility displays a monotonic decline until it reaches the plateau associated with the Pacific segment. There are no signs of elevated volatility when trading closes down in New York. Hence, while volatility clearly increases when each of the main regional segments become active, there is no systematic evidence of enhanced volatility associated with the termination of regional trading.[11] This overall pattern is also consistent with previous evidence reported in Baillie and Bollerslev (1991), Harvey and Huang (1991), and Dacorogna et al. (1993). We now turn towards a discussion of the other systematic calendar features prevalent in high frequency returns.

Although there generally is a close coherence between the naive one-step-ahead volatility forecasts from the daily GARCH model and the cumulative absolute return volatility measure depicted in Figure 2B, there are a few dramatic deviations, most notably exemplified by the trading days 62 and 67. These are Christmas Day and New Year's Day, and both have close to zero quote activity, resulting in imputed intraday returns of near zero. Effectively, they are "weekends", as the low activity renders the intraday volatility computation meaningless. A similar, albeit weaker, manifestation of a low quoting intensity is at work on other U.S. Holidays throughout the sample. The days 41, 98, 137-138, 173, 198, and 243, representing Thanksgiving, President's Day, Easter, Memorial Day, July 4, and Labor Day, are prominent examples. There are also instances of failures in the data transmission that causes gaps of

[10] We later demonstrate that the distinct peaks, at exactly 12:30 and 13:30 GMT, are caused by price movements associated with the release of U.S. macroeconomic news at 8:30 Eastern Standard Time (EST).

[11] We conjecture that this is due to the particular structure of the interbank market. Small transaction costs, the large number of dealers and brokers, and the large size of regular transactions, force participants to continuously adjust their currency positions towards the desired level. Each dealer will thus only have a relatively small inventory imbalance that must be addressed prior to the end of trading. Lyons (1995, 1996) provides a theoretical exposition of the trading process in the interbank market.

several hours in our intraday time series. The most noticeable manifestation of this phenomenon is for day 258. The subsequent analysis explicitly controls for such spurious breaks in the volatility process.

A second type of calendar effects often recognized in high frequency returns is day-of-the-week dependencies. The apparent need to allow for such effects is illustrated in Figure 4, where a set of two-hourly dummies is estimated along with dummies for each of the weekdays. Mondays appear the least volatile, while Thursdays and Fridays are the most volatile.

Third, the GMT time scale used in Figure 3 is dubious due to the observance of Daylight Savings Time in both North America and Europe. If the daily cycles of economic activity and trading in the different regions are underlying determinants of the intraday pattern, then it should differ across the Summer Time and Winter Time regimes. Figure 5 supports this conjecture. The volatility pattern appears translated leftward by exactly one hour between 6:00 and 21:00 GMT (the European and North American segments) during the U.S. Summer Time regime.[12]

Finally, motivated by the apparent importance of market openings and closures, we also consider the possibility that volatility behaves differently in periods leading into, or out of, such market closures. In particular, we find that Friday evenings and Monday mornings appear different from the identical periods on other weekdays, and the following analysis consequently controls for both of these effects.

C. Macroeconomic Announcement Effects

Figure 6 suggests that U.S. announcements released at 8:30 EST, or 12:30 and 13:30 GMT, are the source of the previously observed volatility spikes. It displays the intraday volatility pattern for days that contained scheduled announcements on U.S. macroeconomic data, including the Employment Report, the Merchandise Trade Deficit, the Producer Price Index (PPI), Durable Goods, estimates and revisions to quarterly Gross Domestic Product (GDP), Retail Sales, Housing Starts, Leading Indicators, and Jobless Claims. It is apparent that the releases induce quite dramatic price adjustments. However, while there are signs of elevated volatility for several hours, the main impact seems to be gone within 10-20 minutes. These findings are consistent with the observation of heightened return volatility on days with macroeconomic announcements noted by, e.g., Harvey and Huang (1991) and Ederington and Lee (1993).

Table II displays the 25 largest absolute five-minute returns over the sample, and indicates whether any economic or political events may be identified as contributors to the abrupt price change.

[12] As detailed in the appendix, we effectively delete the Tokyo lunch period by artificially assigning a low return to intervals between 3:00 and 4:45 GMT, thus causing the volatility pattern to appear rectangular over this period. These observations are further "dummied" out in the formal regression analysis conducted below.

The latter exercise is, of course, subjective. Nonetheless, the evidence is striking, with the seven largest, and 15 of the 25 largest, absolute returns directly associated with the release of economic news *in the same or the immediately preceding interval.* Among other events that seemingly induced "jumps" in the DM-$ exchange rate were the "Russia Crisis", involving a military confrontation between Yeltsin and hardliners in the Russian Parliament, the plunge of the U.S. stock market on October 5, the election of Bill Clinton as the next U.S. president, and various tumultuous episodes in the ERM, including the widening of the band to 15%, and the floating of the Swedish Krona on 11/19 that culminated with a devaluation of the Peseta and Escudo the following weekend.

We conclude that scheduled releases occasionally induce large price changes, but the associated volatility shocks appear short-lived. The reason is probably their one-time character. While market participants differ in their interpretation of the news, the market typically settles on a new equilibrium price after a brief period of hectic trading, see, e.g., Goodhart and Figliuoli (1992) and Goodhart et al. (1993). This is contrary to the often more prolonged impact of unscheduled news. Examples include the Russia Crisis and the Stock Market Plunge. Each are related to three separate, large innovations, and appear to exert longer-lasting effects. Announcements may thus constitute news arrivals with a well-defined content and clearcut termination that endows them with a particularly short-lived impact, largely unrelated to the strong volatility persistence observed at the daily level. Nonetheless, they are sufficiently numerous that they induce an appreciable amount of predictable volatility in overall returns.

III. Estimating the Systematic Features of High-Frequency Volatility

The volatility dynamics of high frequency foreign exchange returns are extremely involved. There are pronounced intraday patterns, highly significant, albeit short-lived, announcement effects, and standard volatility clustering, or ARCH, effects at lower frequencies. Moreover, the latter cannot exist exclusively at the lower frequencies. They must necessarily be present in the form of highly persistent components in the intraday volatility process as well; otherwise the aggregation of intraday returns will not accommodate the persistent volatility processes at the daily level, see, e.g., the theoretical results regarding temporal aggregation of ARCH processes in Drost and Nijman (1993) and Drost and Werker (1996), and stochastic volatility processes in Andersen and Bollerslev (1996b), Ghysels et al. (1996), and Meddahi and Renault (1995). Thus, we stipulate that the volatility process is driven by the simultaneous interaction of numerous different components, some associated with economic news releases, some with predominantly predictable calendar effects, and some with highly persistent, unobserved (latent) factors. It is beyond the scope of the present paper to estimate a full-fledged time series model that accounts for the interaction of all these effects.[13] Instead, we demonstrate how formal evaluation of the effects documented in section II may be performed using a simple two-step procedure, where the final step relies on standard regression techniques.

In full generality, our model takes the following form,

$$R_{t,n} - \bar{R}_{t,n} = \sigma_{t,n} \cdot s_{t,n} \cdot Z_{t,n} \tag{1}$$

where $\bar{R}_{t,n}$ is the expected five-minute return, $Z_{t,n}$ is an i.i.d. mean zero, unit variance, error term, $s_{t,n}$ represents the calendar features as well as the scheduled announcement effects, and $\sigma_{t,n}$ denotes the remaining, potentially highly persistent volatility components, that traditionally are captured by ARCH or stochastic volatility models. All the return components are assumed to be independent, and the volatility components are non-negative; i.e., $\sigma_{t,n}$, $s_{t,n} > 0$ for all t,n.[14]

Without additional restrictions, the components of equation (1) are not separately identifiable. However, by squaring and taking logs, we may isolate the calendar and announcement effects, $s_{t,n}$, as the sole explanatory variables,

$$2\log[\,|\,R_{t,n} - \bar{R}_{t,n}\,|\,] - \log\sigma^2_{t,n} = c + 2\log s_{t,n} + u_{t,n} \tag{2}$$

where $c = E[\log Z^2_{t,n}]$, and $u_{t,n} = \log Z^2_{t,n} - E[\log Z^2_{t,n}]$. It is evident that $\log s_{t,n}$, in general, will be stochastic. Each particular release of, say, the Employment Report is unique, with the figures providing a certain innovation relative to prior consensus forecasts. The price and volatility reaction will reflect the size of this innovation (the news content), the dispersion of beliefs across traders, and probably a host of other market conditions at the time of the release. In order to capture these dynamic features directly, one must resort to explicit time series modeling based on a wider information set, including consensus forecasts, recent return innovations etc. Instead, our goal is more modest. We merely assume that the (log-) volatility response, conditional on the type of announcement, the time of the release, and other relevant calendar information, has a well-defined expected value, $E[\log s_{t,n}]$. This average impact is then governed by purely deterministic regressors. Of course, the innovation, $\log s_{t,n} - E[\log s_{t,n}]$, will

---

[13] Payne (1996) demonstrates that direct estimation of a system containing all three factors is feasible, but his stochastic volatility model accommodates only one persistent latent factor. In contrast, Andersen and Bollerslev (1996b) show that the long-run features of the five-minute DM-$ return series analyzed here are consistent with a heterogeneous information arrival interpretation of the volatility process, but only if the number of latent components, endowed with relatively strong volatility persistence, is large.

[14] We clearly lose some information by focusing strictly on a model for the imputed five-minute returns. The aim to utilize all of the "ultra-high" frequency data underlies the recent work of Engle (1996) and Engle and Russell (1996).

typically be highly correlated for the immediate period following new a release, and this will induce serial correlation and heteroskedasticity in the error term of the regression that we develop below. Similarly, we assume that log $\sigma_{t,n}$ is strictly stationary with a finite unconditional mean, E[log $\sigma_{t,n}$].

In order to obtain an operational regression equation, we impose some additional structure. Firstly, we assume that $\bar{R}_{t,n}$ is constant and well approximated by the sample mean, $\bar{R}$. This is innocuous because the standard deviation dwarfs the mean return, implying that the inference is not sensitive to minor misspecification of the conditional mean. Secondly, we utilize an a priori estimate of the return standard deviation, $\hat{\sigma}_{t,n}$, to help control for this source of systematic volatility movements. Thirdly, we impose a parametric representation on the regressor E[log $s_{t,n}$] of the form $f(\theta;t,n)$. Since theory provides no specific guidelines regarding the shape of the intraday pattern, we allow for a flexible functional form, essentially letting the data govern the specification. Our choice is the following

$$f(\theta;t,n) = \mu_0 + \mu_1 \cdot n + \mu_2 \cdot n^2 + \sum_{k=1}^{D} \lambda_k \cdot I_k(t,n) + \sum_{p=1}^{P} (\delta_{c,p} \cdot \cos\frac{p2\pi}{N} n + \delta_{s,p} \cdot \sin\frac{p2\pi}{N} n) \,, \tag{3}$$

where $I_k(t,n)$ is an indicator for the event k during interval n on day t. Apart from the dummy variables, equation (3) is identical to the flexible Fourier functional form proposed by Gallant (1981, 1982), and as such may be given a semi-nonparametric interpretation.[15]

Assembling all the pieces, we obtain the operational regression,

$$\hat{x}_{t,n} \equiv 2\log[\,|R_{t,n} - \bar{R}|\,] - \log\hat{\sigma}^2_{t,n} = \hat{c} + f(\theta;t,n) + \hat{u}_{t,n} \tag{4}$$

where $\hat{c} = E[\log Z^2_{t,n}] + E[\log \sigma^2_{t,n} - \log \hat{\sigma}^2_{t,n}]$ and, ignoring any inconsequential misspecification of the conditional mean, the error process $\{ \hat{u}_{t,n} \}$ is given by,

$$\begin{aligned} \hat{u}_{t,n} = {} & \left(\log s^2_{t,n} - E[\log s^2_{t,n}]\right) + \left(\log\sigma^2_{t,n} - \log\hat{\sigma}^2_{t,n} - E[\log\sigma^2_{t,n} - \log\hat{\sigma}^2_{t,n}]\right) \\ & + \left(\log Z^2_{t,n} - E[\log Z^2_{t,n}]\right) . \end{aligned} \tag{5}$$

The two-step procedure is now apparent. The first step requires calculating $\bar{R}$, providing a reasonable estimator of $\hat{\sigma}_{t,n}$, and specifying the exact form of the announcement dummies and lag lengths to be included in the regressors of equation (3). Thus, the first step provides the observable regressors and

[15] Allowing for the intraday pattern to depend on the overall volatility level for the day, σ_t, appears important for some markets, but was not significant in this context. The more general specification is utilized in Andersen and Bollerslev (1996a).

regressand for (4). The resulting expression constitutes a non-linear regression in the intraday time interval, n, and the event dummies, I_k. It is parameterized by a quadratic (μ-coefficients), a number of sinusoids (δ-coefficients), and dummies (λ-coefficients). It is estimated, in the second step, by ordinary least squares (OLS). Estimation efficiency is enhanced by enforcing the intrinsic periodicity of the intraday pattern (of one day).[16] We refer to equation (4) and the associated OLS procedure as the *FFF-Regression*. Clearly, this two-step method is not fully efficient, but, as argued below, given correct specification of the FFF-regressor in the first step, the parameter estimates are consistent.

The distributional properties of $\hat{x}_{t,n}$ compared to those of $R^2_{t,n}$, represents an obvious advantage of the FFF-procedure, relative to running an equivalent regression in the squared returns. The vast majority of the five-minute returns are very small, but there are instances of quite dramatic moves, usually associated with the arrival of political or economic news. Thus, the five-minute returns has a kurtosis of 21.5 compared to 4.5 for the 12-hour returns. This indicates a serious outlier problem, that is effectively eliminated by the log-transform.[17]

The statistical properties of the FFF-regression is determined by the nature of the error process, $\hat{u}_{t,n}$, in equation (5). It consists of three terms. The last is simple, as it constitutes an i.i.d. process. The first captures the discrepancy of calendar and announcement components from their expected values. Such divergences arise from stochastic components in the intraday "seasonal" or "news" innovations that differ from their expected values. As such, errors from this source are the rule rather than the exception. Nonetheless, if the mean effects are correctly specified and the errors are stationary, this does not affect the consistency of the OLS-estimator. The second term reflects potential misspecification of the estimated volatility component, $\hat{\sigma}_{t,n}$. Given the complexity of this process, it is inevitable that any preliminary estimator is misspecified, so this error term is likely heteroskedastic, serially correlated and perhaps even biased. However, any bias is absorbed in the constant, $\hat{c}$, and will not further affect inference. Moreover, as long as the regressand and the volatility process itself are stationary, this entire error

---

[16] This is done through the restriction $f(\theta;t,n) - \Sigma_1^D \lambda_k I_k(t,n) = f(\theta;t,n+288) - \Sigma_1^D \lambda_k I_k(t+1,n)$. This is, by construction, satisfied for the trigonometric terms in equation (3), so the constraint simply imposes a linear restriction on μ_1 and μ_2. The resulting estimates for μ_1 and μ_2 were always insignificant, and the estimated patterns, with or without inclusion of the quadratic terms, were indistinguishable. Hence, the quadratic was omitted from the analysis throughout by imposing $\mu_1 = \mu_2 = 0$.

[17] In fact, inspection of the regressor series, $\hat{x}_{t,n}$, now suggests a possible "inlier" problem, arising from the low values obtained when taking logs of small positive squared returns. The problem is similar to that encountered when applying the Kalman filter to log-squared returns in order to estimate stochastic volatility models, see, e.g., Harvey, Ruiz and Shephard (1994). However, we explicitly analyzed the data for the presence of unduly influential observations, following the procedure in Davidson and MacKinnon (1993), section 1.6. We also truncated the observations for $\hat{x}_{t,n}$ from below by letting all return observations in the interval (0% , 0.00036%) equal 0% (minus the sample mean) before transforming to $\hat{x}_{t,n}$. It was confirmed in both cases that the presence of inliers did not exert an appreciable impact on the estimated volatility pattern.

component is stationary. We conclude that the OLS-estimator is consistent, while the associated error process will display dependencies of unknown form. Consequently, formal inference requires the use of robust standard errors that are consistent under general heteroskedasticity and autocorrelation.

A final issue concerns the proper choice of the first stage estimator, $\hat{\sigma}_{t,n}$. A simple candidate class may be derived from standard ARCH models, fit at the daily level. For example, the GARCH(1,1) estimates in section II.A are directly applicable, if one stipulates that this volatility component is constant over the trading day. The associated intraday estimates are

$$\hat{\sigma}_{t,n} = \hat{\sigma}_t / N^{1/2}. \tag{6}$$

An alternative is to abandon the estimation of $\sigma_{t,n}$ altogether. This is equivalent to the imposition of a constant value for $\sigma_{t,n}$, e.g., letting the sample mean of the estimated σ_t be denoted $\bar{\sigma}$, we have

$$\hat{\sigma}_{t,n} = \bar{\sigma} / N^{1/2}. \tag{7}$$

In either case, we do not capture the high frequency movements in this component, but, as argued above, the consistency of the FFF-regression is retained. The advantage of the (constant) estimator (7) is that it eliminates any generated regressor problem. On the other hand, it does nothing to alleviate the heteroskedasticity. In contrast, the estimator, (6), does provide a normalization with respect to strong overall movements in volatility, which should improve the efficiency of the second step procedure.[18]

IV. Empirical Results

We report on the empirical findings in two separate sections. The first focuses on the calendar effects, and the second on the announcement effects. Notice, however, that all coefficients were estimated simultaneously, so that the full range of systematic volatility features were controlled for throughout.

A. Estimated Calendar Effects

Decisions regarding the treatment of a number of distinct features in the five-minute return series are necessary prior to estimation of the intraday pattern. We briefly outline our approach, but refer to the appendix for a full exposition. First, we observe that the extreme slowdown in market activity over some

[18] The robustness of the developed FFF-regression is worth reiterating. The nature of conditional heteroskedasticity is left unspecified, and need have to nothing to do with the preliminary estimator, $\hat{\sigma}_{t,n}$. Likewise, the distributional form for the conditional errors is unspecified, except for the existence of second order moments. General stochastic dependencies are allowed in both the calendar and announcement effects. The only caveat is a "generated regressor" problem that may arise from the first step estimates of $\hat{\sigma}_{t,n}$, which may impart a bias in our standard errors, see Pagan (1984). However, we document below that this problem is negligible in the current context, given our choice of first-step estimators.

Holidays as well as the Tokyo lunch period resemble weekends. Since we aim to characterize the overall, average volatility pattern, the systematic lack of reliable return observations over a given interval is an overriding concern. Consequently, we treat these episodes as analogous to weekends, and, effectively, eliminate them from the sample. Each Tokyo lunch period, from 12:00 to 1:45pm local time, each major Holiday, and each interval associated with a failure in data transmission, were assigned the identical low, positive return, and a dummy variable was introduced to account (perfectly) for the returns over these periods. This retains the strict periodicity in the data, while removing any impact from these episodes on the inference. Some regional Holidays involve only subdued, rather than extremely thin, quoting activity, so we introduce a "Holiday" dummy to accommodate these predictable reductions in volatility. There is also some evidence of a slowdown in the periods surrounding the weekends, i.e., early Monday morning in the Pacific zone, and late Friday afternoon in the North American segment. We accommodate them by constrained second order polynomials over the corresponding intervals, resulting in two regression coefficients for each of these periods. The Tokyo market opening effect is captured by a single coefficient that allows for a linear decay in the associated volatility burst. U.S. Daylight Savings Time induces a one-hour parallel shift in the intraday pattern over parts of the day, which is readily accommodated. However, this increases volatility in the earlier part of the day, and this is compensated by lower volatility during the now longer hiatus between the North American and Pacific segments. This is captured by a restricted second order polynomial (one free parameter) over the latter part of the day. Moreover, we incorporate day-of-the-week dummies for all weekdays except Monday. Finally, we need to select the sinusoids to be included in the seminonparametric component of equation (4). The removal of the Tokyo lunch period facilitates approximation of the intraday pattern by means of smooth functional, and we obtain an excellent fit using only four sets of sinusoids, see Andersen and Bollerslev (1996a), Payne (1996), and Kofman and Martens (1996) for earlier representations of this form.

We control for four different types of macroeconomic announcements in this section. The most influential is the Employment Report ("king of kings" among announcements (Carnes and Slifer (1991)), and it is allowed to abide by its own volatility decay rate. The other significant U.S. announcements are included as either "Category 1" (more important) or "Category 2" (less important) releases. The former incorporates GDP figures, trade balance figures, and durable goods purchases, while the latter contains the PPI, retail sales, housing starts, leading indicators, initial jobless claims, factory orders, and German M3-figures. Finally, releases following the biweekly Bundesbank meeting had a major impact on the market, so this effect was also treated separately. Each type of announcement effect is summarized by a single regression coefficient. Interpretation of these point estimates is discussed in the next section.

The estimation results for the full system, using the first step estimator (6), are recorded in the second column of Table III. All coefficients associated with the intraday pattern are highly significant, except for the last sine term.[19] As mentioned, the volatility slowdown over the latter part of the Summer days compensates for increased activity earlier in the day due to Daylight Savings Time. The strong market opening in Tokyo is noteworthy, while pronounced announcement and Holiday effects were expected. In contrast, there is no evidence of a Monday morning effect, once the other calendar effects are taken into account, and the Friday afternoon effect is at best borderline significant. Similarly, there is no indication of a day-of-the-week effect. Although the Friday coefficient is large when judged by the conventional OLS standard errors, the effect is likely an artefact of specific events that happened to occur on Fridays. When evaluated against the robustified standard error the effect is decidedly insignificant.

The above results justify estimation without the day-of-the-week dummies. These estimates are given in Table III, column three. The only qualitative difference is that the Friday afternoon effect now is insignificant at the 5% level. As a last robustness check, we estimated the identical system, imposing the constant daily volatility factor, (7). The results in Table III, column four, confirm that the parameter estimates are largely unchanged, and the qualitative features of the inference unaffected. Thus, the inclusion of $\hat{\sigma}_t$ does not seem to give rise to a practical generated regressors problem.

The intraday volatility pattern, as dictated by the estimates in column three, Table III, are displayed in Figure 7. Both the Tokyo opening effect, and the increased volatility during the overlap in the Asian and European, and subsequently, the European and North American segments are apparent. The Monday morning and Friday afternoon effects also manifest themselves clearly, in spite of being marginally insignificant. The excellent overall fit is evident from Figure 8, which displays the predicted and average absolute realized five-minute returns in the FFF-dimension underlying the estimation.

Arguably, the corresponding fit in the absolute return dimension is a better gauge of the success of the model. To convert the FFF-pattern into absolute returns, note that equations (1) through (4) imply

$$|R_{t,n} - \bar{R}| \quad = \quad N^{-1/2} \cdot \hat{\sigma}_t \cdot \exp(f(\theta;t,n)/2) \cdot \exp(\hat{u}_{t,n}/2) \; . \qquad \textbf{(8)}$$

One-day-ahead intraday forecasts, conditional on $\hat{\sigma}_t$, may therefore be generated by taking the conditional expectation in equation (8), and evaluating $f(\cdot;t,n)$ at the estimated $\hat{\theta}$. If we ignore a potential correlation between $\hat{\sigma}_t$ and the transformed error term, we simply have to obtain the unconditional expectation,

[19] The large differences between the heteroskedasticity and autocorrelation consistent and the OLS standard errors signify the importance of accounting for the effects of the strong volatility clustering and outlying observations when conducting the inference.

E[$\exp(\hat{u}_{t,n}/2)$], which in turn may be estimated by averaging the corresponding expression, $\exp(\hat{u}_{t,n}/2)$, over the relevant residuals in the sample.[20] The unconditional volatility forecasts may be obtained in identical fashion, except that $\bar{\sigma}$ would be used in place of $\hat{\sigma}_t$. This resulting (unconditional) pattern is displayed in Figure 9, and contrasted to the actual average absolute returns. While the more pronounced sensitivity to outliers renders the actual average pattern somewhat jagged, the overall fit is very good.

We end the section by assessing the economic significance of the estimated FFF-coefficients. The point estimates associated with regular dummy variables in equation (3) are readily interpreted. For example, a coefficient of unity is tantamount to the addition, to equation (8), of a multiplicative factor of $\exp(1/2) \approx 1.65$. Thus, volatility for the corresponding interval increases by about 65 percent. Consequently, the Holiday factor amounts to $\exp(-0.712/2) \approx 0.700$, or a reduction in volatility of about 30 percent. This effect applies uniformly to each interval covered by the Holiday dummy. Beyond less important U.S. "Holidays" such as Veterans day and weekdays between Christmas and New Year, these also include regional Holidays in Tokyo, Wellington, Sydney, and London.

Assessment of the remaining calendar and announcement effects is more complicated because the regressors are not simple indicators, but involve prespecified dynamic response patterns. In particular, assuming that event k impacts volatility over N_k intervals, the implied set of regressors are,

$$\sum_{i=0}^{N_k} \lambda(k,i) \cdot I_k(t,n-i) \ .$$

If the announcement affects volatility for an hour or two, there are 13 or 25 separate event-specific coefficients to estimate. Given the limited number of occurrences of each event and the inherent noise in the returns process, this is highly inefficient. Instead, we impose a reasonable decay-structure on the volatility response pattern, and simply estimate the degree to which the event "loads onto" this pattern, by imposing $\lambda(k,i) = \lambda_k \cdot \gamma(i)$, $i = 0, 1, \ldots, N_k$, where $\gamma(i)$ dictates the prespecified pattern. Hence, $\exp(\lambda_k \cdot \gamma(0)/2)$ signifies the immediate response of the absolute returns, while the response at the i'th lag equals $\exp(\lambda_k \cdot \gamma(i)/2)$. The corresponding cumulative response measure is naturally defined by,

$$M(k) = \sum_{i=0}^{N_k} \left[\exp\left(\frac{\lambda_k \cdot \gamma(i)}{2} \right) - 1 \right] \ . \tag{9}$$

[20] Note that any correlation would enhance the predictive power of the daily volatility factor. Thus, the assessment of the explanatory power provided by $\hat{\sigma}_t$ in this context may be deemed a conservative estimate of its true predictive value.

This nonlinear function of the event-specific loading coefficient, λ_k, reflects the impact over the entire response horizon expressed as a multiplicative factor scaled in units of average volatility per interval over the period. The Tokyo market opening, for example, has an immediate response coefficient of 0.65 and a cumulative response measure of 2.12, implying that volatility jumps by 65% at 9am Tokyo time, while more than twice the usual volatility of a five-minute interval is added over the span of the half hour response horizon. However, volatility is low at this point in the trading cycle, averaging about 0.025% per interval, so the full impact is only around 0.053%. Since the (median) cumulative absolute return is about 9% over our sample, this constitutes less than 0.6% of the return variability for a typical day. Although the effect is pronounced and robust - and market observers and traders clearly recognize it - it is thus arguably of limited overall economic importance. A similar calculation shows the economic significance of the early Monday effect to be negligible. While the first few intervals have an estimated 14% reduction in volatility, the total effect amounts to about 0.38% of the daily cumulative absolute returns. In contrast, the late Friday slowdown may exert a considerable effect. Due to Daylight Savings Time, separate estimates are obtained for Summer and Winter, but the reduction in volatility over the last interval of the day is 31% in both cases, with a cumulative impact of 3% to 4% at the daily level.

B. Estimated Announcement Effects

This section reports on our estimation of the volatility responses associated with regularly scheduled macroeconomic announcements in the U.S., Germany and Japan. Extensive experimentation revealed the qualitative features of the average volatility impact to be remarkably similar across most of the announcements, and well approximated by third order polynomials constrained to reach zero at a one hour horizon. In order to allow for simultaneous estimation of the multiple effects, we adopted this pattern as a universal format for the $\gamma(i)$ sequence. The announcements load onto the pattern in accordance with the logic underlying equation (9), except that few releases, notably the Employment Report, follow elongated versions, so that their response horizons go beyond one hour. Apart from this, the only source of variation across the estimated response patterns is the announcement-specific loading coefficients, λ_k.

A summary of the results may be based on the point estimates for the announcement coefficients in Table III, column three. The table includes all releases that were determined to be highly significant, with category I and II consolidating those that have about the same response patterns. The estimated average effects take the form displayed in Figure 10. For comparison, the figure also includes an estimated response pattern for the period following the widening of the ERM-band. Since this decision arguably had the same one-shot character as regularly scheduled announcements, with no additional information to be released subsequently, the market response should share the qualitative features of the standard announcement responses. The remaining estimates are invariant to the inclusion of this event.

The relative size of the response patterns in Figure 10 is as expected. The revision of the EMS-band was a major event, and the large and prolonged volatility response is no surprise. The ranking of the regular announcements reflects the fact that they were presorted according to apparent significance. More interesting is the size of the estimated effects. The coefficients displayed in Figure 10 represent $\lambda(k) \cdot \gamma(i)$ for $i = 0, 1, \ldots, 32$ where, e.g., $\gamma(0) = 2.18869$ and λ_k is given in Table III. For example, the contemporaneous response to an Employment Report is governed by $\lambda_k \cdot \gamma(0) = 1.746 \cdot 2.18869 = 3.822$. From equation (10), this is tantamount to a multiplicative impact on the absolute return of $\exp(3.822/2) = 6.76$, or an instantaneous jump in volatility of about 576%. The corresponding cumulate response from equation (9) amounts to 27.17. Since a conservative estimate of the expected absolute return during 8:30-10:30 EST, absent announcement effects, is around 0.05% per five-minute interval, the overall effect is an elevation of volatility by $27.27 \cdot 0.05\% = 1.3585\%$. Therefore, we find about a 15% (= 1.3585/9 %) average increase in the cumulate absolute return for trading days that contain a scheduled Employment Report. Analogous calculations reveal that the instantaneous volatility nearly triples for category I announcements, almost doubles for category II announcements, and jumps by almost 400% following Bundesbank Meetings, while the cumulative impact represent an increase in the daily cumulative absolute returns of about 3.6%, 2.0%, and 5.1%, respectively. Of course, the Bundesbank Meetings are biweekly rather than monthly, so their overall impact, at least over this sample period, is estimated at close to 2/3 of that of the U.S. Employment Report. Likewise, categories I and II represent multiple monthly releases, so their combined impact is substantial. We again stress that these estimates represent average, or expected, responses. The most surprising releases were associated with a much larger impact. This point is exemplified by the ERM-Band widening, which ranks eleventh on the list of large return innovations in Table II. The event is estimated to have raised the instantaneous absolute five-minute return by 3,000%, and to have increased the cumulative absolute return on August 2, 1993, by 36.6%. There are announcements within each of the four categories that are associated with even larger immediate responses than the ERM-band correction. Thus, some scheduled announcements induce truly spectacular bursts of volatility, although the responses, on average, are decidedly less pronounced.

While the strict categorization in Table III is adequate for general characterization of the announcement effects, it is evident that the categories cover quite diverse events. In order to convey more direct information regarding the importance of each individual type of release, Table IV reports loading coefficients for all U.S. and German announcements investigated in the study. These were obtained by treating each announcement in the manner afforded the Employment Report and Bundesbank Meetings in Table III, i.e., we control for all remaining significant announcements while estimating the marginal impact of the release under investigation. All statistically significant releases are listed in Table IV.A, and ranked according to their estimated impact on the cumulative absolute returns. The results largely confirm our earlier findings. Indeed, the first twelve announcements are the ones controlled for throughout in our estimation procedure. The set of significant U.S. releases also corresponds closely to those identified by Ederington and Lee (1993, 1995a, 1995b) and Payne (1996). Of course, we would expect the relative importance of the releases to differ across markets. For instance, Ederington and Lee (1993) and Jones et al. (1995) find the PPI figures to be almost as important as the employment report for U.S. bond market volatility; see also Goodhart et al. (1993) and DeGennaro and Shrieve (1995) for analyses of high frequency news effects in the U.S. Dollar - British Pound and U.S. Dollar - Japanese Yen foreign exchange markets, respectively.[21] The overwhelming significance of the two German monetary announcements is also interesting, especially in light of the fact that none of the corresponding U.S. monetary announcements have any explanatory power; for evidence pertaining to earlier periods see, e.g., Hardouvelis (1984), Goodhart and Smith (1985), Hakkio and Pearce (1985), and Thornton (1989). However, it is worth recalling the intense scrutiny of German monetary policy over the sample period due to the frictions in the EMS.[22] Moreover, U.S. monetary policy was unusually uncontroversial over the period. For example, there were no changes in the Fed Funds rate over the sample. Only confirmation of our results over a longer sample period will allow us to gauge the robustness of these particular findings. Nonetheless, the results are consistent with the emphasis that the Bundesbank allegedly places on monetary targets as guidelines for its policy decisions.

Table IV.A also includes the only Japanese release of any significance, namely the Japanese GNP figures. Since there are only four annual releases of this statistic, the announcement is of limited overall import, but the statistical significance is noteworthy. The directional response of the exchange rate is consistent with a strengthening of the dollar on positive innovations to the Japanese GNP. It suggests an interpretation that stresses the U.S.-Japanese trade imbalance. Strong growth in Japan would be conducive to imports from the U.S., and a shrinkage of the overall U.S. deficit vis-a-vis Japan. However, the small sample precludes any firm conclusions. For comparison purposes the table also

---

[21] In a related context, Eddelbüttel and McCurdy (1996) find that a simple frequency count of the news headlines on the Reuters screen is positively related to the intradaily DM-$ volatility. Similar correlation measures between equity volatility and headline news-counts are developed in Mitchell and Mulherin (1994), and Berry and Howe (1994).

[22] For a recent discussion of the Bundesbank monetary policy rules see Clarida and Gertler (1996).

includes the German GDP figures. These are estimated to be of about the same economic importance as the Japanese GNP numbers, although the effect is not statistically significant at the 5% level.

The remaining 22 U.S. and 20 German news releases were all individually insignificant, but the mere fact that a majority of the coefficients are positive (15 versus 7 and 11 versus 9, respectively) suggests that on average these announcements contribute positively to the DM-$ volatility, although the economic impact in most instances is negligible. The complete listing is given in Tables IV.B and IV.C.

The qualitative importance of the announcement effects is perhaps best illustrated by observing that they "explain" the significance of weekday dummies. A number of previous studies have noted the importance of allowing for day-of-the-week effects when modelling daily exchange rate movements, see, e.g., McFarland et al. (1982, 1987), So (1987), Hsieh (1988, 1989), and Baillie and Bollerslev (1989). Upon running the FFF-regression, using the volatility estimates from (6), including all calendar effects but *excluding the announcement effects*, we obtain the following coefficients on the weekday dummies for Tuesday through Friday, -0.038 [-0.54] (-1.26), 0.038 [0.53] (1.24), 0.118 [1.68] (3.85), and 0.165 [2.19] (5.09), where the square brackets provide robust t-statistics and the parentheses report standard OLS t-statistics. Thus, ignoring the announcement effects produces economically large day-of-the-week effects, with Friday having estimated excess absolute returns on the order of 8.5%, and Thursday of 6.0%. Moreover, the effect is highly significant based on the conventional heteroskedasticity adjusted OLS standard errors, and the Friday effect remains significant at the 5% level when judged against fully robust standard errors. Of course, Table III demonstrates that this result vanishes, if we account for the announcement effects. The large Thursday and Friday dummies reflect the clustering of scheduled news releases on these weekdays. Given the estimates in Table IV, a back-of-the-envelope calculation indicates the magnitude of the involved effects. For example, Fridays contain all 12 Employment Report releases, as well as 4 trade balance, 2 housing start, 5 CPI, 3 retail sales, 5 PPI, 5 business inventories, 3 durable goods, 3 GDP, 3 factory orders, 5 industrial output/capital utilization, one leading indicator, one jobless claims, and one Bundesbank meeting releases over the year. In total, this increases the average cumulative absolute returns on Fridays by about 5.4%. The unexplained gap of about 3.1% is small, and certainly consistent with random variation. In fact, from Table II it is evident that the most influential releases of PPI, retail sales, and durable goods figures happen to occur on Fridays, and, in addition, there are two distinct episodes of "ERM turmoil" on this weekday. Hence, the enhanced volatility on Fridays is readily "explained", which is consistent with the message obtained from the robust inference. A similar analysis applies to Thursdays. The Bundesbank meetings typically take place on this weekday, resulting in 23 releases. This combined with 51 jobless claims, 6 trade balance, 6 factory orders, 5 retail sales, 4 PPI, 3 CPI, 3 GDP, and 2 housing start releases plus numerous minor announcements explains an average elevation of volatility on the order of 5.0%. The residual 1.0% is comparable to the implied variation across the first three weekdays, and clearly insignificant.

Consequently, there is no evidence of a day-of-the week effect. The implication is that volatility forecasts based on such dummies are biased. For instance, if there are no scheduled announcements on a Friday, forecasts will tend to be inflated by about 7-8%, while volatility for a Friday containing just an Employment Report release, on average, will be underestimated by the same magnitude. If additional announcements are scheduled for the same Friday the downward bias in the forecast is further aggravated.

In summary, macroeconomic announcements have a large impact when they hit the market, with the largest 5-minute returns over the entire sample readily being identified with such public releases. Clearly, for sensible inference around these periods, it is necessary to control for this effect. However, the induced bursts of volatility are short-lived. As such, the overall significance of these announcements for volatility at the daily level is tenuous. In fact, the majority of the releases induce an average excess cumulative absolute returns of around, or less than, 5% of that for a typical trading day. Only the employment report is associated with a substantially higher impact. The following section explores the relative significance of the various effects documented above for explaining the overall volatility.

V. The Relative Importance of the Volatility Components at Different Frequencies

Different market participants are concerned with different features of the volatility process. Market makers, brokers, and money managers engaging in continuous trading or the implementation of dynamic portfolio and hedging strategies are exposed to short-run volatility, and consider information on this dimension vital. Conversely, more passive investors are mostly concerned with lower frequency movements. Likewise, research into the price mechanism or other market microstructure issues focuses on the extreme high-frequency movements, while standard asset pricing models typically are specified and tested at daily or lower frequencies. While the higher and lower frequency characteristics cannot be entirely independent, the difference in perspective will lead to rather wide discrepancies in the assessment of the economic significance of the factors that we have explored above.[23] This section formally evaluates the impact of each component at both the extreme high frequency and the daily level.

The FFF framework allows for a direct assessment of the joint and marginal predictive power

[23] Discrepancies in the investment horizons across different types of traders underlies the motivation behind the Heterogeneous ARCH, or HARCH, class of models proposed by Müller et al. (1996).

achieved by each of the three separate components; the daily volatility factor, the calendar effects, and the announcement effects. The indicator variable, I_σ, is unity if the daily (ARCH-based) volatility factor is included in the construction of a given forecast, as in equation (6), and zero if the forecast is based on a constant daily volatility factor, as in equation (7). Formally, the associated component takes the form,

$$\hat{\sigma}_{t,n} = \hat{\sigma}_t \cdot I_\sigma + \bar{\sigma} \cdot (1 - I_\sigma).$$

Likewise, the indicators, I_c, I_a, and I_h, signify whether calendar, announcement, and Holiday effects are accounted for in the construction of the forecast. The calendar effects, f_c, include the impact of the FFF-sinusoids, the Tokyo open, the Daylight Savings Time, and the early Monday, and late Friday regressors. The announcement effects, f_a, signify the contribution of the four announcement regressors from Table III, while the holiday effects, f_h, refer to the predicted reduction in volatility associated with the Holiday dummy and the control for missing observations. The latter effects were incorporated in all forecasts, so that this source of predictable return variability does not interfere with the interpretation of the results. Based on the set of FFF-regressors in equation (3), it is now straightforward to construct a volatility forecast from equation (8), and identify the unique contribution from each of the three remaining sources of systematic variation. In particular, letting the vector of indicator variables, $I = (I_\sigma, I_c, I_a)$, identify a given model configuration, the set of one-day-ahead absolute return interval forecasts is calculated as,

$$v(I;t,n) = c_0 \cdot \hat{\sigma}_{t,n} \cdot \exp\left[\frac{\hat{f}_c(t,n) \cdot I_c + \hat{f}_a(t,n) \cdot I_a + \hat{f}_h(t,n) \cdot I_h}{2} \right] \tag{10}$$

where $\hat{f}_c(t,n) = f_c(\hat{\theta};t,n)$ and so forth, and $\hat{\theta}$ is estimated conditional on the current variant of the model, as indicated by I. Thus, the parameters are allowed to vary across the designs in a manner that maximizes the explanatory power of the specific components for each configuration.

Table V provides the fraction of the total variation in absolute returns explained by each forecast. These are given as the R^2 from the following regressions of realized cumulative absolute returns, $(t = 1, \ldots, T)$,

$$\sum_{n=1}^{N} |R_{t,n} - \bar{R}| = b_0 + b_1 \cdot \sum_{n=1}^{N} v(I;t,n) + \epsilon_t, \tag{11}$$

and realized five-minute absolute returns, $(t = 1,\ldots,T;\; n = 1,\ldots,N)$,

$$|R_{t,n} - \bar{R}| \quad = \quad b_0 \ + \ b_1 \cdot v(I;t,n) \ + \ \epsilon_t \, . \tag{12}$$

on the corresponding volatility forecasts.

The results are telling. Consider the first data column in Table V, that refers to the degree of explained variation in daily cumulative absolute returns. The complete model accounts for an impressive 60.6% of the total variation. Moreover, this number drops only slightly if we remove the announcement or the calendar components from the forecast. In contrast, the explained variation drops precipitously when the daily volatility factor is omitted. In fact, the benchmark explanatory power, provided by the Holiday effects alone (8%), is only marginally improved by incorporating calendar effects (8.3%) and only slightly improved when allowing for announcement effects (11.4%). In contrast, the daily GARCH volatility factor alone explains 57.8%. The message is clear. The daily volatility forecasts capture broader movements in volatility that generally are independent of calendar effects. This is perhaps not surprising given that the intraday pattern, which accounts for the majority of these features, is annihilated when aggregated to the daily level. More striking is the marginal impact of the announcements at the daily level.[24] Evidently, the conclusion of Ederington and Lee (1993) that announcement effects provide most of the explanatory power for return variability over the trading day is grossly misleading for the current sample period, that is characterized by large fluctuations in the overall level of volatility.[25] Finally, we note that the calendar effects pick up additional explanatory power if day-of-the-week effects are allowed, but remain less important than the announcement effects. This is fully consistent with the weekday dummies substituting (imperfectly) for the news releases.

Turning to the last column of Table V, we find that the explanatory power of the components is reversed when we consider the high-frequency return variability. The overall explained variation has dropped to 15.9%, but more revealing is the fraction explained by the calendar effects alone (8.1%) relative to the announcement effects (4.9%) and the daily volatility factor (3.4%). In other words, the intraday pattern accounts for the majority of the explained variation, while the announcements are sufficiently influential, in spite of the relatively few intervals they affect, that they also exert an appreciable impact, and, finally, the overall predictable movements in daily volatility have only a limited, although not negligible, impact at the five-minute return level.

---

[24] This is consistent with French and Roll (1986) who argue that public information releases account for little of the daily variability in U.S. equity returns.

[25] Subsample analysis reveals that the explained variation drops considerably when the overall level of volatility is more stable. However, the ranking of the effects remains the same across all investigated subsamples.

VI. Concluding Remarks

The volatility process of the interbank DM-$ spot exchange rate market is quite involved. Entirely new phenomena become visible as one proceeds from daily returns to high frequency intraday returns. Nonetheless, it is possible to identify three sets of characteristics that govern the main systematic features of the process. At the high frequency level, the pronounced intradaily volatility pattern is dominant. It accounts for an average variation in the absolute returns of more than 250% across the 24-hour trading cycle (after exclusion of the Tokyo lunch period). The magnitude of this effect overwhelms the predictable day-by-day changes in volatility captured by, e.g., ARCH models, which rarely changes by more than 25% over any 24 hour period. In addition, strong, but short-lived, announcement effects are eminently prevalent at the very highest frequencies.

Our analysis documents that the high frequency calendar and announcement effects may be estimated with a reasonable degree of precision, even without accounting for the broad movements in the daily volatility. This is possible because dramatic overall changes in the level of volatility will affect each interval over the 24-hour cycle in a similar manner. Likewise, if the announcements are distributed evenly over the entire sample, then the average effects should be captured fairly accurately. On the other hand, it is evident that analysis of one-time events must account not only for the intradaily pattern and the possible release of economic or political news within the event window, but also for the overall level of volatility.

Furthermore, when analyzing the broader economic implications of the identified factors, it is evident that the daily volatility factor dominates at frequencies around and lower than one day. Thus, it might be argued that the intraday pattern and announcement effects are of lesser importance, and that the high frequency data are of limited interest outside the area of market microstructure. This conclusion is highly misleading, however. Perhaps the most significant finding to emerge from our study is that the high frequency returns contain extremely valuable information for the measurement of volatility at the daily level. The cumulative absolute returns provide a vastly superior ex-post measure of the underlying daily latent volatility factor than either absolute or squared daily returns. As such, the results encourage the development of new and improved techniques for the estimation and prediction of daily or lower frequency volatility that explicitly incorporates the information in the intraday return observations. Moreover, the intraday returns may provide new insights that are of critical importance for the understanding of the lower frequency return dynamics.

To illustrate the latter point, Figure 11 displays the correlogram for the raw absolute five-minute returns, $|R_{t,n} - \bar{R}|$, as well as the corresponding filtered absolute returns, $\hat{s}_{t,n}^{-1} \cdot |R_{t,n} - \bar{R}|$. The former, depicted in Figure 11.A, is dominated by the strong periodicity at the daily frequency and does not appear particularly informative.[26] Figure 11.B, in contrast, features a strictly positive and slowly declining correlogram. While spikes are visible at the daily frequencies, they are minor and do not distort the overall pattern. This may be interpreted as a testimony to the relative success of our model for $s_{t,n}$ in capturing the systematic calendar and announcement effects. The regularity of the correlogram in Figure 11.B compares favorably to those of similarly filtered absolute returns presented in Andersen and Bollerslev (1996a) and Payne (1996). The main point, however, is that the pronounced hyperbolic rate of decay in the absolute return autocorrelations is indicative of so-called long-memory in the volatility process. This is not consistent with the rate of decay implied by ordinary ARCH models, but points towards a fractionally integrated volatility process, as proposed by Baillie et al. (1996) in the ARCH framework, and Harvey (1994) and Breidt et al. (1995) within the context of stochastic volatility models. Specifically, the estimate of the so-called degree of fractional integration, or d, implied by the fitted hyperbolic decay in Figure 11 equals 0.387, which is in close accordance with the estimates obtained by semi-parametric frequency domain methods in Andersen and Bollerslev (1996b) and Henry and Payne (1996). Thus, the long-memory characteristics appear inherent to the return series, as they manifest themselves, even over relatively short time-spans. This suggests that the source of fractional integration in the volatility is related to the data generating process itself, rather than induced by infrequent structural shifts as suggested by Lamoureux and Lastrapes (1990b). Thus, once we have accounted for the predictable calendar and announcement effects, the high frequency data provide important evidence on the plausibility of two alternative hypotheses that appear observationally equivalent from the perspective of lower frequency returns. Indeed, we expect the information provided by high frequency returns to become increasingly valuable to a broad range of issues in financial economics, both within and beyond the realm of market microstructure.

[26] Both series are adjusted for missing observations, so that, e.g., the Tokyo lunch hour is removed. The daily periodicity is even more pronounced when the lunch time observations are retained.

References

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, No. 1, 169-204.

Andersen, T.G. and T. Bollerslev (1996a), "Intraday Periodicity and Volatility Persistence in Financial Markets," forthcoming ***Journal of Empirical Finance***.

Andersen, T.G. and T. Bollerslev (1996b), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long Run in High-Frequency Returns," Manuscript, Department of Finance, J.L Kellogg Graduate School of Management, Northwestern University.

Andersen, T.G. and T. Bollerslev (1996c), "ARCH Models Provide Good Volatility Forecasts: Answering the Critics," Mimeo, Department of Finance, J.L Kellogg Graduate School of Management, Northwestern University.

Baillie, R.T. and T. Bollerslev (1989), "The Message in Daily Exchange Rates: A Conditional Variance Tale," ***Journal of Business and Economic Statistics***, 7, 297-305.

Baillie, R.T. and T. Bollerslev (1991), "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates," ***Review of Economic Studies***, 58, 565-585.

Baillie, R.T., T. Bollerslev, and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," ***Journal of Econometrics***, 74, 3-30.

Ball, C.A. and W.N. Torous (1984), "The Maximum Likelihood Estimation of Security Price Volatility: Theory, Evidence and Applications to Option Pricing," ***Journal of Business***, 57, 97-112.

Beckers, C.E. (1983), "Variances of Security Price Returns Based on High, Low and Closing Prices," ***Journal of Business***, 56, 97-112.

Berry, T.D. and K.M. Howe (1994), "Public Information Arrival," ***Journal of Finance***, 49, 1331-1346.

Bollerslev, T. and I. Domowitz (1993), "Trading Patterns and Prices in the Interbank Foreign Exchange Market," ***Journal of Finance***, 48, 1421-1443.

Bollerslev, T. and M. Melvin (1994), "Bid-Ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis," ***Journal of International Economics***, 36, 355-372.

Breidt, F.J., N. Crato, and P. de Lima (1995), "On the Detection and Estimation of Long Memory in Stochastic Volatility," Manuscript, Department of Statistics, Iowa State University.

Canina, L. and S. Figlewski (1993) "The Informational Content of Implied Volatility," ***Review of Financial Studies***, 6, 659-681.

Carnes, W.S. and S.D. Slifer (1991), ***The Atlas of Economic Indicators: A Visual Guide to Market Forces and the Federal Reserve***, HarperCollins Publishers.

Clarida, R. and M. Gertler (1996), "How the Bundesbank Conducts Monetary Policy," Manuscript, Department of Finance, New York University.

Clark, P.K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," ***Econometrica***, 41, 135-155.

Cumby, R., S. Figlewski and J. Hasbrouck (1993), "Forecasting Volatility and Correlations with EGARCH Models," ***Journal of Derivatives***, Winter.

Dacorogna, M.M., U.A. Müller, R.J. Nagler, R.B. Olsen and O.V. Pictet (1993), "A Geographical Model for the Daily and Weekly Seasonal Volatility in the Foreign Exchange Market," ***Journal of International Money and Finance***, 12, 413-438.

Davidson, R. and J.G. MacKinnon (1993), ***Estimation and Inference in Econometrics***, Oxford University Press.

Day, T.E. and C.M. Lewis (1992), "Stock Market Volatility and the Information Content of Stock Index Options," ***Journal of Econometrics***, 52, 267-287.

DeGennaro, R.P. and R.E. Shrieves (1995), "Public Information Releases, Private Information Arrival, and Volatility in the Foreign Exchange Market," Manuscript, Department of Finance, University of Tennessee.

Drost, F.C. and T.E. Nijman (1993), "Temporal Aggregation of GARCH Processes," ***Econometrica***, 61, 909-927.

Drost, F.C. and B.J.M. Werker (1996), "Closing the GARCH Gap: Continuous Time GARCH Modeling," ***Journal of Econometrics***, 74, 31-57.

Eddelbüttel, D. and T.H. McCurdy (1996), "The Impact of News on Foreign Exchange rates: Evidence from Very High Frequency Data," Manuscript, Department of Economics, Queen's University.

Ederington, L.H. and J.H. Lee (1993), "How Markets Process Information: News Releases and Volatility," ***Journal of Finance***, 48, 1161-1191.

Ederington, L.H. and J.H. Lee (1995a), "The Creation and Resolution of Market Uncertainty: The Impact of Macroeconomic Announcements on Implied Volatility," Manuscript, Department of Finance, University of Oklahoma.

Ederington, L.H. and J.H. Lee (1995b), "Volatility Prediction: ARCH, Announcement, and Seasonality Effects," Manuscript, Department of Finance, University of Oklahoma.

Engle, R.F. (1996), "The Econometrics of Ultra-High Frequency Data," Manuscript, Department of Economics, University of California, San Diego.

Engle, R.F., T. Ito and W.-L. Lin (1990), "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market," ***Econometrica***, 58, 525-542.

Engle, R.F. and J.R. Russell (1996), "Forecasting the Frequency of Changes in Quoted Foreign Exchange Prices with the ACD Model," ***Journal of Empirical Finance***, forthcoming.

Epps, T.W. and M.L. Epps (1976), "The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distributions Hypothesis," ***Econometrica***, 44, 305-321.

Figlewski, S. (1995), "Forecasting Volatility Using Historical Data," Manuscript, Department of Finance, New York University.

French, K.R. and R. Roll (1986), "Stock Returns Variances: The Arrival of Information and the Reaction of Traders", ***Journal of Financial Economics***, 17, 5-26.

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., P.E. Rossi and G.E. Tauchen (1992), "Stock Prices and Volume," ***Review of Financial Studies***, 5, 199-242.

Garman, M.B. and M.J. Klass (1980), "On the Estimation of Security Price Volatilities from Historical Data," ***Journal of Business***, 53, 67-78.

Ghysels, E., A. Harvey and E. Renault (1996), "Stochastic Volatility," forthcoming in G.S. Maddala (ed.), ***Handbook of Statistics***, Vol. 14, ***Statistical Methods in Finance***, North Holland, Amsterdam.

Goodhart, C.A.E. and L. Figliuoli (1992), "The Geographical Location of the Foreign Exchange Market: A Test of the 'Island' Hypothesis," ***Journal of International and Comparative Economics***, 1, 13-28.

Goodhart, C.A.E., S. Hall, S. Henry and B. Pesaran (1993), "News Effects in a High Frequency Model of the Sterling-Dollar Exchange Rate," ***Journal of Applied Econometrics***, 8, 1-13.

Goodhart, C., T. Ito, and R. Payne (1996), "One Day in June 1993: A Study of the Working of Reuters' Dealing 2000-2 Electronic Foreign Exchange Trading System," in ***The Microstructure of Foreign Exchange Markets*** (J. Frankel, A. Galli and A. Giovannini, eds.), University of Chicago Press.

Goodhart, C.A.E. and M. O'Hara (1996), "High Frequency Data in Financial Markets: Issues and Applications," forthcoming ***Journal of Empirical Finance***.

Goodhart, C.A.E. and R. Smith (1985), "The Impact of News on Financial Markets in the United Kingdom," ***Journal of Money, Credit and Banking***, 17, 507-511.

Guillaume, D.M., M.M. Dacorogna, R.R. Davé, U.A. Müller, R.B. Olsen and O.V. Pictet (1995), "From the Bird's Eye to the Microscope: A Survey of New Stylized Facts of the Intradaily Foreign Exchange Markets," Manuscript, Olsen & Associates, Zürich, Switzerland.

Hakkio, C. and D. Pearce (1985) "The Reaction of Exchange Rates to Economic News," ***Economic Inquiry***, 23, 621-635.

Hamao, Y., R.W. Masulis and V. Ng (1990) "Correlations in Price Changes and Volatility across International Stock Markets," ***Review of Financial Studies***, 3, 281-307.

Hardouvelis, G. (1984) "Market Perceptions of Federal Reserve Policy and the Weekly Monetary Announcements," ***Journal of Monetary Economics***, 14, 225-240.

Harris, L. (1986) "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns," ***Journal of Financial Economics***, 16, 99-117.

Harvey, A.C. (1994), "Long-Memory in Stochastic Volatility," Manuscript, London School of Economics.

Harvey, A.C., E. Ruiz and N. Shephard (1994), "Multivariate Stochastic Variance Models," ***Review of Economic Studies***, 61, 247-264.

Harvey, C.R. and R.D. Huang (1991), "Volatility in the Foreign Currency Futures Market," ***Review of Financial Studies***, 4, 543-569.

Henry, M. and R. Payne (1996), "An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility," Manuscript, Financial Markets Group, London School of Economics.

Hogan, K. and M.T. Melvin (1994), "Sources of Heat Waves and Meteor Showers in the Foreign Exchange Markets," ***Journal of International Economics***, 37, 239-247.

Hsieh, D.A. (1988), "The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983," ***Journal of International Economics***, 24, 129-145.

Hsieh, D.A. (1989), "Modeling Heteroskedasticity in Daily Foreign Exchange Rates," ***Journal of Business and Economic Statistics***, 7, 307-317.

Ito, T., R.K. Lyons, and M.T. Melvin (1996), "Is There Private Information in the FX Market? The Tokyo Experiment," Manuscript, Haas School of Business, University of California, Berkeley.

Ito, T. and V.V. Roley (1987), "News From U.S. and Japan: Which Moves the Yen/Dollar Exchange Rate?" ***Journal of Monetary Economics***, 19, 255-277.

Jones, C., G. Kaul, and M.L. Lipson (1994), "Transactions, Volume and Volatility," ***Review of Financial Studies***, 7, 631-651.

Jones, C., O. Lamont and R. Lumsdaine (1995), "Public Information and the Persistence of Bond Market Volatility," Manuscript, Department of Economics, Princeton University.

Jorion, P. (1995), "Predicting Volatility in the Foreign Exchange Market," ***Journal of Finance***, 50, 507-528.

Kofman, P. and M. Martens (1995), "Interaction Between the London and New York Stock Markets During Common Trading Hours," Manuscript, Department of Econometrics, Monash University, Australia.

Kunitomo, N. (1992), "Improving the Parkinson Method of Estimating Security Price Volatilities," ***Journal of Business***, 65, 295-302.

Lamoureux, C.G. and W.D. Lastrapes (1990a), "Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects," ***Journal of Finance***, 45, 221-229.

Lamoureux, C.G. and W.D. Lastrapes (1990b) "Persistence in Variance, Structural Change and the GARCH Model," ***Journal of Business and Economic Statistics***, 8, 225-234.

Latane, H.A. and R.J. Rendleman (1976), "Standard Deviations of Stock Price Ratios Implied in Option Prices," ***Journal of Finance***, 31, 359-381.

Lyons, R. (1995), "Tests of Microstructural Hypotheses in the Foreign Exchange Market," ***Journal of Financial Economics***, 39, 321-351.

Lyons, R. (1996), "A Simultaneous Trade Model of the Foreign Exchange Hot Potato," forthcoming ***Journal of International Economics***.

McFarland, J.W., R.R. Pettit and S.K. Sung (1982), "The Distribution of Foreign Exchange Price Changes: Trading Day Effects and Risk Measurement," ***Journal of Finance***, 37, 693-715.

McFarland, J.W., R.R. Pettit and S.K. Sung (1987), "The Distribution of Foreign Exchange Price Changes: Trading Day Effects and Risk Measurement - A Reply," ***Journal of Finance***, 42, 189-194.

Meddahi, N. and E. Renault (1995), "Aggregations and Marginalizations of GARCH and Stochastic Volatility Models," Manuscript, GREMAQ.

Mitchell, M.L. and J.H. Mulherin (1994), "The Impact of Public Information on the Stock Market," ***Journal of Finance***, 49, 923-950.

Müller, U.A., M.M. Dacorogna, R.D. Davé, R.B. Olsen, O.V. Pictet, and J.E. Von Weizsäcker (1996), "Volatilities of Different Time Resolutions: Analyzing the Dynamics of Market Components," forthcoming ***Journal of Empirical Finance***.

Müller, U.A., M.M. Dacorogna, R.B. Olsen, O.V. Pictet, M. Schwarz and C. Morgenegg (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. (1992), "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," ***Journal of Econometrics***, 52, 61-90.

Nelson, D.B. and D. Foster (1995), "Filtering and Forecasting with Misspecified ARCH Models II: Making the Right Forecast with the Wrong Model," ***Journal of Econometrics***, 67, 303-335.

Nelson, D.B. and D. Foster (1996), "Continuous Record Asymptotics for Rolling Sample Variance Estimators," ***Econometrica***, 64, 139-174.

Newey, W.K. and K.D. West (1987), "A Simple, Positive Semi-Definite, Heteroskedasticity Consistent Covariance Matrix," ***Econometrica***, 55, 703-708.

Pagan, A.R. (1984) "Econometric Issues in the Analysis of Regressions with Generated Regressors," ***International Economic Review***, 25, 221-247.

Parkinson, M. (1980) "Extreme Value Method for Estimating the Variance of the Rate of Return," ***Journal of Business***, 53, 61-65.

Payne, R. (1996) "Announcement Effects and Seasonality in the Intraday Foreign Exchange Market," Manuscript, Financial Markets Group, London School of Economics.

Rogers, L.C.G. and S.E. Satchell (1991) "Estimating Variance from High, Low and Closing Prices," ***Annals of Applied Probability***, 1, 504-512.

So, J.C. (1987), "The Distribution of Foreign Exchange Price Changes: Trading Day Effects and Risk Measurement - A Comment," ***Journal of Finance***, 42, 181-188.

Tauchen, G.E. and M. Pitts (1983) "The Price Variability-Volume Relationship on Speculative Markets," ***Econometrica***, 51, 485-505.

Thornton, D. (1989), "The Effect of Unanticipated Money on the Money and Foreign Exchange Markets," ***Journal of International Money and Finance***, 8, 573-587.

West, K.D. and D. Cho (1995), "The Predictive Ability of Several Models of Exchange Rate Volatility," ***Journal of Econometrics***, 69, 367-391.

Wood, R.A., T.H. McInish and J.K. Ord (1985), "An Investigation of Transaction Data for NYSE Stocks," ***Journal of Finance***, 25, 723-739.

Xu, X. and S.J. Taylor (1995), "Conditional Volatility and the Informational Efficiency of the PHLX Currency Options Market," ***Journal of Banking and Finance***, 19, 803-821.

Zhou, B. (1996), "High Frequency Data and Volatility in Foreign Exchange Data," ***Journal of Business and Economic Statistics***, 14, 45-52.

Appendix

A.1 Regional trading segments, Holidays, and data gaps

We begin by formally defining the regional trading segments. This classification is used to assign dummy variables to the intervals affected by regional Holidays. The observance of Daylight Savings Time in Europe and North America, at periods that do not fully coincide, induce us to operate with four separate categories. Furthermore, the classification is not exhaustive, in the sense that there are periods which do not belong to any specific regional segment. This is immaterial since it is only used to specify periods that are affected significantly by regional Holidays in one of the market centers. The following daily Greenwich Mean Time (GMT) trading zone definitions are used:

	Year Round
Wellington (New Zealand):	20:55-22:00
Sydney (Australia):	20:55-00:00
Tokyo (Japan):	00:00-06:00

	09/27-10/23	**10/26-03/26**	**03/26-04/02**	**04/05-09/24**
London (Europe):	07:00-15:00	07:00-16:00	06:00-15:00	06:00-15:00
Europe-N.America Overlap:	11:30-15:00	12:30-16:00	12:30-15:00	11:30-15:00
New York (N.America):	11:30-20:30	12:30-20:30	12:30-20:30	11:30-20:30

Regional Holidays affect the entire trading segment, except for certain minor U.S. Holidays, where an appreciable drop in quoting and trading activity only takes place after the London market closes. The following Holiday periods were identified from the quote intensity as well as the Reuter's news tape.

	Dates	**Time Period**	**Occasion**
United States	10/12	11:30-20:30	Columbus Day
	11/11	16:00-20:30	Veterans Day
	11/26	12:30-20:30	Thanksgiving
	12/21-01/01	All Day	Christmas/New Year
	01/18	16:00-20:30	King's Birthday
	02/15	12:30-20:30	President's Day
	04/08	15:00-20:55	Easter Begins
	04/09	All Day	Easter
	04/12	20:55-20:30	Easter Ends
	05/30	11:30-20:30	Memorial Day
	07/05	11:30-20:30	July 4
	09/06	11:30-20:30	Labor Day

Tokyo - Dates: 11/03, 11/23, 01/15, 02/11, 04/29, 05/03, 09/15, 09/23

Wellington - Dates: 10/26, 01/25, 06/07

Sydney - Dates: 10/05, 01/26, 04/26, 06/14

London - Dates: 05/03, 05/30, 08/30

We also checked for slowdowns associated with regional Holidays in a number of additional countries, including Hong Kong, Taiwan, Singapore, Germany, and Switzerland, but no clear signs of an effect could be detected, so these Holidays were not included in the analysis.

All five-minute intervals, covered by the Holiday periods listed above, were assigned one of two different dummies. The "Holiday"-dummy refers to periods of reduced activity, where reliable returns may nonetheless be obtained. An interpretation is that this corresponds to lower levels of general economic activity, where less relevant economic news are generated. The "Market Closure"-dummy refers to periods where the quoting intensity is so low as to render return calculations unreliable. Among the above Holidays, the following are allocated to the latter "Market Closure" category:

	Dates	**Time Period**	**Occasion**
Market Closures:	10/12	15:00-20:30	Columbus Day
	11/26	16:00-20:30	Thanksgiving
	12/22	20:30-20:55	Christmas
	12/23-12/25	All Day	Christmas
	12/28	21:00-23:00	Christmas
	12/31	17:00-20:55	New Year
	01/01	All Day	New Year
	02/15	16:00-20:30	President's Day
	04/08	20:30-20:55	Easter
	04/09	All Day	Easter
	04/12	20:55-20:30	Easter
	05/30	06:00-20:30	Memorial Day
	07/05	11:30-20:30	July 4
	09/06	11:30-20:30	Labor Day

The trading restrictions in Japan over the sample period precludes reliable assessment of the properties of the return series over the local lunch period. It effectively corresponds to a "weekend" in the midst of the trading day. Formally, we define a market closure each day during

Tokyo Lunch-Time: 03:00-04:45

Finally, we identified some apparent failures in the data transmission which result in lengthy gaps in the quote series. All of the affected intervals were treated as market closures. The specific periods are:

	Dates	Time Period
Data Gaps:	10/21	01:18-05:37
	10/28-29	22:16-01:15
	11/17	01:30-05:39
	12/16	01:15-05:12
	01/08	00:33-06:20
	02/10	01:35-06:27
	02/22	04:52-06:40
	05/21	16:41-21:00
	09/26-27	21:57-06:07

The market closures present a modeling dilemma, since we want to eliminate these observations, but also want to retain the strict periodicity associated with the intradaily and weekly features of the high frequency return series. We solve this by artificially assigning a very low, positive return (standardized by an overall daily volatility factor) to all these intervals, and then removing (zeroing out) all regressors except the market closure-dummy from these intervals. This implies that the dummy "explains" the low returns (near) perfectly, while the inference regarding all other features of the return series is unaffected.

A.2 Regressors for constrained calendar and announcement volatility response patterns

In order to accommodate the overall impact through a parsimonious representation that also allows for efficient inference, the reported estimates for the announcement and calendar effects are based on the imposition of an a priori structure on the volatility response pattern. In particular, assuming that the feature in question affects volatility from interval n_0 to $n_0 + n_1$, the impact over the event window, $\tau = 0, 1, \ldots, n_1$, may then be represented by a polynomial specification,

$$p(\tau) = c_0 + c_1 \cdot \tau + \ldots + c_P \cdot \tau^P.$$

Of course, for $P = n_1$ this would effectively imply the estimation of a dummy variable for each of the $\bar{N} \equiv n_1 + 1$ event intervals. However, the use of a lower order polynomial afford a great degree of flexibility along with a significant reduction in the dimensionality of the parameter space. Furthermore, sensible constraints on the response pattern, including smoothness, are readily imposed in terms of the polynomial representation. For example, the requirement that the impact reflects a gradual movement away from the standard pattern is imposed by enforcing $p(0) = 0$. This simply annihilates the constant, i.e., $c_0 = 0$. Another desired property may be that the effect slowly fades, which is obtained by imposing $p(\bar{N}) = 0$. Substituting $\tau = \bar{N}$ into $p(\tau)$, solving for c_P, and inserting the resulting expression for c_P back into $p(\tau)$, leads to a restricted polynomial with one less parameter,

$$p(\tau) = c_0 \cdot [\,1 - (\tau / \bar{N})^P\,] + c_1 \cdot [\,1 - (\tau / \bar{N})^{P-1}\,] \cdot \tau + \ldots + c_{P-1} \cdot [\,1 - (\tau / \bar{N})\,] \cdot \tau^{P-1}.$$

We can now classify a number of our calendar and all of our announcement regressors through the choice of polynomial order, P, the response horizon, or $\bar{N}$, and the endpoint constraints imposed on $p(\tau)$. The following specifications underlie the results reported in the paper:

Tokyo Market Opening:	$\bar{N} = 6$,	$P = 1$,	$p(\bar{N}+1) = 0$,
Late Summer Day Slowdown:	$\bar{N} = 60$,	$P = 2$,	$p(0) = p(\bar{N}+1) = 0$,
Early Monday Effect:	$\bar{N} = 17$,	$P = 2$,	$p(\bar{N}+1) = 0$,
Late Friday Effect:	$\bar{N} = 46$ (58),	$P = 2$,	$p(0) = 0$,
EMS-Band Widening:	$\bar{N} = 30$,	$P = 3$,	$p(\bar{N}+1) = 0$,
Employment Report:	$\bar{N} = 24$,	$P = 3$,	$p(\bar{N}+1) = 0$,
All Other Announcements:	$\bar{N} = 12$,	$P = 3$,	$p(\bar{N}+1) = 0$.

The above representations leave one free parameter for the Tokyo market opening and the Summer slowdown, and two free parameters for the weekend effects denoted "Monday early" and "Friday late". The "Friday late" coefficients are identical in Summer and Winter, but the effects lasts an additional hour during Summer due to Daylight Savings Time. Finally, there are three announcement effect parameters, but as explained in section IV.A, we further restrict this pattern by imposing the common structure,

$$p_k(\tau) = \lambda_k \cdot p_0(\tau)\,,$$

where $p_k(\tau)$ refers to the polynomial for event type k, and $p_0(\tau)$ denotes a fixed response pattern. Specifically, we calibrated the pattern by fitting all three parameters for a set of announcements of about equal significance, resulting in a benchmark pattern that resembles the one associated with Category I releases. Concretely, $(c_0, c_1, c_2) = (2.18868, -0.64101, 0.07663)$. This uniquely identifies $p_0(\tau)$, and $p_k(\tau)$ thus has only one free "loading" parameter, λ_k. Of course, this procedure only strictly applies for response horizons corresponding to $\bar{N} = 12$. In order to retain the benchmark pattern for larger $\bar{N}$, we let the τ-variable progress only by a $(12/\bar{N})$-fraction of a unit per five-minute interval, rather than a full interval. This "stretches" the event time scale so that it conforms to the desired horizon.

Finally, we apply the corresponding "time-deformation" procedure to the sinusoids in the U.S. Summer Time intraday pattern in order to compensate for the one hour leftward shift from 7:00 to 6:00 GMT. This elongation of the intraday pattern is implemented over 19:55-00:00 GMT.

Table Notes

Table I: Ex-Post Return Volatility Measures and GARCH(1,1) Correlations
The table displays the correlations between forecasts of the daily DM-$ return standard deviations (Panel A) or the daily return variances (Panel B) with alternative measures of the ex-post return variability. The daily return standard deviation and variance for each weekday of the one-year sample, October 1, 1992 - September 29, 1993, is obtained from a MA(1)-GARCH(1,1) model estimated using daily data on the DM-$ spot exchange rate over the longer sample period from March 14, 1979 through September 29, 1993. The measures of ex-post return variability for the first two entries in Panel A and the first entry in Panel B, are constructed from percentage returns based on interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT have been excluded, resulting in a total of 74,880 return observations. The last two entries in each of the panels are based on ex-post return variability measures constructed from daily continuously compounded DM-$ returns over the one-year sample. The returns denoted R_t are calculated from the spot exchange rate observed at 12:00 GMT, consistent with the definition used for the longer daily sample, while the preceding entries use the exchange rates observed at 21:00 GMT, which is consistent with the definition of the trading day used for the five-minute return sample.

Table II: Largest Absolute Five-Minute Returns
The table reports the largest absolute five-minute returns calculated from the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. The absolute returns are obtained from interpolated five-minute logarithmic average bid-ask quotes. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT have been excluded, resulting in a total of 74,880 return observations. The 288 intraday returns per 24-hour trading day are numbered, starting at the interval 20:55-21:00 GMT, and ending with the interval 20:50-20:55 GMT. For each interval, we have subjectively indicated whether any economic or political event appears to have contributed to the large absolute five-minute return.

Table III: Flexible Fourier Form Regressions
The table reports the parameter estimates, with robust standard errors in square brackets and regular OLS standard errors in parentheses, for the regression of logarithmic squared demeaned five-minute DM-$ returns on deterministic regressors capturing calendar and announcement effects. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The robust standard errors reflect a Newey and West (1987) type correction incorporating 289 lags. The regression equation takes the form

$$\hat{x}_{t,n} \equiv 2\cdot\log[|R_{t,n} - \bar{R}|] - \log\hat{\sigma}^2_{t,n} = \hat{c} + f(\theta;t,n) + \hat{u}_{t,n} ,$$

where $R_{t,n}$ denotes the five-minute returns for interval n on day t, $\bar{R}$ the sample mean of the five-minute returns, $\hat{\sigma}^2_{t,n}$ is an a priori estimate of the overall daily level of the five-minute return standard deviation, $\hat{u}_{t,n}$ is a mean zero error term, and $f(\theta;t,n)$ represents the deterministic calendar and announcement regressors. The volatility estimates, $\hat{\sigma}_{t,n}$, for interval n on day t, are obtained from an estimated MA(1)-GARCH(1,1) model fit to a longer daily sample of DM-$ spot exchange rates from March 14, 1979 through September 29, 1993. Denoting the daily return standard deviation estimate by $\hat{\sigma}_t$, the daily volatility factor is captured by $\hat{\sigma}_{t,n} \equiv N^{-1/2}\cdot\hat{\sigma}_t$. The "Daily Volatility Excluded" column indicates that $N^{-1/2}\cdot\bar{\sigma}$ is used in place of $\hat{\sigma}_{t,n}$, where $\bar{\sigma}$ denotes the sample mean of $\hat{\sigma}_t$. The $f(\theta;t,n)$ function is given by

$$f(\theta;t,n) = \mu_0 + \sum_{k=1}^{D} \lambda_k \cdot I_k(t,n) + \sum_{p=1}^{4} (\delta_{c,p}\cdot\cos\frac{p2\pi}{N}n + \delta_{s,p}\cdot\sin\frac{p2\pi}{N}n) ,$$

During the U.S. Summer Time, the sinusoids are translated leftward by one hour and an additional restricted second order polynomial allows for a volatility slowdown between 19:00 and 24:00 GMT. The $I_k(t,n)$ regressors indicate either regular dummy variables (in the case of Holidays or weekdays) or a prespecified volatility response pattern associated with a calendar related characteristic or an announcement. A separate linear volatility decay is allowed for the Tokyo open, 00:00-00:35 GMT. Similarly, a restricted second order polynomial adapts to the volatility slowdown around the weekends, i.e., early Monday morning, 21:00-22:30 GMT, and late Friday, 17:00-21:00 GMT (U.S. Winter Time) or 16:00-21:00 GMT (U.S. Summer Time). Finally, the volatility decay pattern following announcements are restricted to last one hour (13 intervals), except for the Employment Report pattern which lasts two hours (25 intervals). All of the response patterns are approximated by a third order polynomial restricted to reach zero at the end of the response horizon. The announcement coefficients measure the extent to which the absolute returns load onto this pattern following the announcement. Category I comprises U.S. announcements on GDP, the trade balance, and durable goods, while Category II covers U.S. releases of PPI, retail sales, housing starts, leading indicators, jobless claims, and factory orders, and the German M3 figures.

Table IV: Estimated Announcement Effects

The table gives the parameter estimates with robust standard errors in square brackets associated with the specific announcements obtained from regressions of the logarithmic squared demeaned five-minute DM-$ returns on a set of deterministic regressors allowing for calendar and other announcement effects. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The regression takes the form

$$\hat{x}_{t,n} \equiv 2\cdot\log[|R_{t,n} - \bar{R}|] - \log\hat{\sigma}^2_{t,n} = \hat{c} + f(\theta;t,n) + \hat{u}_{t,n} ,$$

where $R_{t,n}$ denotes the five-minute returns for interval n on day t, $\bar{R}$ the sample mean of the five-minute returns, $\hat{\sigma}^2_{t,n}$ is an a priori estimate of the overall daily level of the five-minute return standard deviation, $\hat{u}_{t,n}$ is a mean zero error term, and $f(\theta;t,n)$ represents the deterministic calendar and announcement regressors. The volatility estimates, $\hat{\sigma}_{t,n}$, for interval n on day t, are obtained from an estimated MA(1)-GARCH(1,1) model fit to a longer daily sample of DM-$ spot exchange rates covering the period from March 14, 1979 through September 29, 1993. Denoting the daily return standard deviation estimates by $\hat{\sigma}_t$, the daily volatility factor is captured by $\hat{\sigma}_{t,n} \equiv N^{-1/2}\cdot\hat{\sigma}_t$. The $f(\theta;t,n)$ function is given as

$$f(\theta;t,n) = \mu_0 + \sum_{k=1}^{D} \lambda_k \cdot I_k(t,n) + \sum_{p=1}^{4} (\delta_{c,p}\cdot\cos\frac{p2\pi}{N}n + \delta_{s,p}\cdot\sin\frac{p2\pi}{N}n) ,$$

During U.S. Summer Time, the sinusoids are translated leftward by one hour and an additional restricted second order polynomial allows for a volatility slowdown between 19:00 and 24:00 GMT. The $I_k(t,n)$ regressors indicate either regular dummy variables (for Holidays or weekdays) or prespecified volatility response patterns associated with a calendar feature or an announcement. A separate linear volatility decay is allowed for the Tokyo open, 00:00-00:35 GMT. Similarly, a restricted second order polynomial adapts to the volatility slowdown around weekends, early Monday morning, 21:00-22:30 GMT, and late Friday, 17:00-21:00 GMT (U.S. Winter Time) or 16:00-21:00 GMT (U.S. Summer Time). The volatility decay pattern following announcements are restricted to lasts one hour (13 intervals), except for the U.S. Employment Report pattern which lasts for two hours (25 intervals). All of the response patterns are approximated by a third order polynomials restricted to reach zero at the end of the response horizon. The reported coefficients measure the extent to which the absolute returns load onto this pattern following the announcement. Beyond the specific announcement under investigation, all of the regressions allow for the independent influence of the U.S. Employment Report, the Bundesbank Meeting, and Category I and II announcements. Category I includes U.S. announcements on GDP, the trade balance, and durable goods. Category II covers U.S. releases of PPI, retail sales, housing starts, leading indicators, jobless claims, and factory orders, and German M3 figures. If the individual announcement under investigation is a member of one of these categories, the announcement is dropped from the category. The instantaneous jump in volatility measures the estimated increase in the five-minute absolute return for the interval where the announcement is made, while the estimated total cumulative absolute return induced by the announcement over the assumed response horizon is measured relative to the median daily cumulative absolute return of 9.0%.

Table V: Explained Variation for Alternative Absolute Return Forecasts

The table reports the R^2's from OLS regressions of realized daily cumulative absolute DM-$ returns, or realized five-minute absolute DM-$ returns, on alternative one-day-ahead absolute return forecasts. For the daily cumulative absolute returns, the regressions takes the form

$$\sum_{n=1}^{N} |R_{t,n} - \bar{R}| \;=\; b_0 \;+\; b_1 \sum_{n=1}^{N} v(I;t,n) \;+\; \epsilon_t ,$$

where v(I;t,n) denotes the relevant absolute return forecast for interval n on day t, and ϵ_t is an error term. The regressions for the five-minute absolute returns are calculated as

$$|R_{t,n} - \bar{R}| \;=\; b_0 \;+\; b_1 \cdot v(I;t,n) \;+\; \epsilon_t .$$

The five-minute returns are based on interpolated logarithmic average bid-ask quotes for the DM-$ spot exchange rate from October 1, 1992 through September 29, 1993. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 five-minute observations. The daily cumulative absolute returns are aggregated from 21:00 GMT to 21:00 GMT the following day. The corresponding one-day-ahead five-minute absolute return forecasts are obtained as

$$v(I;t,n) \;=\; c_0\, \hat{\sigma}_{t,n}\, \exp\!\left(\frac{\hat{f}_c(t,n)\, I_c \;+\; \hat{f}_a(t,n)\, I_a \;+\; \hat{f}_h(t,n)}{2} \right)$$

where $\hat{\sigma}_{t,n} = N^{-1/2} \cdot [\hat{\sigma}_t \cdot I_\sigma + \bar{\sigma} \cdot (1 - I_\sigma)]$ represents an estimate of the benchmark return volatility of the interval, while $\hat{f}_c(t,n)$, $\hat{f}_a(t,n)$ and $\hat{f}_h(t,n)$ denote the estimated calendar, announcement and Holiday effects from a regression of normalized, log-squared, demeaned five-minute DM-dollar returns on calendar, announcement and Holiday regressors. The functional form of the forecast equation translates the estimates into the absolute return dimension. The indicator variables I_σ, I_c, and I_a signify whether the features associated with a given effect are included in the construction of the forecast. For example, the indicator vector $I \equiv (I_\sigma, I_c, I_a) = (0,1,1)$ corresponds to the model whose coefficients are listed in the last column of Table III; that is, the daily volatility factor is assumed to be constant. Estimates for the time-varying daily return standard deviations over the one-year sample, $\hat{\sigma}_t$, are obtained from a MA(1)-GARCH(1,1) model fit to a longer daily sample of DM-$ returns covering the period from March 14, 1979 through September 29, 1993. The sample mean of $\hat{\sigma}_t$ is denoted by $\bar{\sigma}$. The "*" indicates the allowance for day-of-the-week dummies among the calendar effects, while these are excluded otherwise.

Table I

Ex-Post Return Volatility Measures and GARCH(1,1) Correlations

Panel A: Gaussian MA(1)-GARCH(1,1) Estimates of σ_t

$\Sigma_{n=1}^{N} \lvert R_{t,n} \rvert$	$(\Sigma_{n=1}^{N} R_{t,n}^2)^{1/2}$	$\lvert \Sigma_{n=1}^{N} R_{t,n} \rvert$	$\lvert R_t \rvert$
0.672	0.618	0.046	0.086

Panel B: Gaussian MA(1)-GARCH(1,1) Estimates of σ_t^2

$\Sigma_{n=1}^{N} R_{t,n}^2$	$(\Sigma_{n=1}^{N} R_{t,n})^2$	R_t^2
0.660	0.066	0.107

Table II

Largest Absolute Five-Minute Returns

Absolute Return	Date	Interval	Weekday	Event
1.244	10/02	188	Friday	Employment Report
0.897	06/04	188	Friday	Employment Report
0.648	11/19	200	Thursday	Jobless Claims Housing Starts
0.637	03/04	189	Thursday	Bundesbank Meeting
0.581	09/03	188	Friday	Employment report
0.580	06/11	188	Friday	Retail Sales Producer Price Index
0.573	10/02	189	Friday	Employment Report
0.530	09/21	234	Tuesday	Russia Crisis
0.529	11/20	37	Friday	ERM Turmoil
0.527	01/29	200	Friday	Durable Goods
0.517	08/02	36	Monday	ERM Band Revision
0.510	10/05	243	Monday	U.S. Stock Market Plunge
0.503	09/21	233	Tuesday	Russia Crisis
0.501	03/05	200	Friday	Employment Report
0.498	09/16	197	Thursday	Industrial Output
0.494	08/31	188	Tuesday	Gross Domestic Product
0.480	07/02	188	Friday	Employment Report
0.478	10/05	240	Monday	U.S. Stock Market Plunge
0.463	10/05	229	Monday	U.S. Stock Market Plunge
0.458	11/04	39	Wednesday	U.S. Presidential Election
0.455	09/21	235	Tuesday	Russia Crisis
0.449	08/19	188	Thursday	Jobless Claims Trade Balance
0.441	10/23	195	Friday	ERM Turmoil
0.439	04/22	198	Thursday	Bundesbank Meeting
0.434	10/27	200	Tuesday	Gross Domestic Product

Table III

Flexible Fourier Form Regressions

Parameter	Full System	Day-of-Week Effect Excluded	Day-of-Week, Daily Volatility Excluded
$\mu_0 + \hat{c}$	-1.77	-1.76	-1.85
	[-32.8]	[-69.2]	[-56.3]
	(-79.4)	(-155.6)	(-162.0)
$\delta_{c,1}$	-0.12	-0.13	-0.13
	[-4.41]	[-4.58]	[-4.78]
	(-8.27)	(-8.62)	(-8.77)
$\delta_{c,2}$	-0.13	-0.13	-0.13
	[-4.93]	[-5.09]	[-5.10]
	(-8.16)	(-8.30)	(-8.26)
$\delta_{c,3}$	-0.28	-0.28	-0.29
	[-11.8]	[-12.0]	[-11.4]
	(-18.4)	(-18.5)	(-18.5)
$\delta_{c,4}$	0.14	0.14	0.14
	[8.10]	[8.01]	[8.10]
	(10.6)	(10.6)	(10.6)
$\delta_{s,1}$	-0.62	-0.62	-0.62
	[-24.1]	[-23.8]	[-23.4]
	(-38.6)	(-38.5)	(-38.4)
$\delta_{s,2}$	-0.21	-0.21	-0.21
	[-10.4]	[-10.2]	[-10.2]
	(-14.3)	(-14.1)	(-14.0)
$\delta_{s,3}$	0.17	0.18	0.17
	[8.64]	[8.76]	[8.55]
	(11.9)	(12.1)	(11.8)
$\delta_{s,4}$	-0.01	-0.01	-0.01
	[-0.68]	[-0.46]	[-0.68]
	(-0.91)	(-0.62)	(-0.94)
Summer Slowdown	-1.14	-1.15	-1.08
	[-5.91]	[-5.95]	[-5.06]
	(-10.6)	(-10.7)	(-9.93)
Tokyo Opening	0.59	0.58	0.59
	[8.96]	[8.91]	[9.06]
	(9.81)	(9.76)	(9.76)

Table III, Continued

Parameter	Full System	Day-of-Week Effect Excluded	Day-of-Week, Daily Volatility Excluded
Holiday	-0.698 [-5.76] (-13.86)	-0.712 [-6.28] (-15.25)	-0.703 [-6.87] (-14.93)
Employment Report	1.755 [10.38] (11.11)	1.746 [10.47] (11.09)	1.739 [8.82] (10.95)
Category I Announcement	0.997 [7.23] (8.35)	0.991 [7.28] (8.33)	0.992 [7.48] (8.26)
Category II Announcement	0.627 [6.71] (8.64)	0.620 [7.03] (8.65)	0.619 [6.89] (8.56)
Bundesbank Meeting	1.465 [6.20] (10.01)	1.457 [6.20] (9.99)	1.492 [6.43] (10.13)
Monday Early	-0.301 [-0.26] (-0.36)	-0.368 [-0.32] (-0.44)	-0.529 [-0.45] (-0.63)
	0.001 [0.001] (0.001)	0.069 [0.047] (0.065)	0.208 [0.14] (0.19)
Friday Late	-0.609 [-2.04] (-3.36)	-0.412 [-1.43] (-2.37)	-0.437 [-1.39] (-2.49)
	0.068 [0.49] (0.88)	0.011 [0.08] (0.14)	0.017 [0.12] (0.22)
Tuesday	-0.065 [-0.95] (-2.17)	-	-
Wednesday	0.006 [0.08] (0.19)	-	-
Thursday	0.050 [0.74] (1.65)	-	-
Friday	0.096 [1.31] (2.99)	-	-

Table IV.A

Important Announcement Effects

Announcement	Coefficient [Robust t-Stat]	Instantaneous Jump in Volatility (Percent)	Impact in Percent of Daily Cum. Abs. Return
Employment Report	1.75 [11.5]	576	15.1
Advance Report on Durable Goods	1.27 [5.75]	303	5.17
Bundesbank Meeting	1.46 [9.74]	392	5.11
Merchandise Trade	0.889 [4.24]	164	3.12
Gross Domestic Product	0.836 [3.43]	150	2.87
Producer Price Index	0.703 [3.67]	116	2.31
Retail Sales	0.670 [2.86]	108	2.17
German M3	0.872 [4.77]	160	2.13
Leading Indicators	0.624 [3.55]	98	1.99
Housing Starts	0.515 [2.29]	76	1.59
Factory Orders	0.481 [2.05]	69	1.46
New Jobless Claims	0.334 [3.02]	44	0.968
Japanese Gross National Product	0.600 [2.40]	93	0.949
German Gross Domestic Product	0.506 [1.43]	74	0.931

Table IV.B

Less Important U.S. Announcements

Announcement	Coefficient	Robust t-Stat
U.S. Treasury Report	0.338	1.62
Consumer Confidence (Conference Board)	0.273	1.20
Consumer Price Index	0.236	1.02
Construction Spending	0.211	0.954
Car Sales	0.091	0.709
Business Inventories	0.124	0.704
Housing Completions	0.070	0.430
Import Prices	0.076	0.374
University of Michigan Survey	0.043	0.315
Current Account Deficit	0.084	0.314
Industrial Output / Capital Utilization	0.067	0.282
Non-Farm Productivity	0.035	0.154
M2 Figures	0.017	0.134
Personal Income	0.030	0.102
Real Earnings	0.005	0.021
Reserve Assets	-0.012	-0.062
House Sales	-0.041	-0.150
Minutes from FOMC Meeting	-0.177	-0.613
Capital Spending Survey	-0.261	-0.689
NAPM Survey	-0.205	-0.777
Consumer Installment Credit	-0.375	-1.02
Wholesale Sales	-0.181	-1.06

Table IV.C

Less Important German Announcements

Announcement	Coefficient	Robust t-Stat
Wholesale Turnover	0.322	1.56
Retail Sales	0.124	0.668
Consumer Price Index (All States Tallied)	0.072	0.668
East German Consumer Price Index	0.122	0.647
East German Industrial Orders	0.152	0.630
Industrial Orders	0.110	0.465
Producer Price Index	0.085	0.423
Wholesale Prices	0.054	0.295
Current Account	0.032	0.181
Consumer Price Index (First State)	0.010	0.051
Business Insolvencies	0.002	0.010
Employment Report	-0.004	-0.022
Import Prices	-0.011	-0.049
Consumer Price Index (Final)	-0.089	-0.427
East German Employment	-0.075	-0.438
East German Producer Price Index	-0.092	-0.503
Industrial Output	-0.176	-0.702
Capital Account	-0.238	-1.14
East German Industrial Output	-0.245	-1.18
Consumer Price Index (Preliminary)	-0.563	-1.84

Table V

Explained Variation for Alternative Absolute Return Forecasts

Design	Daily Cumulative Absolute Returns	Five-Minute Absolute Returns
Complete Model $(I_\sigma, I_c, I_a) = (1,1,1)$	0.606	0.159
No Announcements $(I_\sigma, I_c, I_a) = (1,1,0)$	0.579	0.113
No Calendar Effects $(I_\sigma, I_c, I_a) = (1,0,1)$	0.603	0.084
Only Daily Volatility $(I_\sigma, I_c, I_a) = (1,0,0)$	0.578	0.034
No Daily Volatility $(I_\sigma, I_c, I_a) = (0,1,1)$	0.119	0.124
Only Announcements $(I_\sigma, I_c, I_a) = (0,0,1)$	0.114	0.049
Only Calendar Effects $(I_\sigma, I_c, I_a) = (0,1,0)$	0.083	0.081
Calendar + Day-of-Week $(I_\sigma, I_c, I_a) = (0,1,0)^*$	0.107	0.083
Only Holiday Effects $(I_\sigma, I_c, I_a) = (0,0,0)$	0.080	0.002

Figure Notes

Figure 1: Daily GARCH(1,1) Volatility Forecasts
The figure plots the one-step-ahead conditional standard deviation forecasts from a MA(1)-GARCH(1,1) model for the daily DM-$ spot exchange rate from October 1, 1992 through September 30, 1993, for a total of 260 non-weekend days. The model is estimated with data over the longer sample period from March 14, 1979 through September 29, 1993.

Figure 2: Daily GARCH(1,1) Volatility Forecasts versus Ex-Post Return Variability Measures
The figures plot the conditional return standard deviation forecasts from a MA(1)-GARCH(1,1) model for the daily DM-$ returns from October 1, 1992 trough September 29, 1993, along with the corresponding realized absolute daily returns in Panel A and the corresponding cumulative absolute five-minute returns in Panel B. All series have been normalized to average unity over the one-year sample.

Figure 3: Intraday Volatility Pattern
The figure plots the average absolute five-minute DM-$ return for each five-minute interval, starting with the interval 20:55-21:00 GMT and ending at 20:50-20:55 GMT. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. All 260 weekdays are employed in calculating the averages.

Figure 4: Daily and Weekly Volatility Patterns
The figure displays the estimated average absolute five-minute DM-$ returns obtained from a regression on two-hour and day-of-week dummies. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. The two-hour intervals start out at 20:55-22:55 GMT and end at 18:55-20:55 GMT.

Figure 5: Intradaily U.S. Summer and Winter Time Volatility Patterns
The figure plots the average absolute five-minute DM-$ return for each five-minute interval, starting with 20:55-21:00 GMT and ending at 20:50-20:55 GMT. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The Tokyo lunch period, 3:00-4:45 GMT, is assigned an artificially low returns. All 145 weekdays during the U.S. Summer Time and the 115 weekdays during U.S. Winter Time are employed.

Figure 6: U.S. Announcement Day Volatility
The figure plots the average absolute five-minute DM-Dollar return for each five-minute interval, starting with 20:55-21:00 GMT and ending at 20:50-20:55 GMT for days with regularly scheduled U.S. macroeconomic announcements. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The Tokyo lunch period, 3:00-4:45 GMT, is assigned an artificially low return. Figure A is based on U.S. Summer Time, while Figure B graphs the average across the U.S. Winter Time announcement days. These days each contain at least one release, at 8:30 Eastern Standard Time, of one of the following U.S. macroeconomic announcements: the employment report, the merchandise trade deficit, the producer price index, the advance durable goods report, estimates or revisions to the gross domestic product, retail sales, housing starts, leading indicators, and new jobless claims.

Figure 7: Flexible Fourier Form Fit
The figure graphs the fit to the average logarithmic squared, normalized and demeaned five-minute DM-$ returns across the 24-hour weekday trading cycle. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The intervals starts with 20:55-21:00 GMT and ends at 20:50-20:55 GMT. The Tokyo lunch period, 3:00-4:45 GMT, is artificially assigned low returns, so this part of the pattern is not estimated. The fit is based on four sets of sinusoids, dummies for the Tokyo open period, 00:00-00:35 GMT, and constrained

second order polynomials for early Monday and late Friday, as well as the latter part of the U.S. Summer Time trading day. Separate estimates for U.S. Summer and Winter Time are reported in figures A and B.

Figure 8: Average Intradaily Log-Volatility Fit
The figure graphs the fit to the average logarithmic squared, normalized and demeaned five-minute DM-dollar returns across the 24-hour weekday trading cycle plotted against the corresponding averaged sample values. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The GMT axis starts with the 20:55-21:00 GMT interval and ends at 20:50-20:55 GMT. The Tokyo lunch period, 3:00-4:45 GMT, is artificially assigned low returns, so this part is not fitted. The fit is based on four sets of sinusoids, dummies for the Tokyo open, 00:00-00:35 GMT, and constrained second order polynomials for the latter part of the U.S. Summer Time trading day, as well as early Monday and late Friday. The latter "weekend effects" are not indicated on the figures. The Summer and Winter Time averages in figures A and B are based on 145 and 115 weekdays, respectively.

Figure 9: Average Intradaily Absolute Return Fit
The figure graphs the fit to the average absolute five-minute DM-dollar returns across the 24-hour weekday trading cycle plotted against the corresponding averaged sample values. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The GMT axis starts with the 20:55-21:00 GMT interval and ends at 20:50-20:55 GMT. The Tokyo lunch period, 3:00-4:45 GMT, is artificially assigned low returns, so this part is not fitted. The fit is based on a Flexible Fourier Form regression of logarithmic squared, normalized and demeaned returns onto four sets of sinusoids, dummies for the Tokyo open, 00:00-00:35 GMT, and constrained second order polynomials for the latter part of the U.S. Summer Time trading day, as well as early Monday and late Friday. The latter "weekend effects" are not indicated on the figures. The Summer and Winter Time averages in figures A and B are based on 145 and 115 weekdays, respectively.

Figure 10: Dynamic Announcement Response Patterns
The figure graphs the relative strength and duration of the estimated dynamic log-volatility response pattern of the five-minute DM-$ returns following the release of macroeconomic announcements. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT are excluded, resulting in a total of 74,880 return observations. The response pattern estimates are obtained from a Flexible Fourier Form regression of the logarithmic squared, normalized and demeaned returns onto four sets of sinusoids, dummies for the Tokyo market opening, and constrained second order polynomials for the latter part of the U.S. Summer Time trading day, as well as early Monday and late Friday.

Figure 11: Absolute Return Correlograms
The figures display the autocorrelations for demeaned raw and filtered five-minute absolute returns. The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the DM-$ spot exchange rate over the October 1, 1992 through September 29, 1993 sample period. Quotes from Friday 21:00 GMT through Sunday 21:00 GMT, along with the Tokyo lunch period, 3:00-4:45 GMT, are not included, resulting in 267 weekday return observations, for a total of 69,420 five-minute returns. Additional minor corrections were also made for extremely low quoting activity during Holidays and gaps in the data series. The filtered returns in figure B are obtained by standardizing the raw demeaned absolute returns by the estimated volatility impact of calendar, Holiday and announcement effects.

Figure 1: Daily GARCH(1,1) Volatility Forecasts

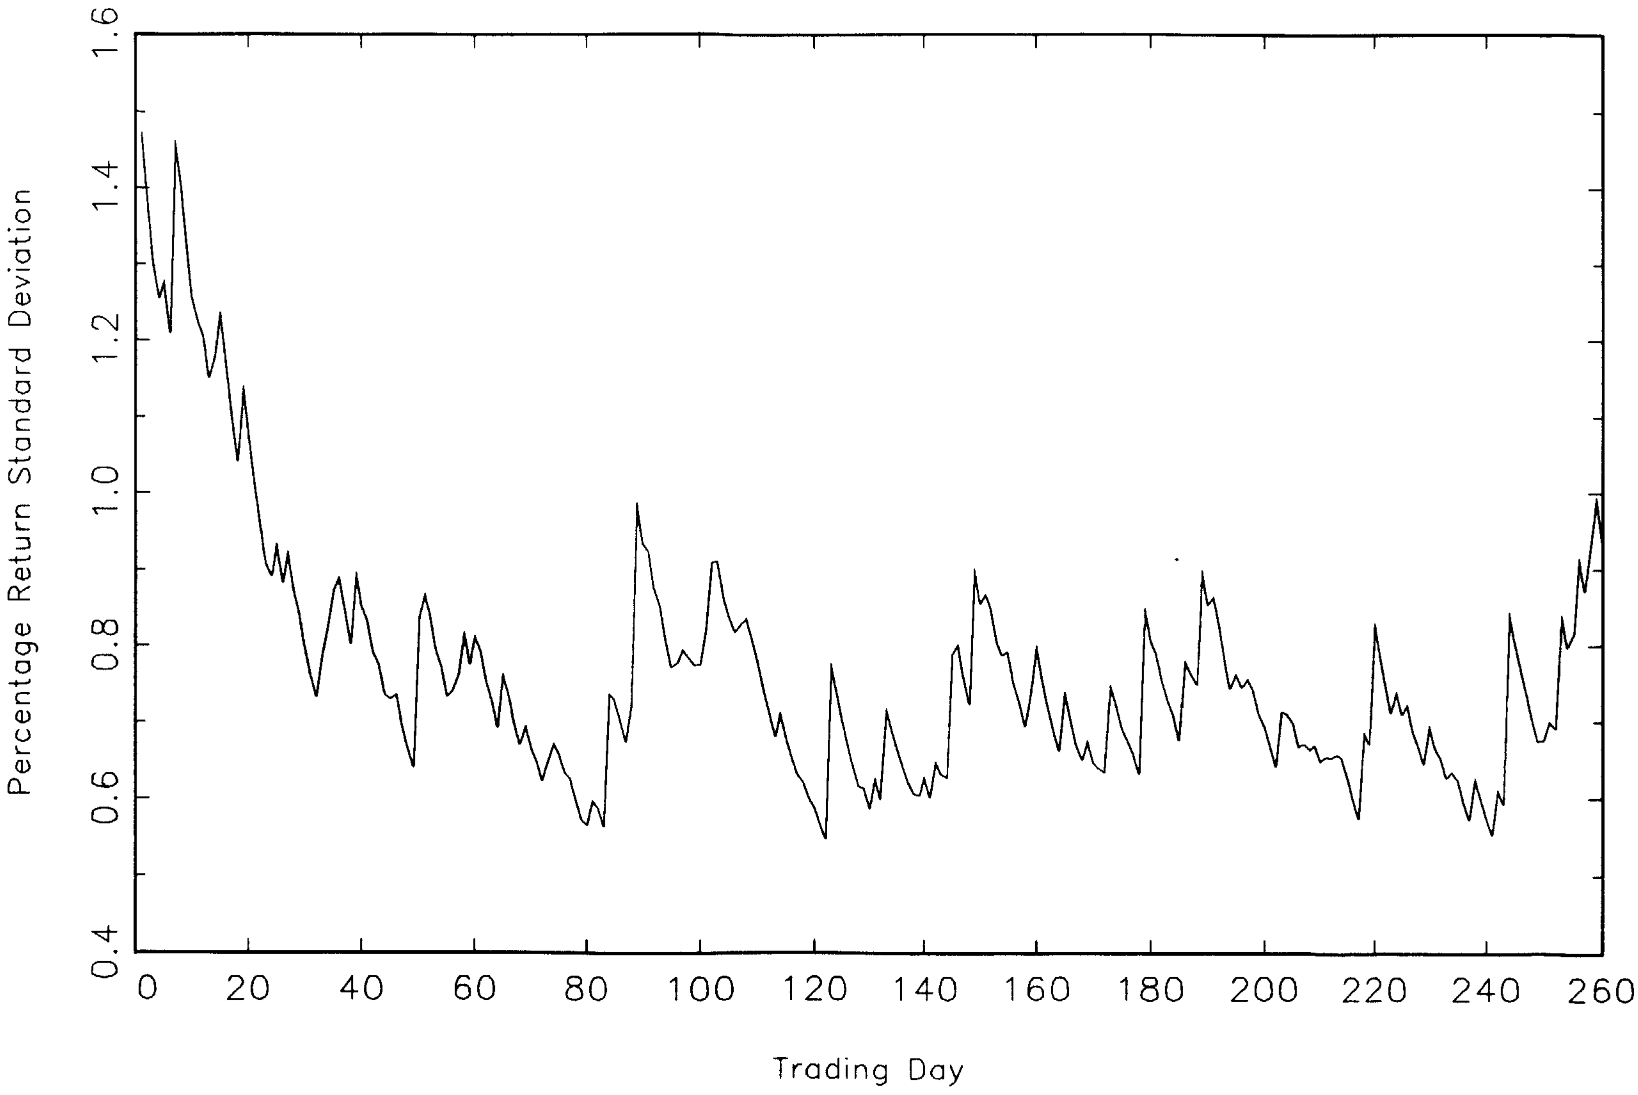

Figure 2A: Daily GARCH(1,1) Volatility Forecasts versus Ex-Post Return Variability Measures

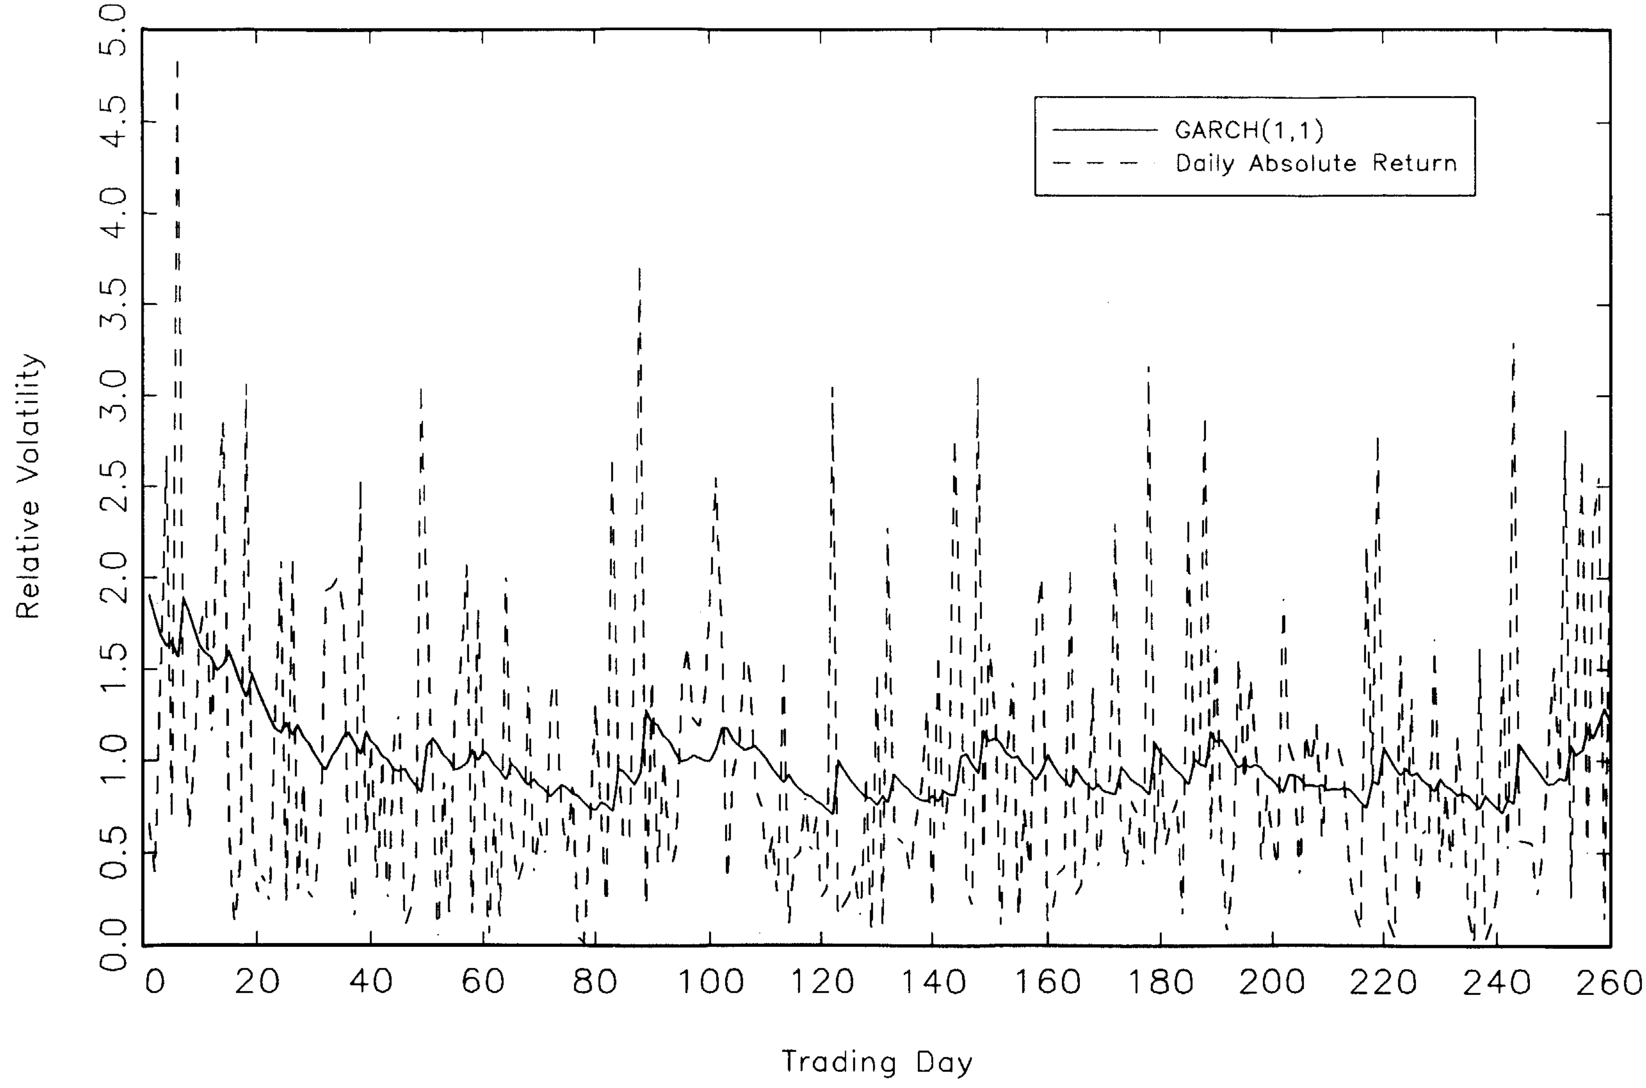

Figure 2B: Daily GARCH(1,1) Volatility Forecasts versus Ex-Post Return Variability Measures

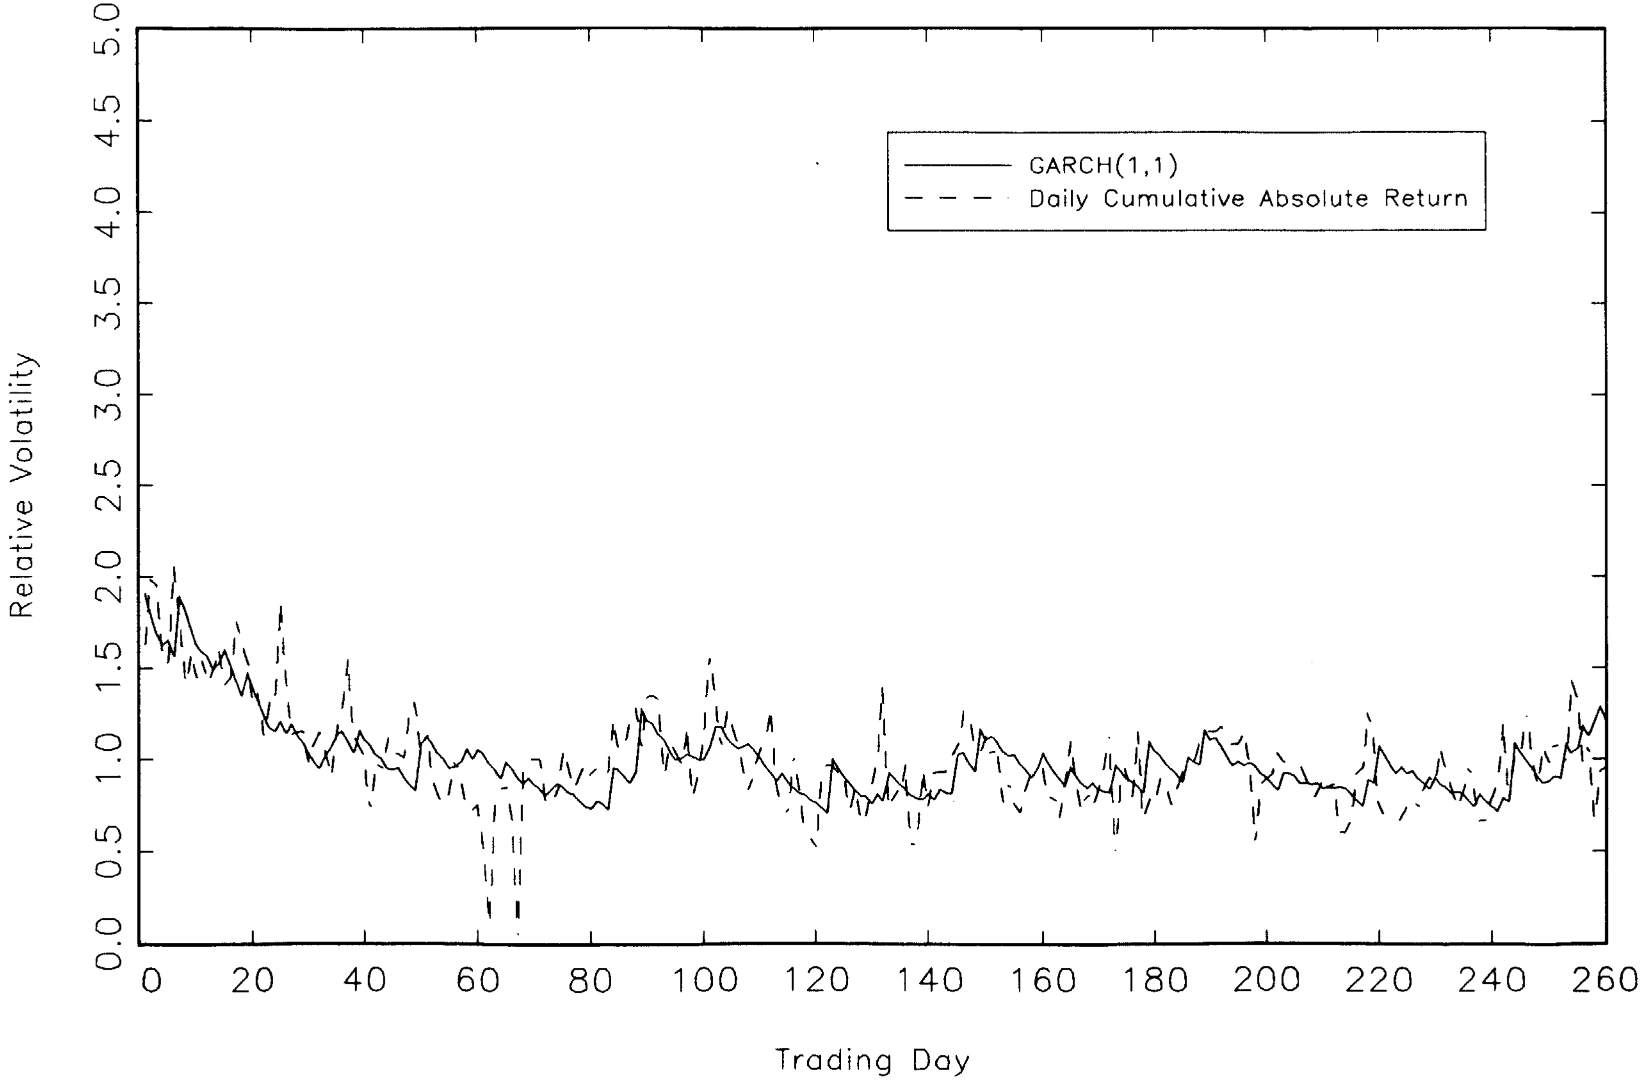

Figure 3: Intraday Volatility Pattern

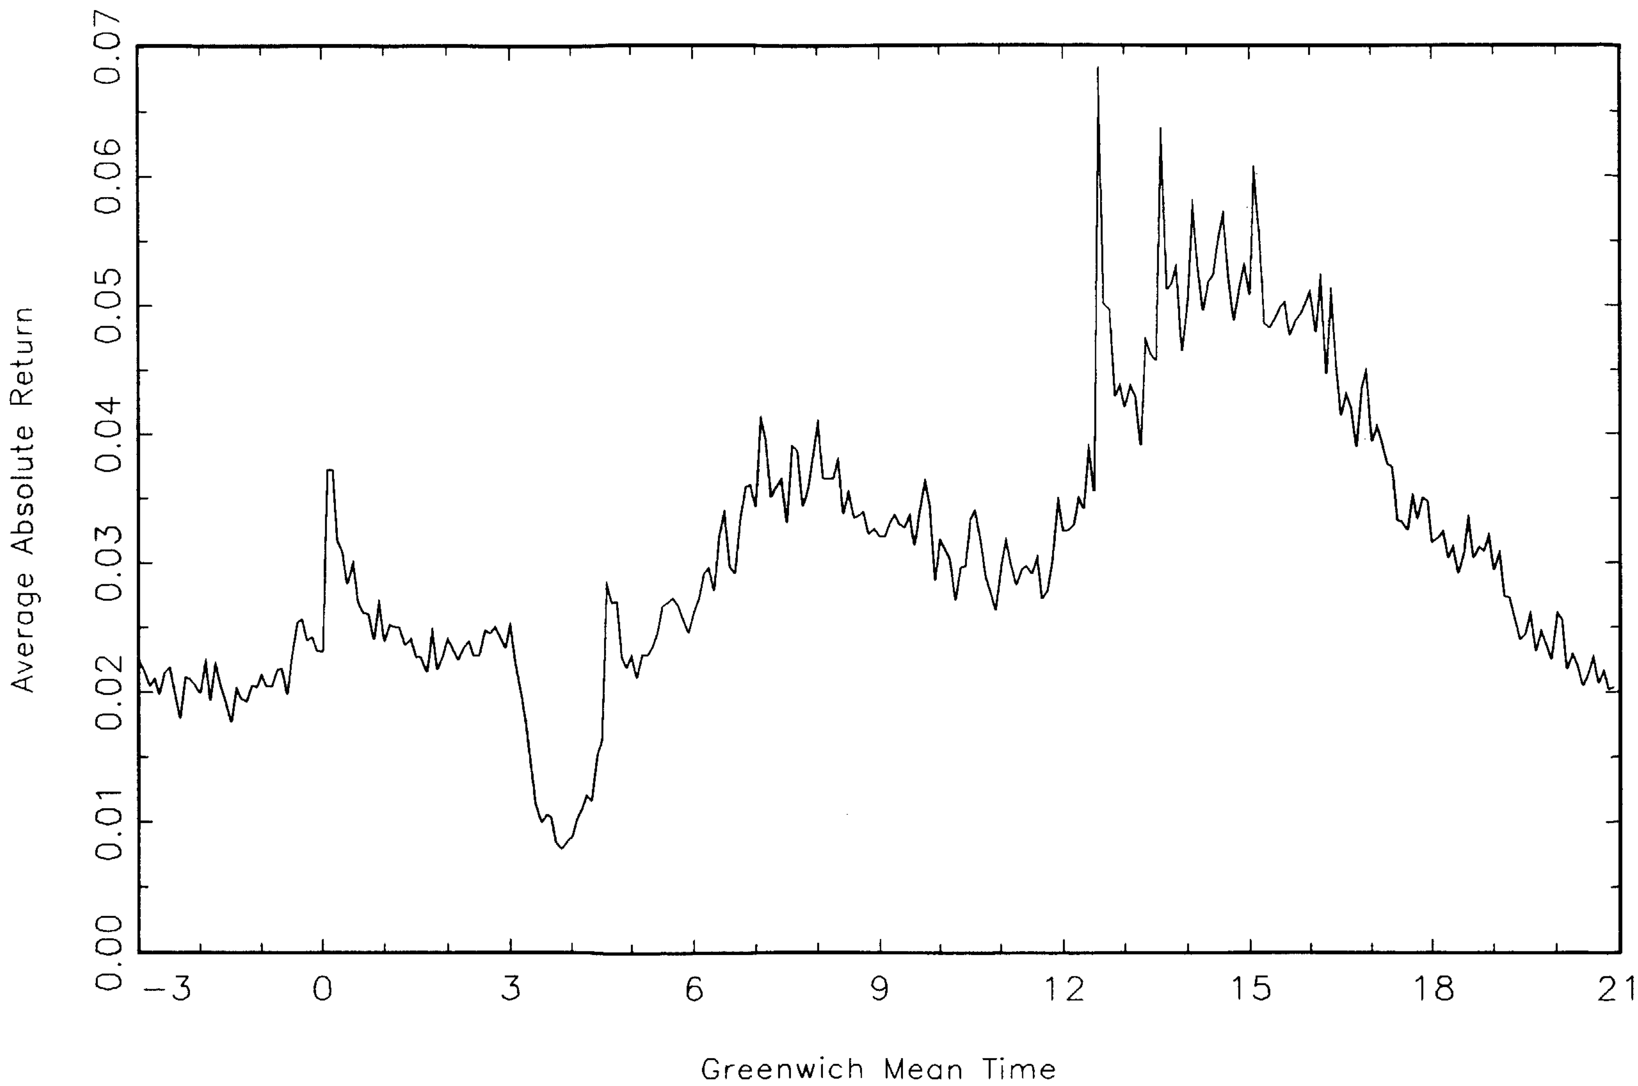

Figure 4: Daily and Weekly Volatility Pattern

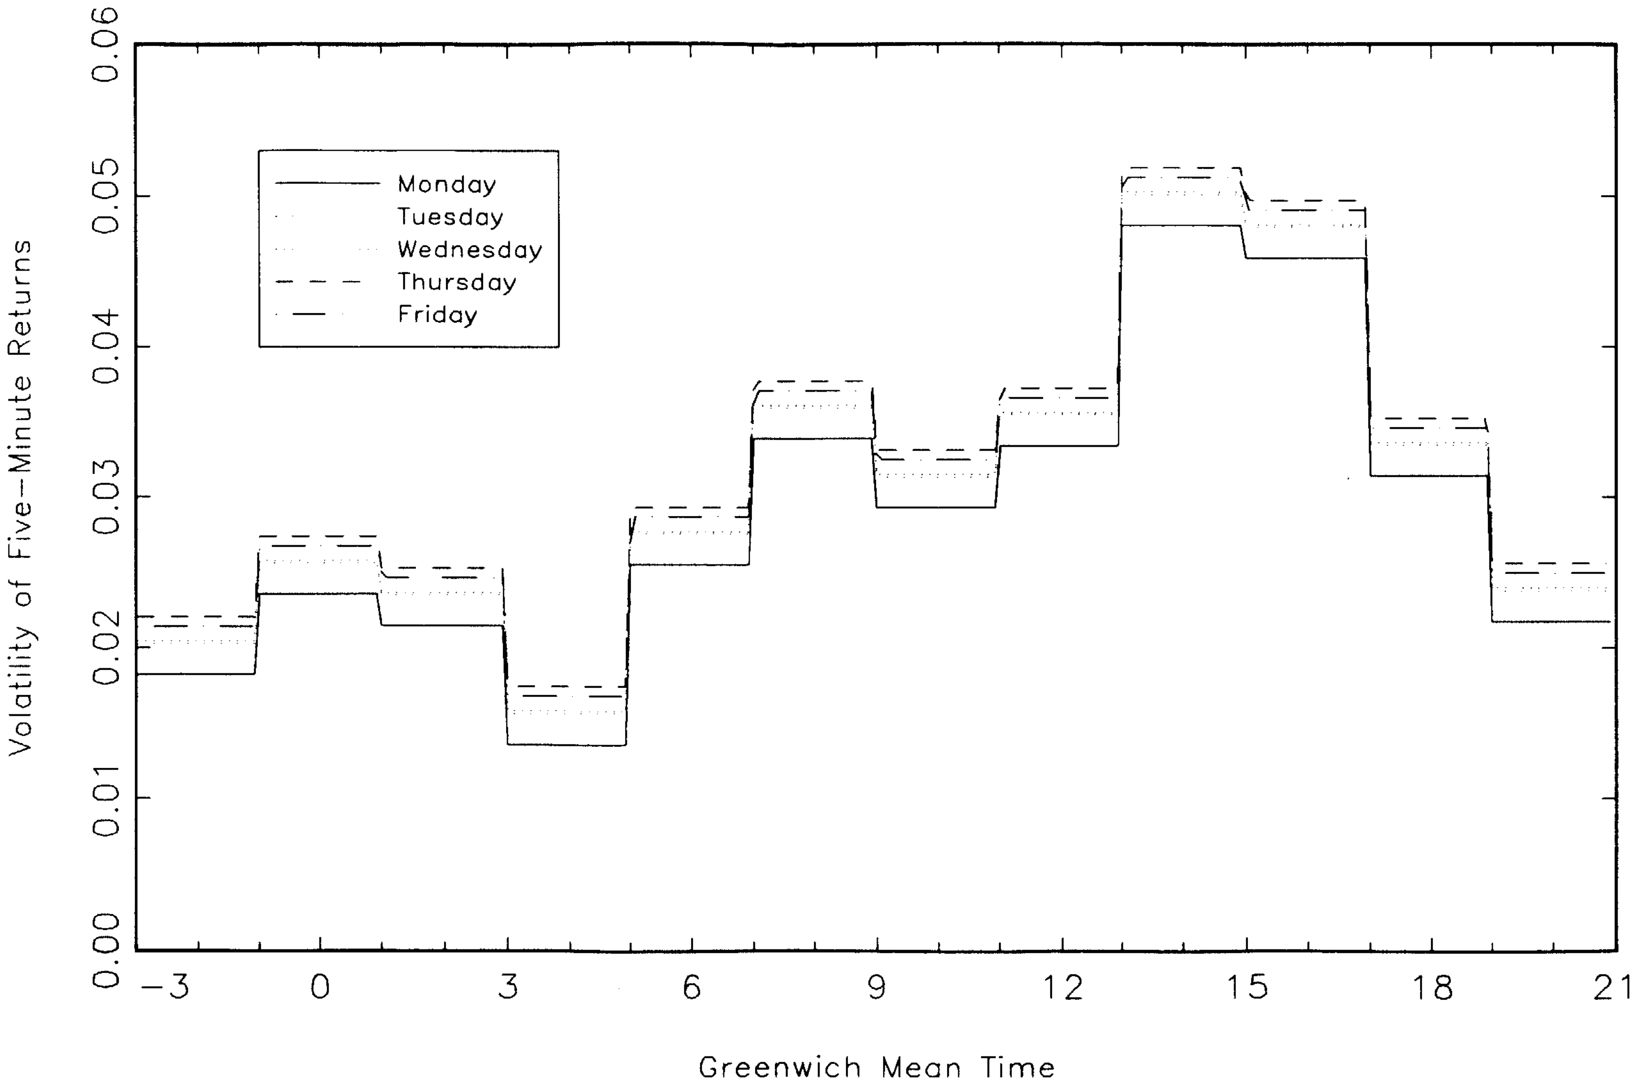

Figure 5: Intradaily U.S. Summer and Winter Time
Volatility Patterns

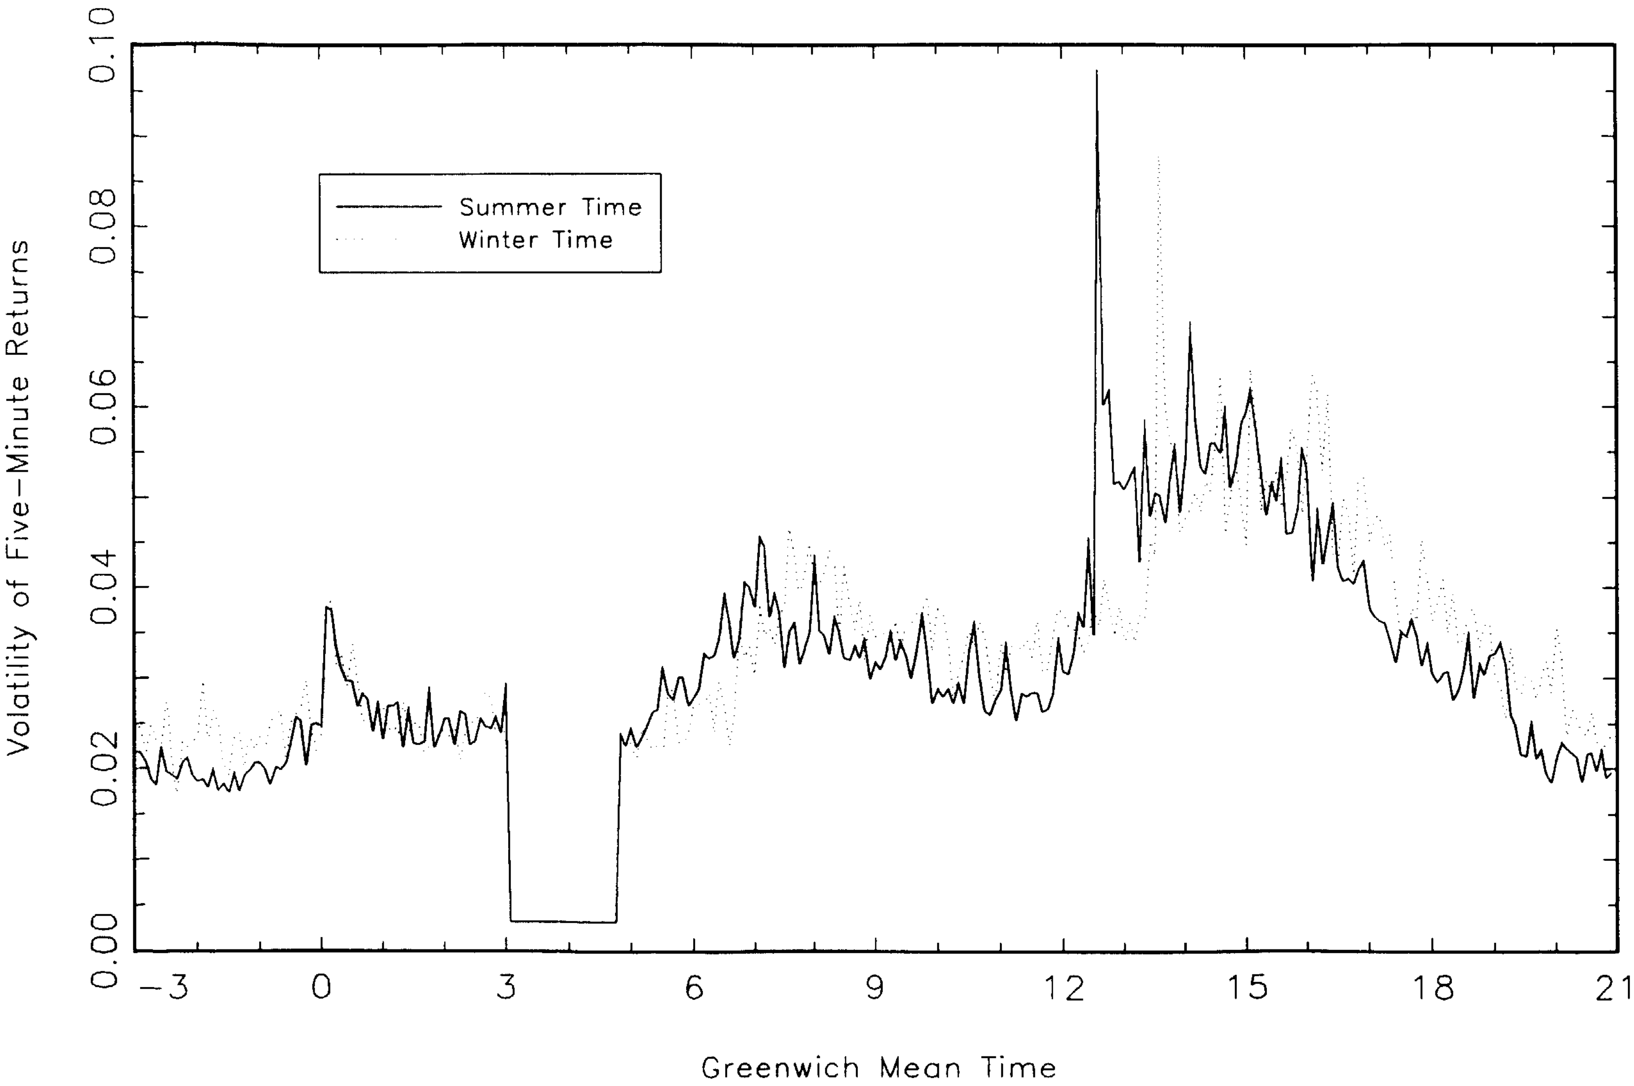

Figure 6A: U.S. Announcement Day Volatility
Summer Time

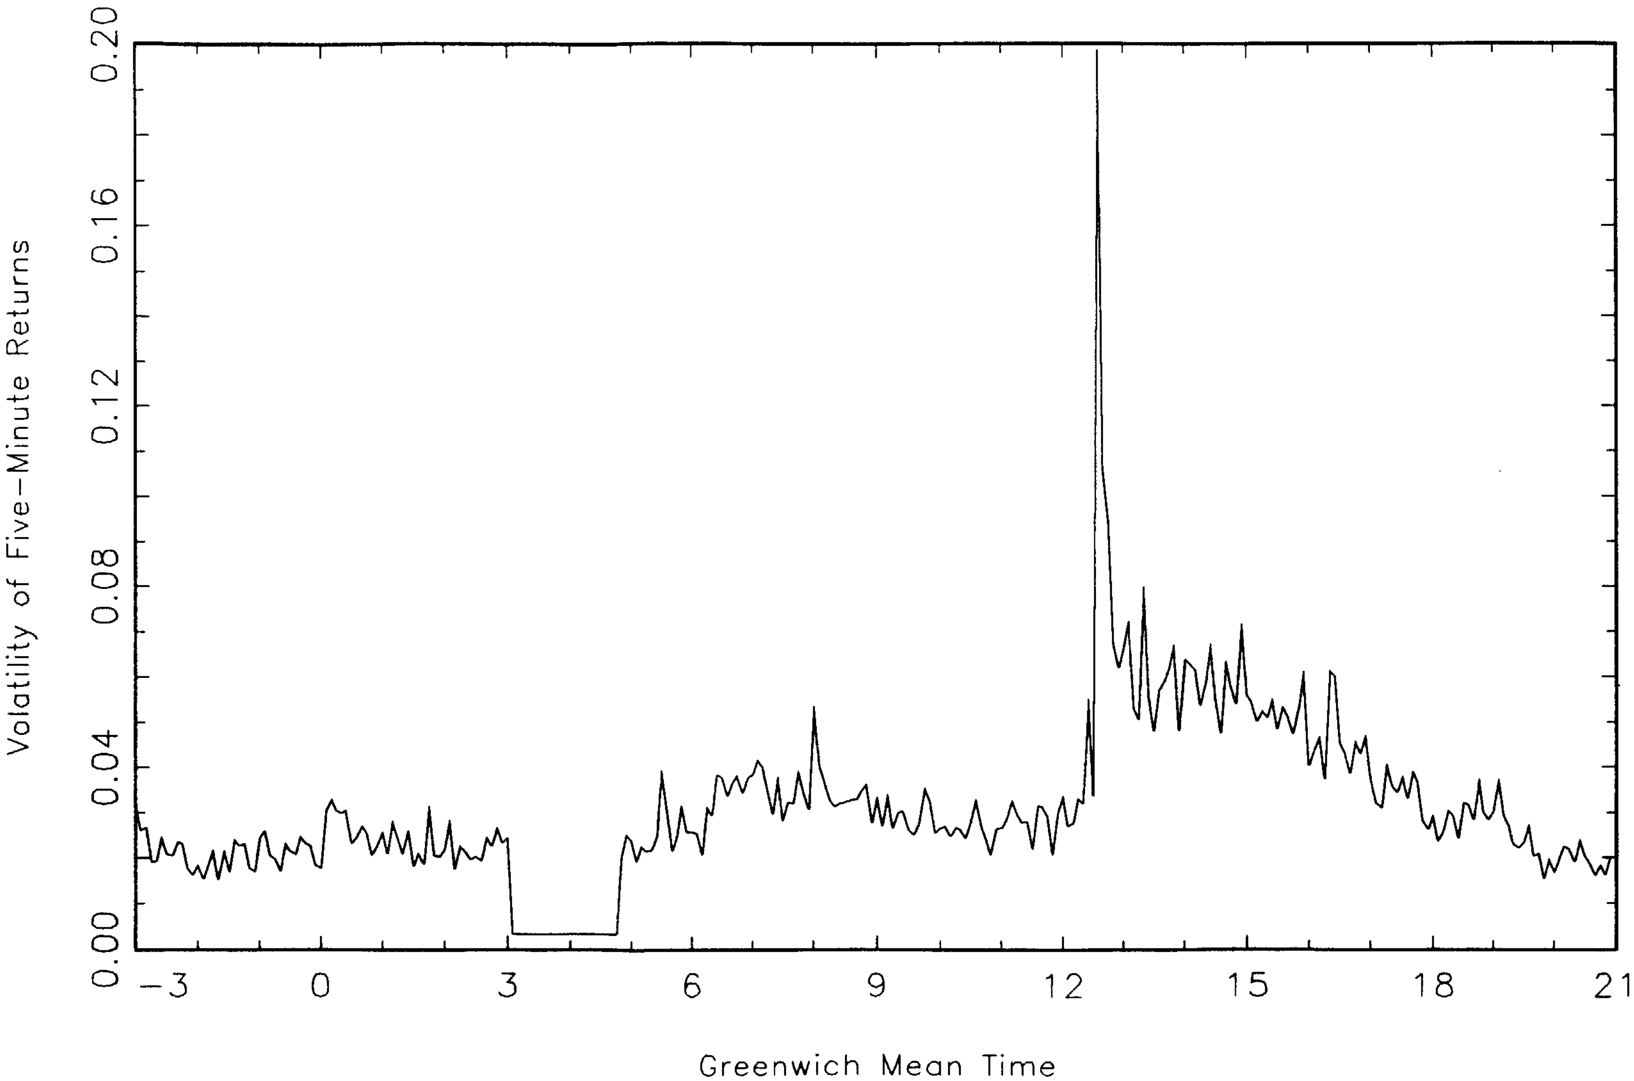

Figure 6B: U.S. Announcement Day Volatility
Winter Time

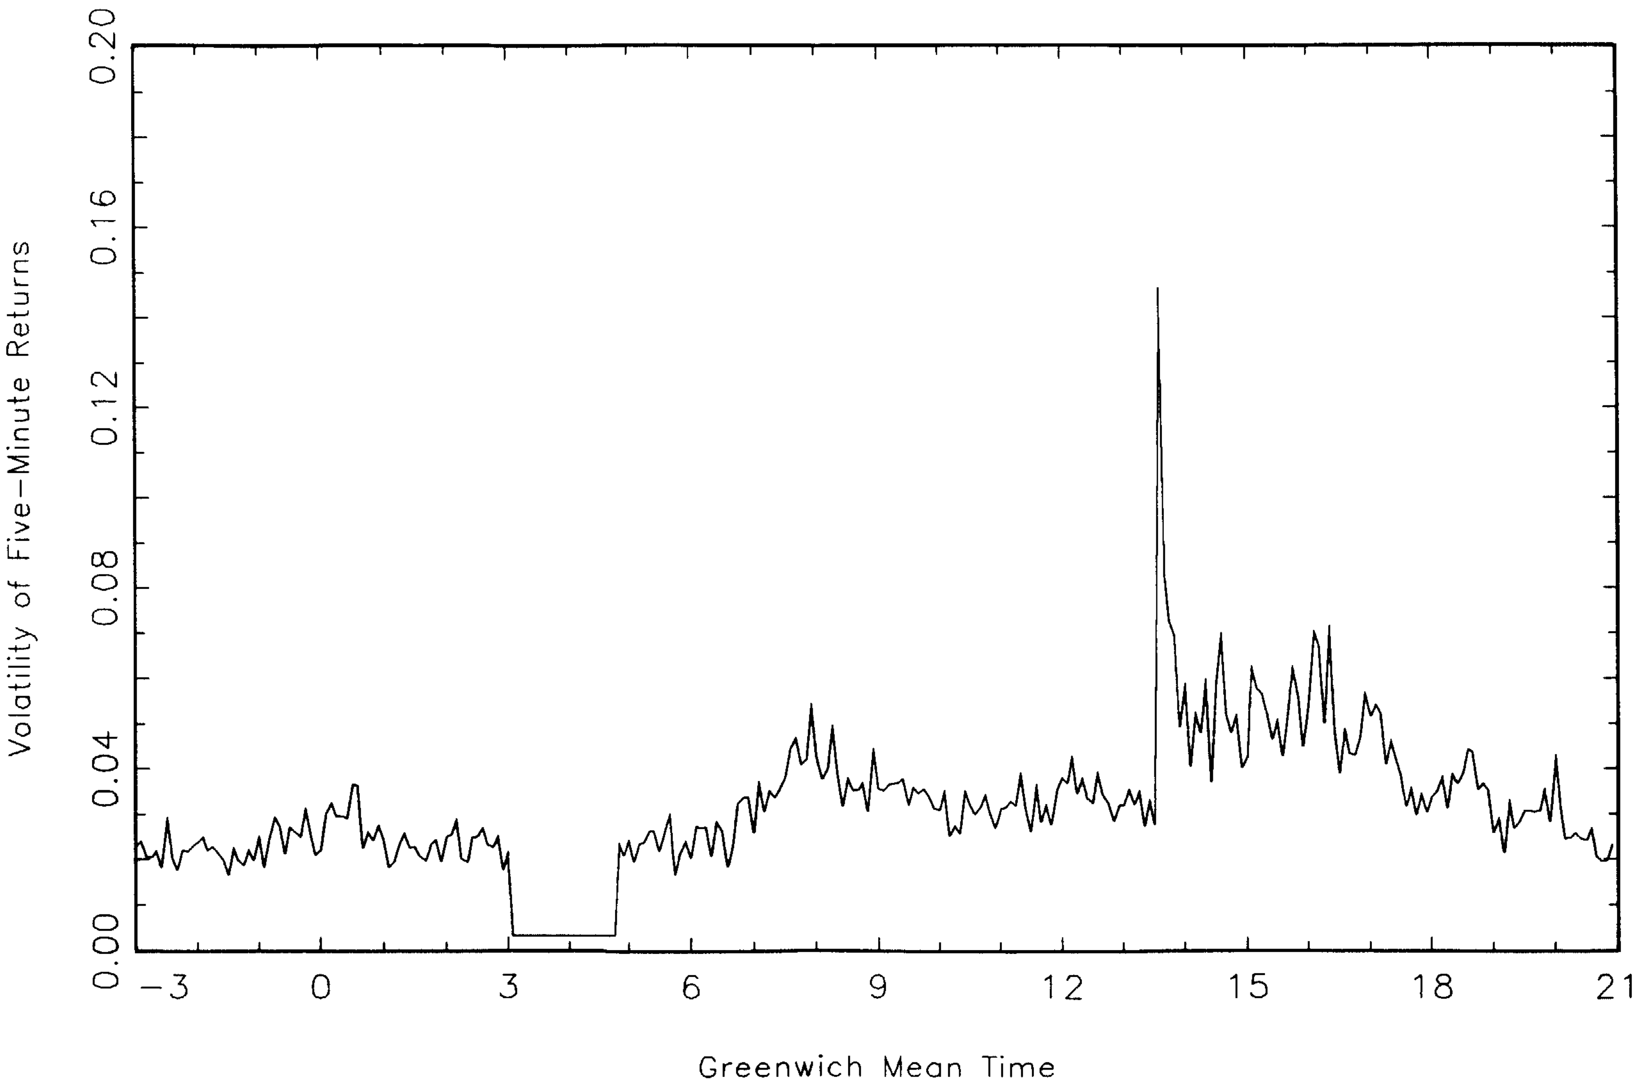

Figure 7A: Flexible Fourier Form Fit
Summer Time

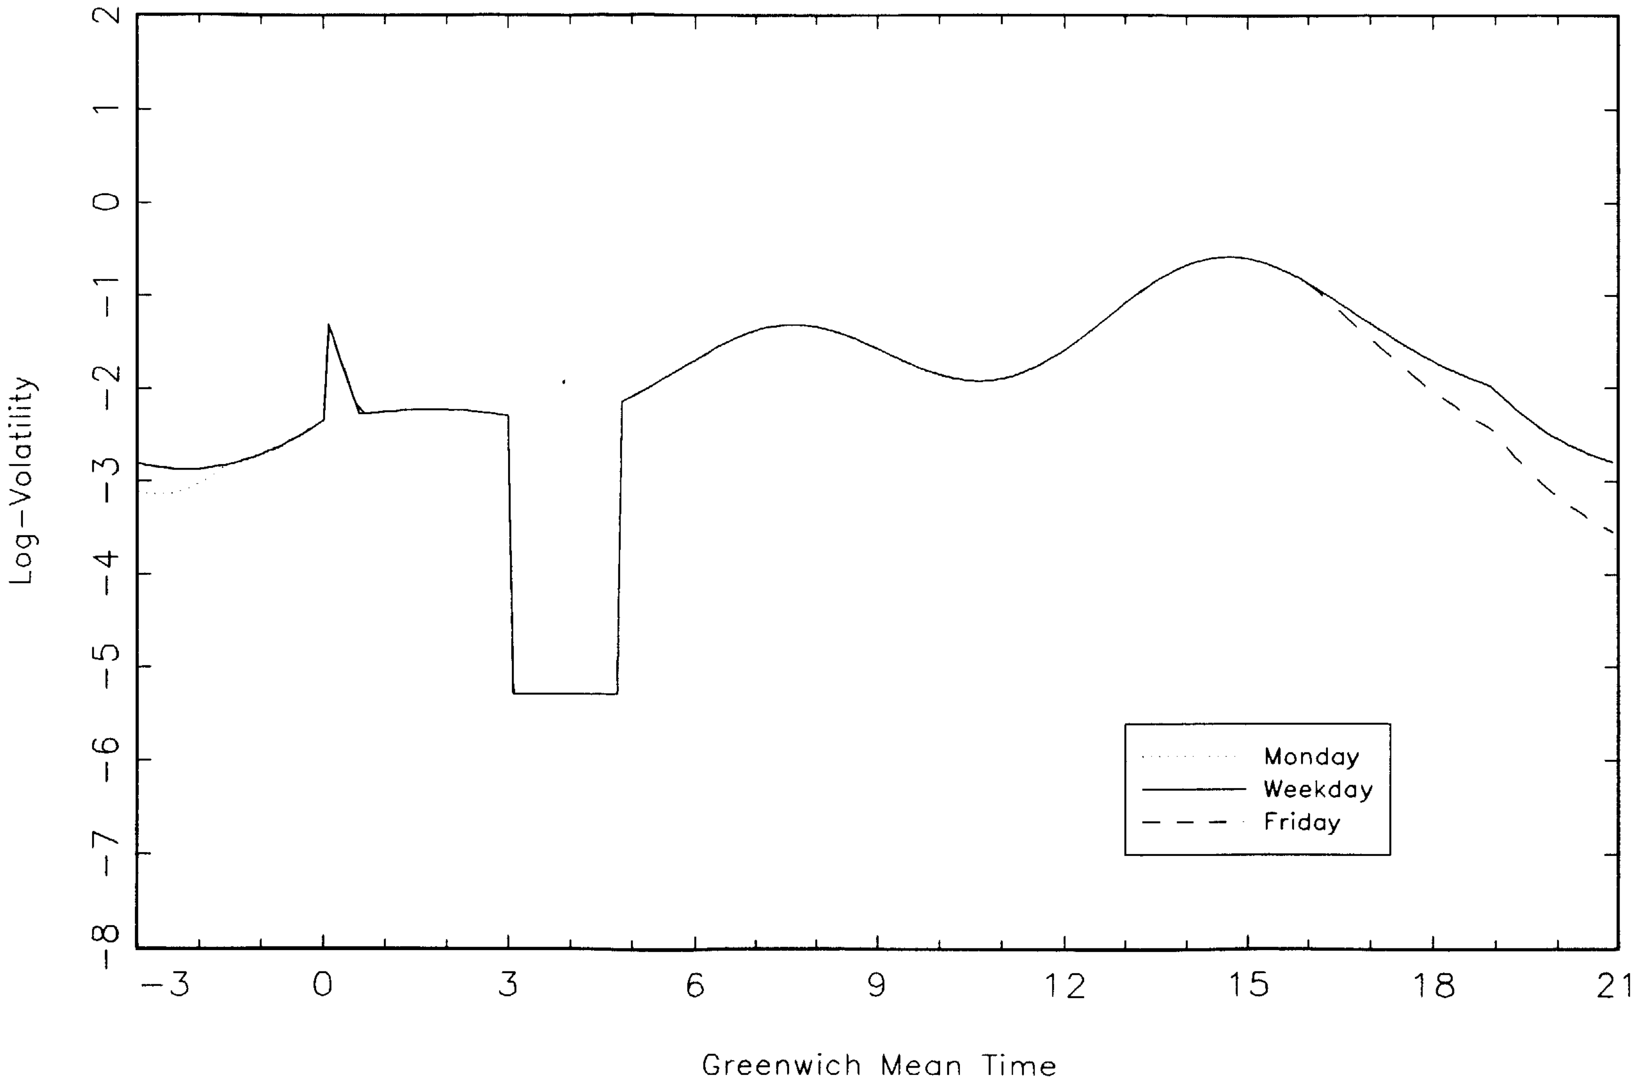

Figure 7B: Flexible Fourier Form Fit
Winter Time

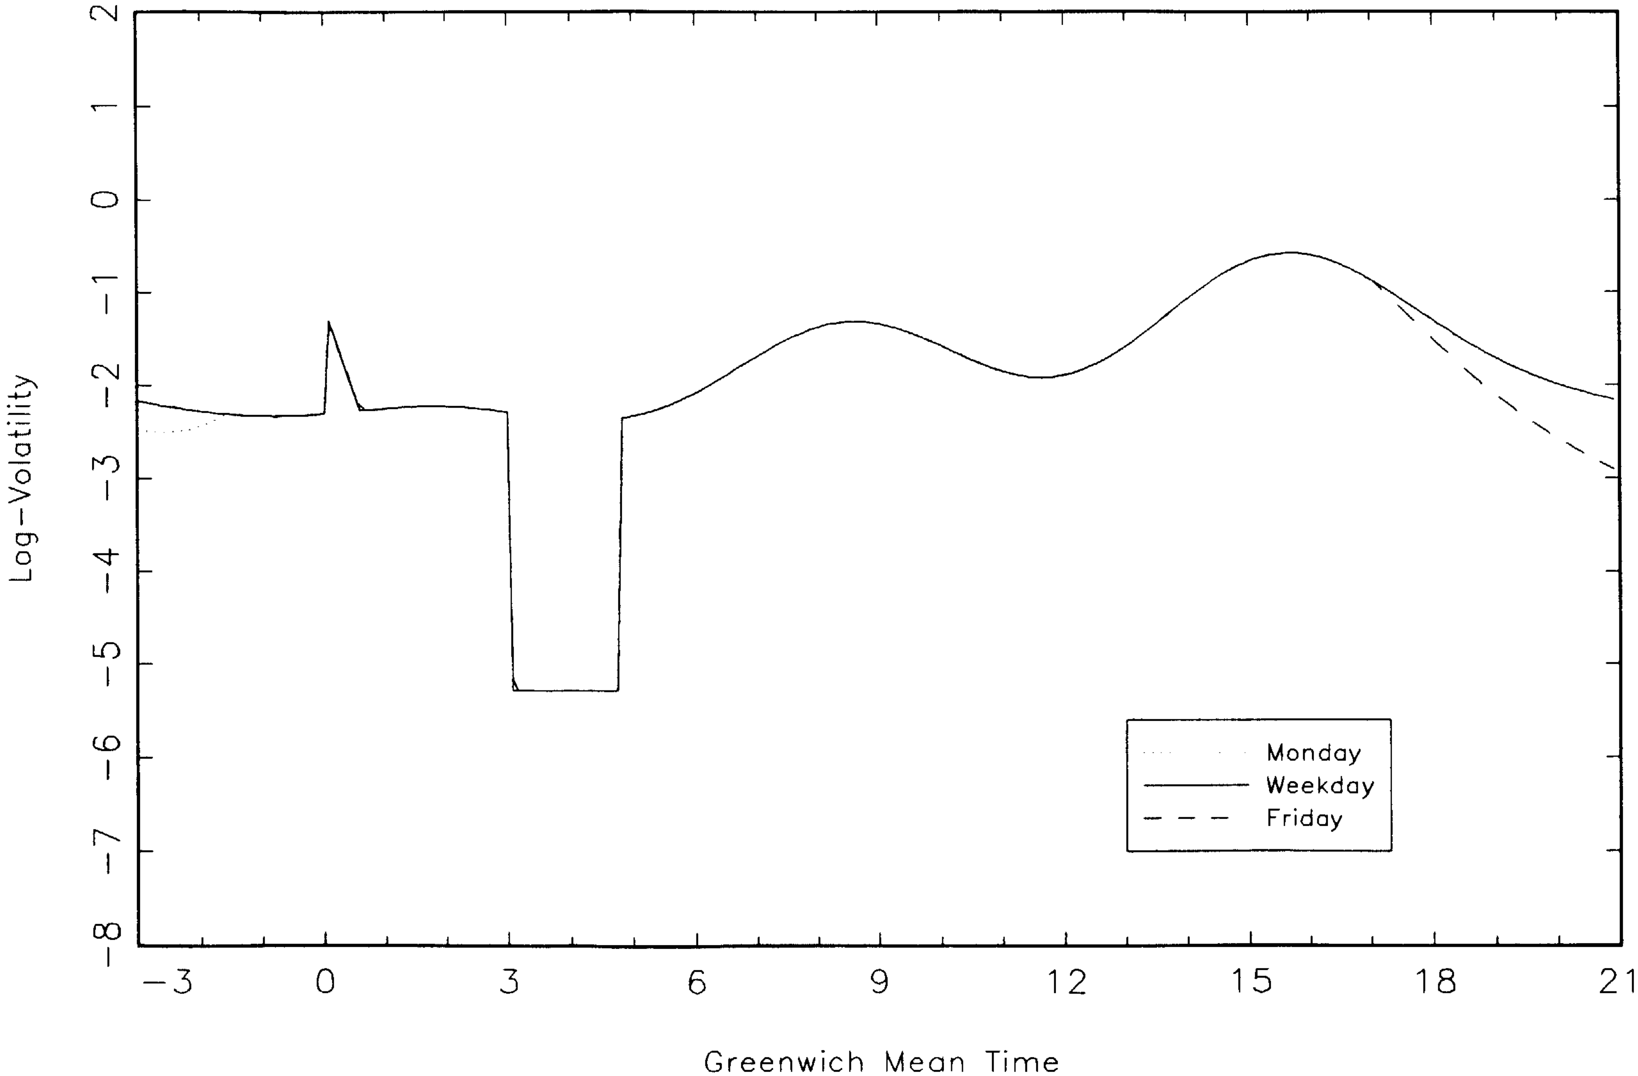

Figure 8A: Average Intradaily Log–Volatility Fit
Summer Time

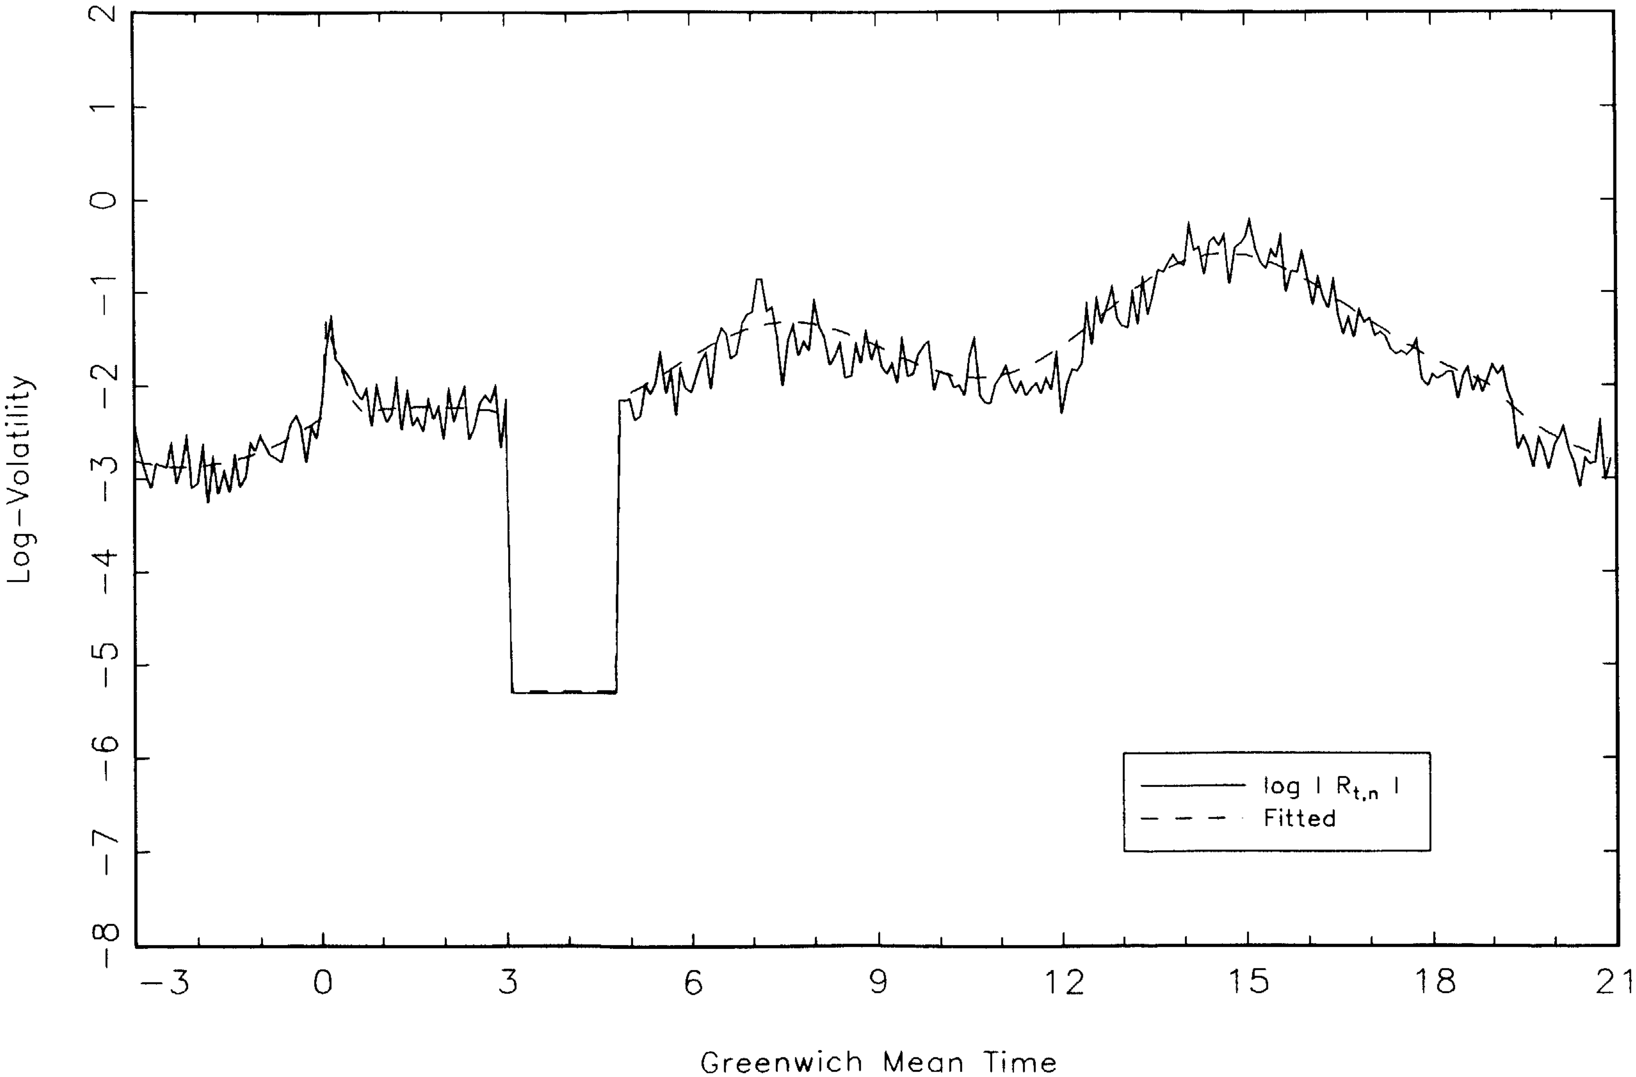

Figure 8B: Average Intradaily Log-Volatility Fit
Winter Time

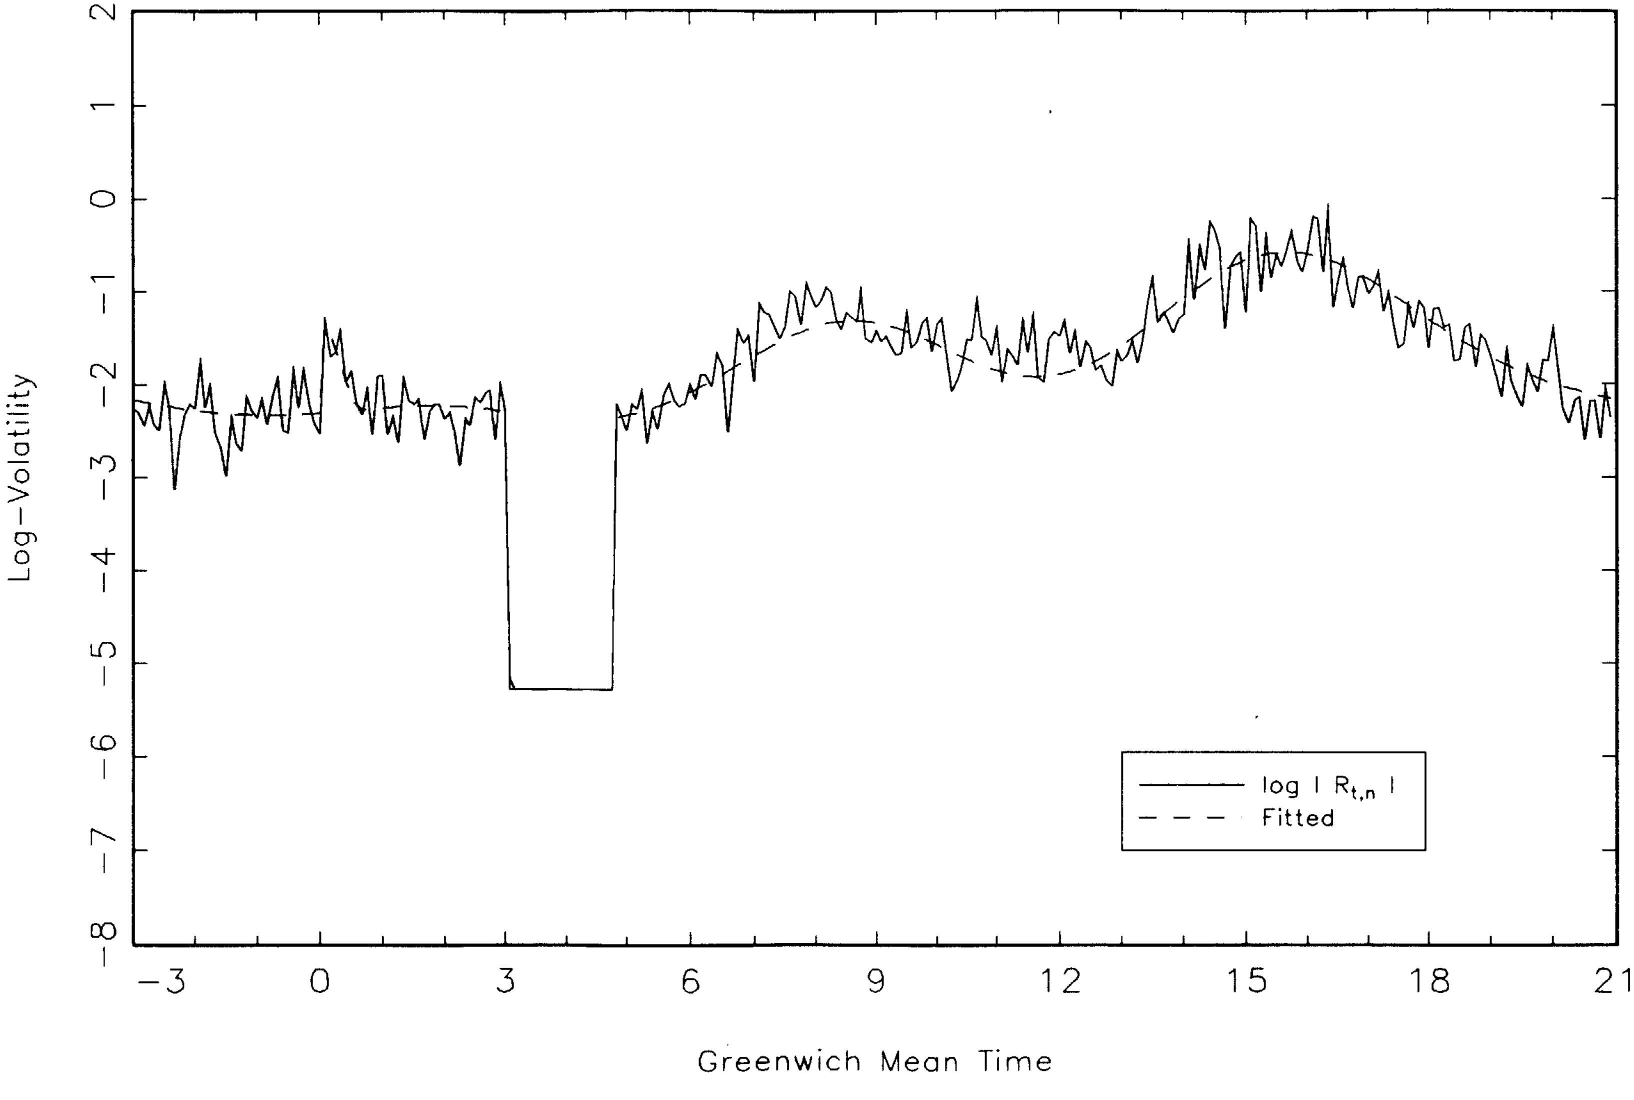

Figure 9A: Average Intradaily Absolute Return Fit
Summer Time

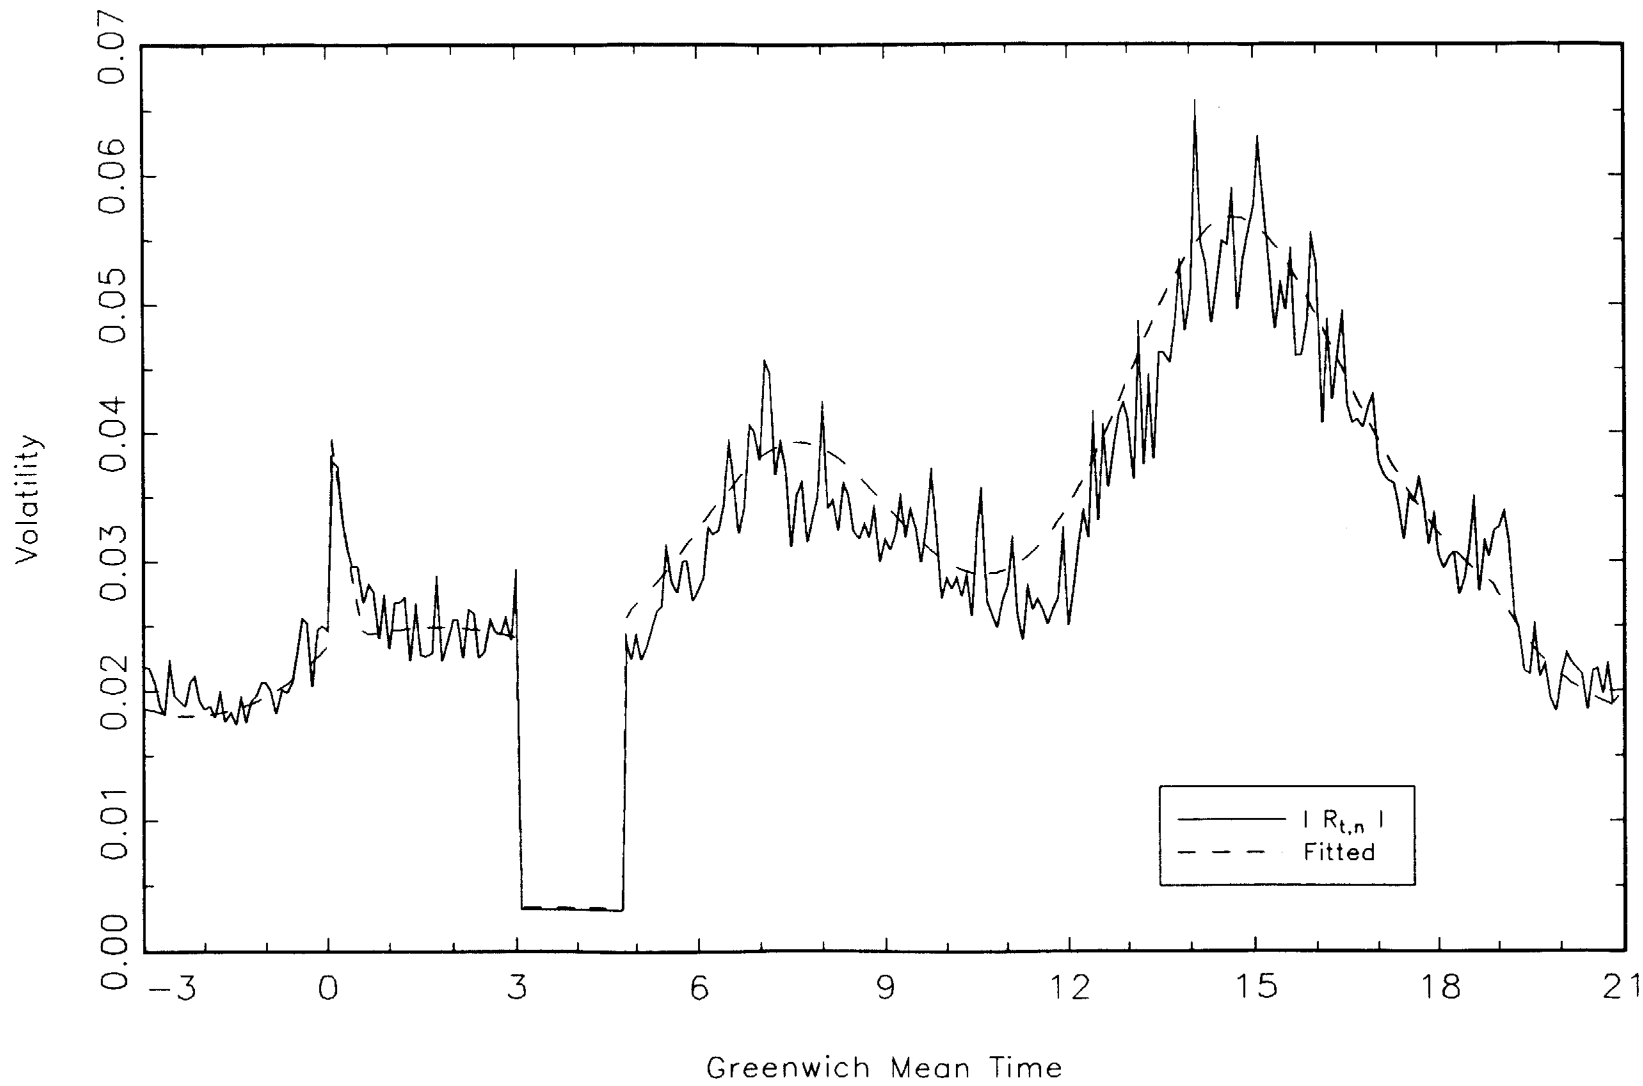

Figure 9B: Average Intradaily Absolute Return Fit
Winter Time

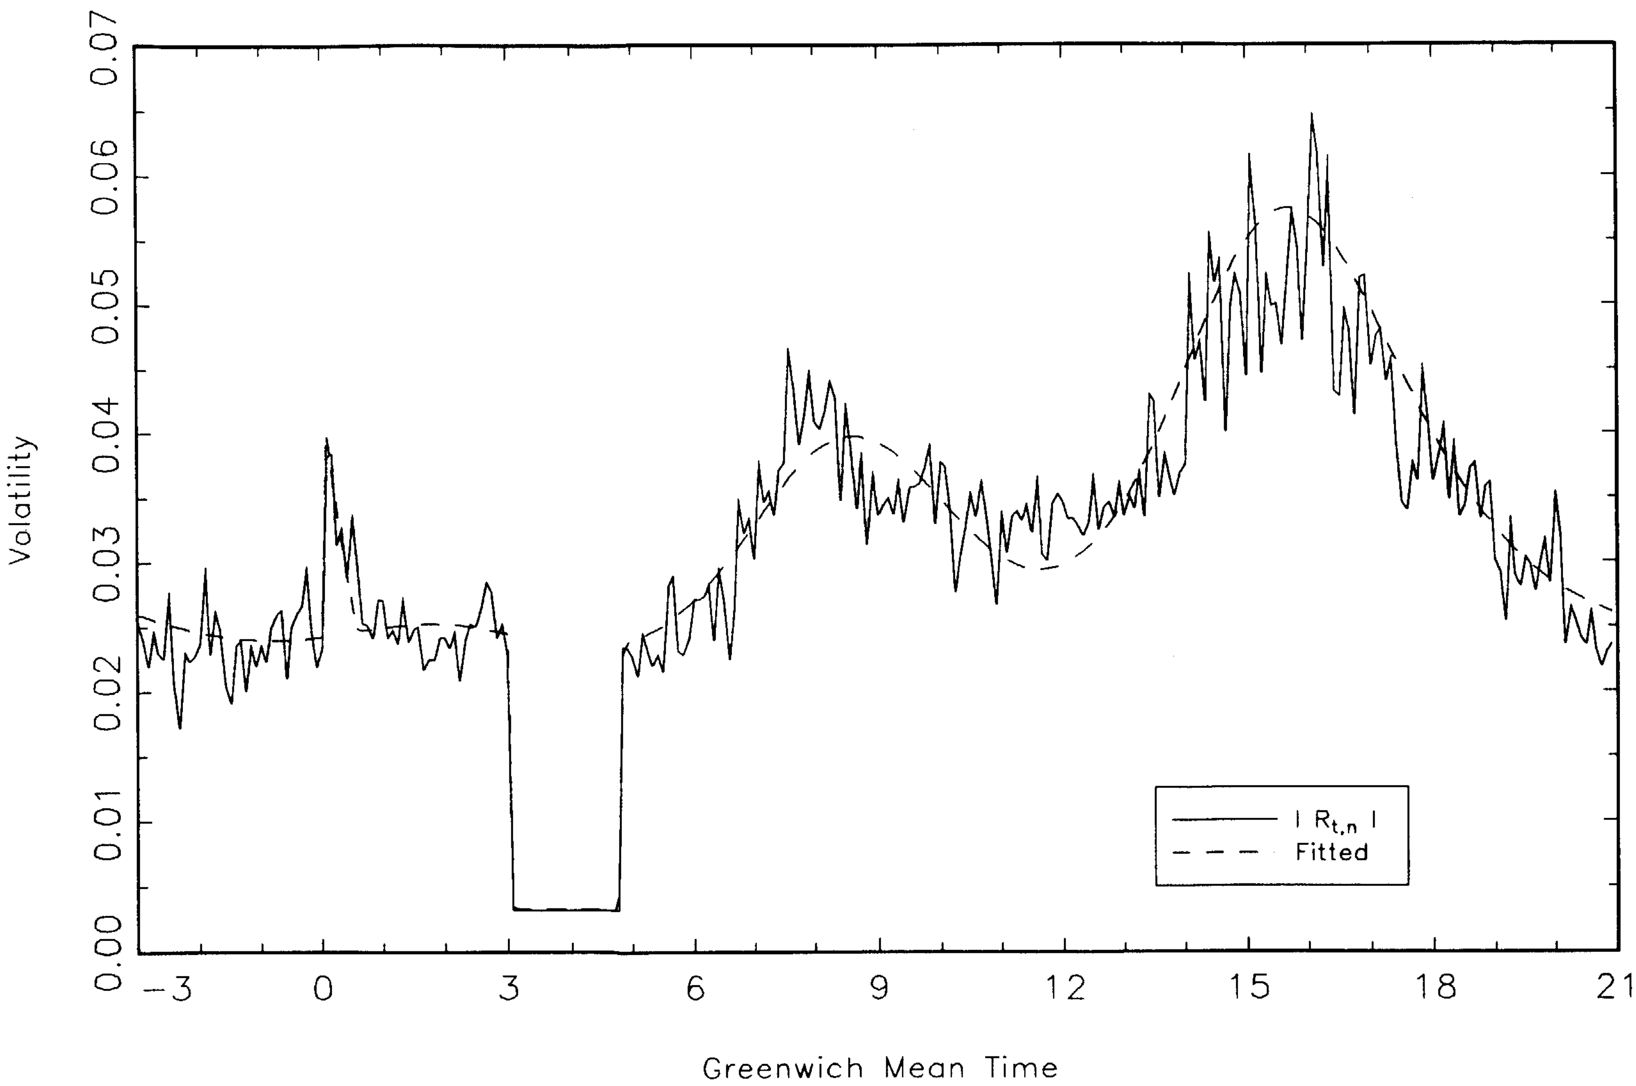

Figure 10: Dynamic Announcement Response Patterns

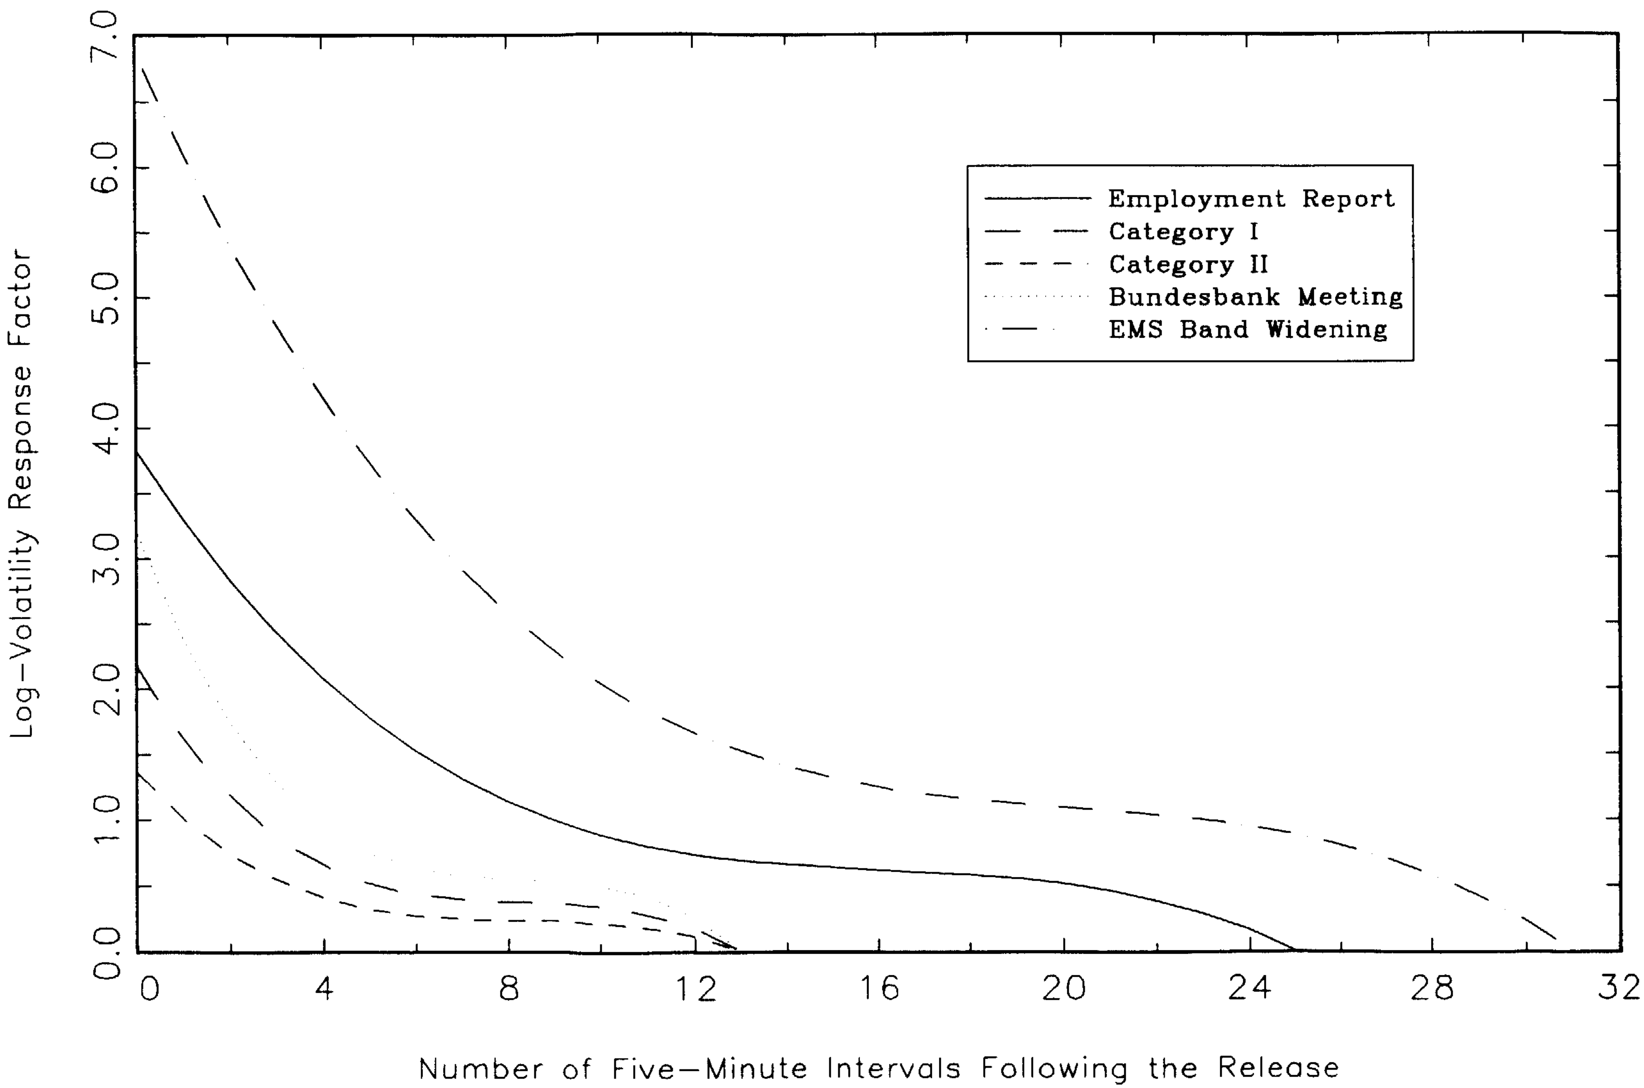

Figure 11A: Absolute Return Correlograms

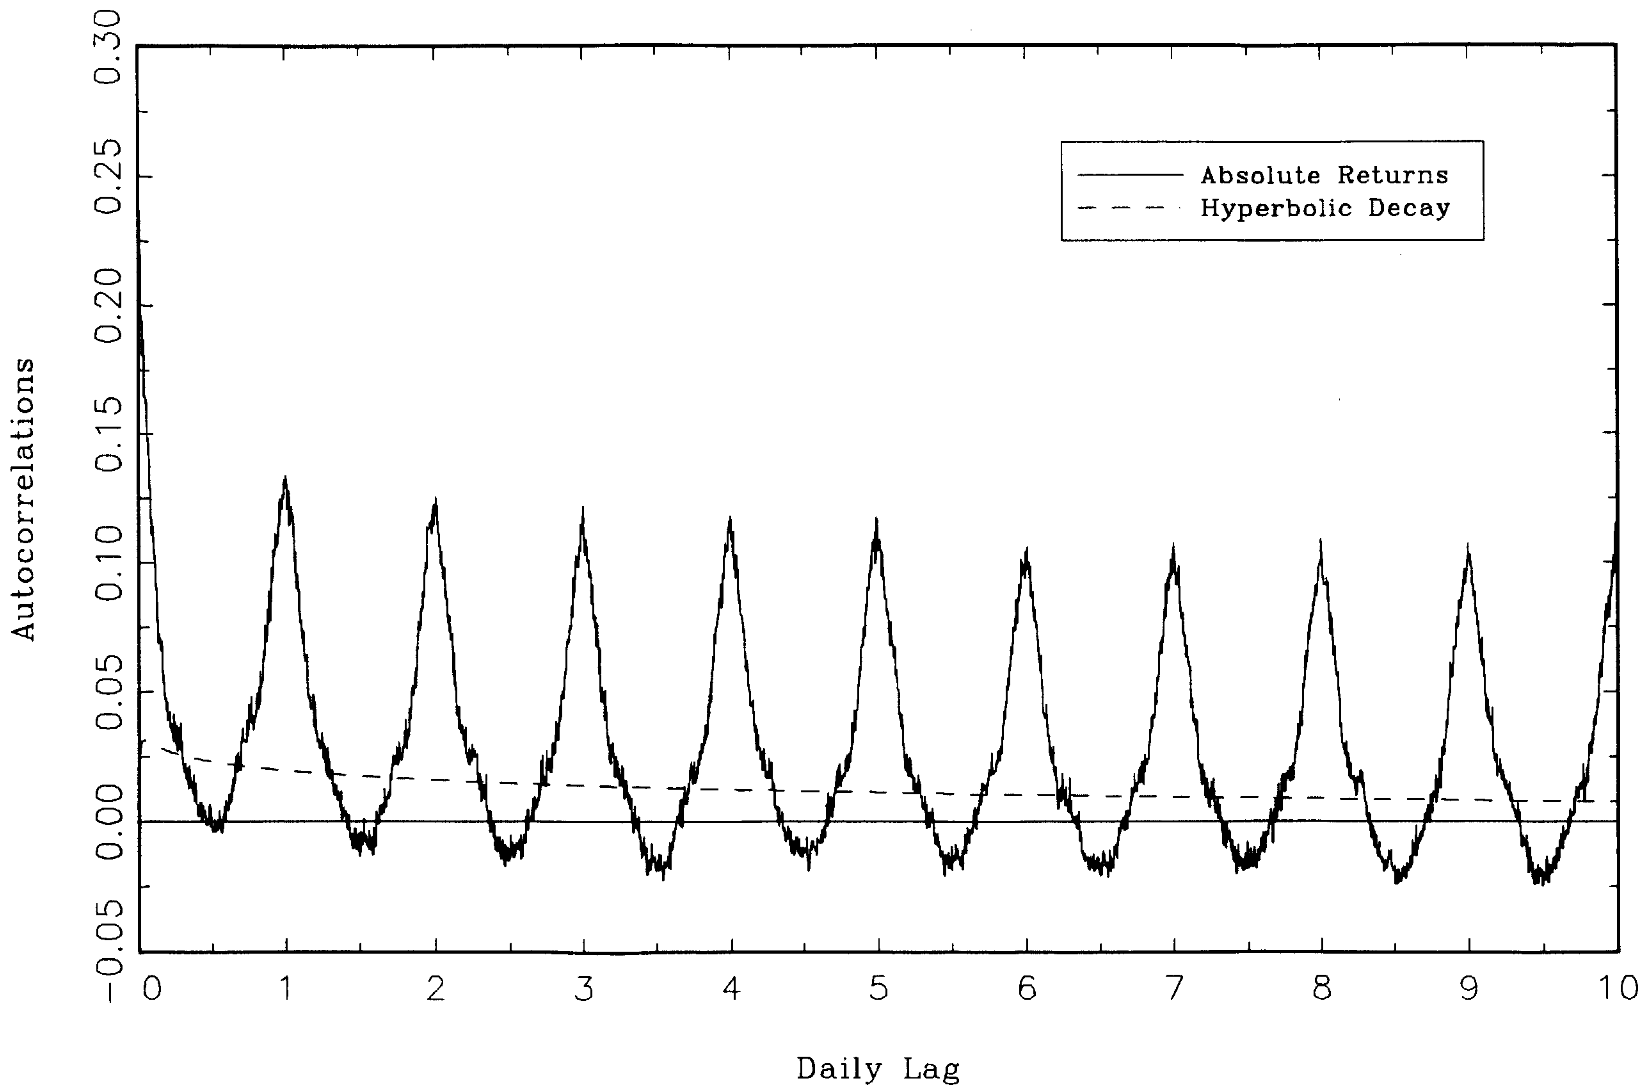

Figure 11B: Absolute Return Correlograms