

Deutsche Mark–Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies

TORBEN G. ANDERSEN and TIM BOLLERSLEV*

ABSTRACT

This paper provides a detailed characterization of the volatility in the deutsche mark–dollar foreign exchange market using an annual sample of five-minute returns. The approach captures the intraday activity patterns, the macroeconomic announcements, and the volatility persistence (ARCH) known from daily returns. The different features are separately quantified and shown to account for a substantial fraction of return variability, both at the intraday and daily level. The implications of the results for the interpretation of the fundamental “driving forces” behind the volatility process is also discussed.

THE EFFICIENT MARKET HYPOTHESIS contends that financial asset prices provide rational assessments of fundamental values and future payoffs. By implication, price changes should reflect the arrival and processing of all relevant new information. However, the proportion of return variability that can be ascribed to public news announcements is low. For example, Cutler, Poterba, and Summers (1989) find that even the largest daily price changes in the stock market generally cannot be associated with any significant economic news.¹ Moreover, proxies for the flow of public information, e.g., the number of news items released by Reuter’s News Service (Berry and Howe (1994)) or Dow Jones & Company (Mitchell and Mulherin (1994)) explain little of the overall variation. Thus, if private information is ruled out, proponents of rational price formation have to assert the existence of important news items

*Andersen is at the J.L. Kellogg Graduate School of Management, Northwestern University. Bollerslev is with the Department of Economics, University of Virginia. 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, University of Chicago, University of Houston, University of Windsor, Pennsylvania State University, 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 comments from Mike Fishman, John Heaton, Bob Hodrick, Takatoshi Ito, Ken Kavajecz, Richard Lyons, Bruce Mizrahi, an anonymous referee, and the editor, René Stulz. Needless to say, we remain fully responsible for the content.

¹ The evidence is not much more encouraging for commodities whose aggregate supply is relatively transparent (Roll (1984)), or, perhaps less surprisingly, at the individual firm level, where private information may be more prevalent (Roll (1988)).

that may be discerned by market agents, but not by economists. This is possible, but hardly convincing, and French and Roll (1986) provide strong evidence to the contrary. They document that return volatility is higher during trading than nontrading periods, independent of the flow of public news. They ascribe this to private information that is incorporated into prices through trading. One may, however, reiterate the question: What type of private information can individuals obtain that is relevant to the pricing of broad assets, while not being readily available to other investors and, even *ex post*, not being discernible to outside observers?

Furthermore, as stressed by Goodhart and O'Hara (1996), any satisfactory explanation for the behavior of volatility based on the private information hypothesis must be able to accommodate the striking empirical regularities in return volatility, not only over regular trading versus nontrading periods, but also within the trading day, within the trading week and over holiday periods. Moreover, a full account of the process governing price variability must also confront the pronounced volatility clustering, or ARCH effects, that are evident at the interday level. The one common observation across these dimensions is that "market activity" is strongly correlated with price variability. Trading volume, return volatility, and bid-ask spreads are highest around the open and close of trading; return variability per unit of time is higher over trading than nontrading periods; trading volume and spreads are particularly high on days with large return innovations; and public information releases—which theoretically may induce price jumps without any trading—are typically associated with extremely heavy volume.

Unfortunately, most empirical work has studied each of the above phenomena—the intraday and intraweekly patterns (calendar effects), the announcements (public information effects), and the interday volatility persistence (ARCH effects)—in isolation. This is ultimately not satisfactory. First, the similarity of the basic attributes across these dimensions suggests that there is a common component to the behavior which should be explicitly accounted for in any rationalization of the volatility process. Second, real-time decision-making requires that the agents recognize the various factors that exert a significant impact on current price movements. Since the factors display widely different stochastic properties, a decomposition of the contemporaneous effects is necessary in order to forecast volatility, even for the near future. Third, to judge the economic significance of the findings, it is important to gauge the influence of the features at different horizons. For example, the allocational role of the market, arguably, hinges on the information conveyed by prices at daily or lower frequencies, while studies of the market mechanism or institutions must pay close attention to the higher frequency patterns.

This paper takes a step in the direction called for above by providing a comprehensive characterization of the volatility process in the deutsche mark (DM)-dollar foreign exchange market based on a one-year sample of five-minute returns extracted from quotes on the Reuters interbank network. Although our emphasis is on the empirical identification of the various fac-

tors and their relative impact at different frequencies, we also stress the implications for theories of rational price fluctuations. The spot DM–dollar market constitutes a near ideal setting. It is the world's largest market measured by turnover, it is highly liquid, and it has low transaction costs. Moreover, it is linked globally through computerized information systems, so news is transmitted almost instantaneously to all participants. Further, it is a 24-hour market composed of sequential and partially overlapping trading in regional centers worldwide, so it has no definite closures, except those generated endogenously by the market. This allows for the study of the volatility process over periods that would be nontrading intervals under centralized market structures, thus including all public announcements, both in the United States and Germany, most of which fall outside the trading hours of the organized stock exchanges.

Our main findings are as follows: First, contrary to Cutler et al. (1989), but consistent with Ederington and Lee (1993), we show that the largest returns appear linked to the release of public information, and, in particular, certain macroeconomic announcements. Nonetheless, we conclude that these effects are secondary when explaining overall volatility. The major announcements dominate the picture immediately following the 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. Thus, high-frequency data are critical for identification of news that impacts the market, but these rather spectacular episodes do not explain what generally “drives” markets. Second, the most significant U.S. announcements, namely the employment report, gross domestic product, trade balance figures, and durable goods orders, are all related to the real economy, while the important German announcements, the Bundesbank meetings and M3 supply figures, are monetary. This may reflect differences in the (perceived) central bank reaction functions, or that, over the sample period, monetary policy in the United States was stable, while in Germany it was highly controversial. Third, the clustering of public information releases on certain weekdays explains the day-of-week effect. Hence, if announcement effects are ignored, day-dummies provide biased predictions of excess volatility on specific weekdays. For days with important scheduled announcements, such dummies fail to capture the full elevation in volatility, but the volatility will be exaggerated when no announcements are pending. Fourth, the significant calendar effects include a distinct intraday volatility pattern, reflecting the daily activity cycle of the regional centers, as well as strong holiday, weekend, Daylight Saving Time, and Tokyo market opening effects. As a group, the calendar effects are the most important determinant of overall volatility at the highest frequencies. Fifth, standard ARCH effects permeate intraday returns. Not only are these effects identifiable, but the high-frequency data provide novel insights into this long-run component of volatility. For example, return variability over a given day is typically measured by the daily absolute (squared) return, but we find that procedures exploiting the entire sequence of intraday returns provide much improved measurements. In addition, the interday ARCH component displays dependencies of the so-called

long-memory variety. This is an important result. Such features have been documented at daily and lower frequencies, but their origin has been much debated. For instance, it has been argued that infrequent exogenous shifts in the volatility process induces long-memory characteristics. Using daily data, these explanations are difficult to distinguish empirically. However, the presence of long-memory characteristics in high-frequency returns, covering a short time span, indicates that this feature is intrinsic to the system and not an artifact of exogenous shocks.

The empirical findings also have important implications for our understanding of the forces that determine the intensity of price movements. First, the pronounced activity pattern in intraday volatility suggests a significant role for the trading process itself. For trading to be informative, there must be an element of asymmetric information present in the market. Lyons (1995, 1997) develops a model of the interbank market in which dealers extract (private) signals from their customer (nondealer) order flow. This is rational if order flows speak to economic developments that only belatedly show up in statistics such as the trade balance and international capital flows. However, the aggregate net customer order flow is the variable of primary interest, and dealers only get to interpret the fraction of overall orders they receive individually. Because dealers also have a strong inventory control motive for trade, interdealer trade will not reveal, or perfectly aggregate, the information in order flows. In such settings, multiple trading rounds may be informative, even in the absence of additional news (see, e.g., Grundy and McNichols (1989), Brown and Jennings (1989), and Foster and Viswanathan (1996)). Moreover, most models stipulate that the precision of the information held by the different type of agents is common knowledge. If this assumption is relaxed, rational models again predict that the trading process provides a useful tool for inference about the structure and quality of the dispersed economy-wide information; see, e.g., Jacklin, Kleidon, and Pfleiderer (1992) and Romer (1993) for mostly theoretical arguments, and Peiers (1997) for an empirical investigation into Bundesbank interventions in the foreign exchange market.

Second, Hsieh and Kleidon (1996) document that return volatility and bid-ask spreads extracted from quotes pertaining to specific regional segments of the interbank market obey the usual u-shaped patterns that have been rationalized by clustering of informed trading (see, e.g., Admati and Pfleiderer (1988)). However, other regional segments, in the midst of their trading sessions, do not display any trace of a u-shape at the corresponding point in time. Hence, the regional u-shapes do not reflect particularly informative trading during the opening or closing hours, but may, instead, constitute a rational response to the abrupt changes in dealer exposure that occur as dealers periodically withdraw from the market place (see, e.g., Brock and Kleidon (1992) and Hong and Wang (1995)). In addition, the necessity of getting a "feel" for the market before engaging in active trading is cited as an explanation for the early parts of the regional u-shape by Hsieh and

Kleidon (1996). This is, of course, consistent with the importance of learning directly from the trading process.

Third, public information arrivals induce abrupt price changes, but the average price move is typically attained within minutes. Still, volatility and trading volume remain elevated for several hours (see, e.g., Kim and Verrecchia (1991)). If agents have identical information sets and interpret news similarly, the protracted response pattern is hard to explain, and as such provides yet another argument in favor of models with heterogeneously informed agents.

Finally, the long-memory characteristics of the volatility process speak to the potential importance of long-run persistence in the fundamental forces governing the price process. It is hard to imagine that one may rationalize such features without relying on strong dependence in the processes determining the rate of information flow and/or (private) order flow, and the associated degree of uncertainty regarding the state of the underlying economic system. On the other hand, it is clear that a study of high-frequency returns over a one-year sample can only address the issue indirectly. The origin of longer run volatility persistence remains an important topic for future research.

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. A robust regression procedure for estimation of the calendar and announcement effects is developed in Section III, and Section IV summarizes the empirical evidence with an emphasis on the instantaneous as well as the cumulative impact of each component. Section V assesses the explanatory power of the volatility factors for intraday and daily return variability. Section VI concludes. An appendix, available at the *Journal of Finance* web site, contains further discussion of the technical aspects of the estimation procedure.

I. Data Sources and Construction

Our primary data set consists of five-minute returns for the deutsche mark–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 are constructed from the DM–\$ exchange rate quotes that appear 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

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 five-minute return series being slightly less than the average quoted spread (see Bollerslev and Melvin (1994)).

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 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 do 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 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 formal model in Section III. 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 (1997a) who demonstrate that the strong intraday volatility pattern interferes with, and garbles, the time series structure of interday volatility. Only by explicitly modeling the intraday pattern is it possible to recover meaningful volatility dynamics. Nonetheless, given the poor forecasting performance of ARCH models reported in a number of recent studies, the question remains as to whether ARCH effects are significant at the highest frequencies. This section documents, to the contrary, that the ARCH fea-




Figure 1. Daily GARCH(1,1) volatility forecasts. The figure plots the one-step-ahead conditional standard deviation forecasts from an MA(1)-GARCH(1,1) model for the daily deutsche mark–dollar spot exchange rate from October 1, 1992 through September 30, 1993, for a total of 260 nonweekend days. The model is estimated with data over the longer sample period from March 14, 1979 through September 29, 1993.

tures exert a pronounced and systematic impact on the high-frequency volatility process. In addition, we show that the forecasting performance of volatility models is much better than the recent evidence suggests. The key point is that these studies rely on daily returns to gauge the realized return variability. These measures are, however, extremely noisy, and when improved measurements are obtained from intraday data, the forecast performance improves dramatically.

For concreteness, we explore the relation between one-step-ahead volatility forecasts generated by an 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 and a half months, 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

³ Although 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)).

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 discuss 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 intraday returns defined above, $R_{t,n} \equiv \log(P_{t+n/N}) - \log(P_{t+(n-1)/N})$, where $t = 1, 2, \dots, T$, and $n = 1, 2, \dots, N$. In empirical applications to high-frequency data, it is often assumed that the mean return is constant,

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

while time variation in the corresponding volatility process is allowed,

$$E|R_{t,n} - \mu| = c \cdot \sigma_{t+n/N},$$

where $c = (\pi/2N)^{1/2}$, and μ is approximately zero. One common approach used for evaluation of daily (or lower) frequency volatility estimates, say $\hat{\sigma}_t$, relies on direct comparison with the corresponding realized absolute returns, ($t = 1, 2, \dots, T$),

⁴ In 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 percent versus each other. This created intense speculation that other currencies would leave the EMS, and the volatility in October 1992 reflects 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 percent 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 percent 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, ERM 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 percent to 15 percent for all currencies except the Dutch guilder vis-à-vis the DM.

$$|R_t| = \left| \sum_{n=1}^N R_{t,n} \right| = |\log(P_t) - \log(P_{t-1})|. \quad (1)$$

Studies adopting this approach include Cumby, Figlewski, and Hasbrouck (1993), Figlewski (1995), Jorion (1995), and West and Cho (1995), among others. However, realized absolute or squared daily returns provide a very noisy measure of the underlying latent 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 volatility might instead be based on the sum of the intraday absolute returns, i.e.,

$$\sum_{n=1}^N |R_{t,n}|. \quad (2)$$

The measure in equation (2) is referred to as the *cumulative absolute returns* in the discussion below.⁵

Under simplifying assumptions it is readily shown that the return variability measure provided by equation (2) is much more efficient than the one in equation (1); see, e.g., Appendix A.⁶ The practical implication of this result is illustrated in Figure 2, which displays the two alternative volatility measures along with the corresponding forecasts from Figure 1.⁷ The weak relation between the GARCH forecasts and the absolute return measure in equation (1) is evident in Figure 2a. The daily absolute returns are scattered almost arbitrarily around the predicted values, reflecting the inherent noise in daily returns. Table I underscores the 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 0.107.⁸ Clearly, a regression of the ex post

⁵ A similar measure is used by Hsieh (1991) in calculating daily stock return standard deviations from fifteen-minute returns. There also is a large literature attempting to extract additional information about volatility from sources other than the daily returns; for example, Garman and Klass (1980), Parkinson (1980), and Rogers and Satchell (1991) explored the information content in daily observations on high and low prices, while Latané and Rendleman (1976), Day and Lewis (1992), Canina and Figlewski (1993), Jorion (1995), and Xu and Taylor (1995) studied the implied volatility extracted from option prices, and Clark (1973), Epps and Epps (1976), Tauchen and Pitts (1983), Lamoureux and Lastrapes (1990a), Gallant, Rossi, and Tauchen (1992), Andersen (1994, 1996), and Jones, Kaul, and Lipson (1994) investigated trading volume.

⁶ As mentioned earlier, the Appendix may be obtained from the *Journal of Finance* web site.

⁷ Note that the GARCH volatility estimates rely solely on the *preceding* squared daily returns and the parameter estimates obtained over the longer sample. Because 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.

⁸ The results are robust to the exclusion of holidays. Calculating the correlation between $|R_t|$ and the GARCH volatility estimate for nonholidays, rather than for the full sample, results in changes to the third decimal of the correlation number only.





Figure 2. Daily GARCH(1,1) volatility forecasts versus ex post return variability mea-
sures. The figures plot the conditional return standard deviation forecasts from an MA(1)-
GARCH(1,1) model for the daily deutsche mark-dollar returns from October 1, 1992 through
September 29, 1993, along with the corresponding realized absolute daily returns (Figure 2a)
and the corresponding cumulative absolute five-minute returns (Figure 2b). All series have
been normalized to average unity over the one-year sample.

Table I
Ex Post Correlations between Forecasts of the Daily
Deutsche Mark–Dollar Return Standard Deviation (Panel A)
or the Daily Return Variance (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 to September 29, 1993, is obtained from a MA(1)-GARCH(1,1) model estimated using daily data on the deutsche mark–dollar (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, $\sum_{n=1}^N |R_{t,n}|$ and $(\sum_{n=1}^N R_{t,n}^2)^{1/2}$, and the first entry in Panel B, $\sum_{n=1}^N R_{t,n}^2$, 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 are excluded, resulting in a total of 74,880 return observations. The last two entries in each of the panels, $|\sum_{n=1}^N R_{t,n}|$ and $|R_t|$, and $(\sum_{n=1}^N R_{t,n})^2$ and R_t^2 , respectively, 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; 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.

Panel A: Gaussian MA(1)-GARCH(1,1) estimates of σ_t			
$\sum_{n=1}^N R_{t,n} $	$(\sum_{n=1}^N R_{t,n}^2)^{1/2}$	$ \sum_{n=1}^N R_{t,n} $	$ R_t $
0.672	0.618	0.046	0.086
Panel B: Gaussian MA(1)-GARCH(1,1) estimates of σ_t^2			
$\sum_{n=1}^N R_{t,n}^2$	$(\sum_{n=1}^N R_{t,n})^2$	R_t^2	
0.660	0.066	0.107	

volatility measure on the GARCH forecasts has negligible explanatory power, with an explained variability of, at best, approximately $(0.107)^2 \approx 1.1$ percent. Given this evidence and the inadequacies of ARCH models when applied directly to intraday returns, it is perhaps understandable that many studies ignore such volatility estimates.

The fallacy of this approach, however, is evident from Figure 2b. The cumulative absolute returns from equation (2) 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 cumulative absolute returns is as high as 0.672. In other words, about $(0.672)^2 \approx 45.2$ percent of the variation in the sum of absolute intraday returns is predicted by the daily forecasts generated by a simple GARCH model.⁹ Furthermore, this variation is at the daily level. It is impossible to explain this phenomenon by the intraday volatility pattern since this is an-

⁹ This R^2 goes beyond 50 percent if simple adjustments are made for holidays with predictably low return volatility.




Figure 3. Intraday volatility pattern. The figure plots the average absolute five-minute deutsche mark-dollar (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 to 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.

nihilated when aggregated over the entire trading day. Thus, ignoring the ARCH volatility forecasts implies that a large component of *predictable* return variability is excluded from the high-frequency analysis. Clearly, a misleading picture may emerge if we fail to control for this source of common variation across the intraday returns.

B. Calendar Effects

Although there is very little evidence of predictability in the conditional mean of the five-minute DM-\$ returns, the series displays pronounced intraday volatility and activity patterns.¹⁰ 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

¹⁰ 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). This explanation is confirmed by the analysis of actual transaction prices over seven hours in Goodhart, Ito, and Payne (1996).

GMT, and 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 9 a.m. 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 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.¹¹ Although yen-\$ dealings comprise the largest portion of the Asian foreign exchange market, the yen-\$ and DM-\$ markets are intimately linked through a triangular arbitrage relationship. Thus, it is perhaps not surprising that the effects of the Tokyo market opening resemble those documented for equity markets by Wood, McInish, and Ord (1985) and Harris (1986). During lunch, 3:00–4:30 GMT, the Tokyo segment shuts down and the overall market typically approaches a standstill. This results in another market opening effect immediately thereafter. Ignoring this lunch effect, we may loosely identify a u-shaped 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 it constitutes the most active trading period and more relevant economic news may hit the market during this part of the daily cycle. Although the rough outlines of a u-shape are again visible,¹² the latter part of the pattern may, as before, 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 thus no signs of elevated volatility when trading closes down in New York. Hence, although volatility increases when each of the main regional segments becomes active, there is no direct evidence of enhanced volatility associated with the termination of regional trading.¹³ This overall pattern is consistent with the evidence reported in Baillie and Bollers-

¹¹ Prior to December 1994, The Committee of Tokyo Foreign Exchange Market Customs prohibited all authorized foreign exchange trading in Japan prior to 9 a.m., between 12:00–1:30 p.m., and after 3:30 p.m. local time (see Ito, Lyons, and Melvin (1996)).

¹² 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).

¹³ Hsieh and Kleidon (1996) show that volatility computed from quotes specific to a given region may display a u-shape for reasons unrelated to informed trading. Because regional quoting intensity is low at the early and late stages of regional trading, market volatility is primarily determined by quotes emanating from the more active segments, rendering the regional u-shape irrelevant. However, if all other regions display thin quoting activity at this point in time, the regional pattern can impact the properties of returns computed from quotes on the overall market. This provides an alternative explanation of the Tokyo market opening effect.

slev (1991), Harvey and Huang (1991), and Dacorogna et al. (1993). We now turn to a discussion of the other systematic calendar features prevalent in the 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 cause gaps of 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 effect 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 Saving 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.¹⁴

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 Eastern Standard Time (EST), or 12:30 GMT, are the source of the previously observed volatility spikes during the U.S. Summer Time regime. It displays the in-

¹⁴ As detailed in the Appendix (on the *Journal of Finance* web site), 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.




Figure 4. Daily and weekly volatility patterns. The figure displays the estimated average absolute five-minute deutsche mark–dollar (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.

traday volatility pattern for Summer days that contain 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, although there are signs of elevated volatility for several hours, the main impact seems to be gone within 10–20 minutes. Similar effects are evident for announcements at 13:30 GMT during the U.S. Winter Time regime (not displayed). 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. The latter exercise is, of course, subjective. Nonetheless, the evidence is striking, with the 7 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




Figure 5. Intradaily U.S. Summer and Winter Time volatility patterns. The figure plots the average absolute five-minute deutsche mark-dollar (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 return. All 145 weekdays during the U.S. Summer Time and the 115 weekdays during U.S. Winter Time are employed.

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 percent, 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. Market participants may have different information sets, and thus differ in their interpretation of the news, but 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 which both 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 clear-cut termination that endows them with a




Figure 6. U.S. announcement day volatility. The figure plots the average absolute five-minute deutsche mark–dollar (DM–\$) 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 during the U.S. Summer Time regime. 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. 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.

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. Modeling the Systematic Features of High-Frequency Volatility

The volatility dynamics of high-frequency foreign exchange returns are 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, as the aggregation of intraday returns would not be able to accommodate the persistent volatility processes present at the daily

Table II
Largest Absolute Five-Minute Returns Calculated
from the Deutsche Mark-Dollar Spot Exchange Rate
from October 1, 1992 through September 29, 1993

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 are 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 subjectively indicate whether any economic or political event appears to contribute to the large absolute five-minute return.

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

level.¹⁵ Thus, we stipulate that the volatility process is driven by the simultaneous interaction of numerous components, some associated with economic news releases, some with predominantly predictable calendar effects, and some with persistent, unobserved (latent) factors. We demonstrate how formal evaluation of the effects documented in Section II may be performed

¹⁵ See the theoretical results regarding temporal aggregation for ARCH in Drost and Nijman (1993) and Drost and Werker (1996), and stochastic volatility in Andersen and Bollerslev (1996b) and Ghysels, Harvey, and Renault (1996).

using a simple two-step procedure, where the final step relies on standard regression techniques.

In full generality, our model takes the following form,

$$R_{t,n} - \bar{R}_{t,n} = \sigma_{t,n} \cdot s_{t,n} \cdot Z_{t,n} \quad (3)$$

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 nonnegative; i.e., $\sigma_{t,n}, s_{t,n} > 0$ for all t, n .¹⁶

Without additional restrictions, the components of equation (3) are not separately identifiable. By squaring and taking logs, we may isolate the calendar and announcement effects, $s_{t,n}$, as the sole explanatory variables,

$$2 \log[|R_{t,n} - \bar{R}_{t,n}|] - \log \sigma_{t,n}^2 = c + 2 \log s_{t,n} + u_{t,n} \quad (4)$$

where $c = E[\log Z_{t,n}^2]$, and $u_{t,n} = \log Z_{t,n}^2 - E[\log Z_{t,n}^2]$. 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 consensus forecasts. The price and volatility reaction will reflect 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.¹⁷ 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 typically be highly correlated for the immediate period following a new release. This will induce serial correlation and heteroskedasticity in the error terms of the regression that we develop below, and we are careful to accommodate such features. Finally, we impose the analogous (weak) restriction that $\log \sigma_{t,n}$ is strictly stationary and has a finite unconditional mean, $E[\log \sigma_{t,n}]$.

¹⁶ We clearly lose some information by focusing strictly on a model for the imputed five-minute returns. The recent work of Engle and Russell (1997) is motivated by the desire to utilize all of the “ultrahigh” frequency data.

¹⁷ Payne (1996) shows 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 (1997b) 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.

In order to obtain an operational regression equation, we impose some additional structure. First, 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. Second, 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. Third, 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 that adapts well to the smooth cyclical pattern. Our choice is the following

$$f(\theta, t, n) = \mu_0 + \sum_{k=1}^D \lambda_k \cdot I_k(t, n) + \sum_{p=1}^P \left(\delta_{c,p} \cdot \cos \frac{p2\pi}{N} n + \delta_{s,p} \cdot \sin \frac{p2\pi}{N} n \right), \quad (5)$$

where $I_k(t, n)$ is an indicator for the event k during interval n on day t , θ denotes the full parameter vector to be estimated, and μ_0 , λ_k , $\delta_{c,p}$, and $\delta_{s,p}$ are fixed coefficients. Apart from the dummy variables, equation (5) is a Fourier flexible form (FFF) and may be given a semi-nonparametric interpretation.¹⁸

Assembling all the pieces, we obtain the operational regression,

$$\hat{x}_{t,n} \equiv 2 \log[|R_{t,n} - \bar{R}|] - \log \hat{\sigma}_{t,n}^2 = \hat{c} + f(\theta, t, n) + \hat{u}_{t,n} \quad (6)$$

where $\hat{c} = E[\log Z_{t,n}^2] + E[\log \sigma_{t,n}^2 - \log \hat{\sigma}_{t,n}^2]$ and the error process $\{\hat{u}_{t,n}\}$ is stationary.

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 (5). Thus, the first step provides the observable regressand and regressors for equation (6). The resulting expression constitutes a nonlinear regression in the intraday time interval, n , and the event dummies, I_k . It is parameterized by a number of sinusoids (δ -coefficients), and dummies (λ -coefficients). It is estimated, in the second step, by ordinary least squares (OLS). We refer to equation (6) and the associated OLS procedure as the *FFF regression*. This two-step method is not fully efficient, but, as argued in Appendix B, given correct specification of the first-step FFF regressor, the parameter estimates are consistent. An important advantage of the

¹⁸ The FFF is introduced by Gallant (1981). The trigonometric terms obey a strict periodicity of one day, as desired. One may add a quadratic function in the intraday interval, n , but these terms are not significant and thus simply omitted. Allowing for the intraday pattern to depend on the overall volatility level for the day, σ_t , appears important for some markets, but is not significant in this context. The more general specification is utilized in Andersen and Bollerslev (1997a).

regression specified in terms of $\hat{x}_{t,n}$, relative to, say, $R_{t,n}^2$, is that the log-transform effectively eliminates the extreme outliers in the five-minute return series, rendering the regression much more robust.

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

$$\hat{\sigma}_{t,n} = \hat{\sigma}_t / N^{1/2}. \quad (7a)$$

Alternatively, the temporal variation in $\sigma_{t,n}$ could be ignored altogether, as in the estimator,

$$\hat{\sigma}_{t,n} = \bar{\sigma} / N^{1/2}, \quad (7b)$$

where $\bar{\sigma}$ denotes the sample mean of $\hat{\sigma}_t$. 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 (7b) is that it eliminates any generated regressor problem. On the other hand, it does nothing to alleviate the heteroskedasticity. In contrast, the estimator, (7a), does provide a normalization with respect to the strong overall movements in volatility, which should improve the efficiency of the second step procedure, and allow for more accurate volatility forecasts, as documented further in Section V below.¹⁹

IV. Empirical Results

As in Section II, we report on the empirical findings in three separate sections. The first focuses on calendar, the second on announcement, and the third on ARCH effects. However, all coefficients are estimated simultaneously, so the full range of volatility features is controlled for throughout.

A. 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 Appendix C for a more

¹⁹ The robustness of the developed FFF regression is worth reiterating. The nature of conditional heteroskedasticity is left unspecified, and need not have anything 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.

extensive treatment. First, we observe that the extreme slowdown in market activity over some holidays as well as the Tokyo lunch period resembles weekends. Because we aim to characterize the overall, average volatility pattern, the systematic lack of reliable return observations over a given interval during the day 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 1:45 p.m. local time, each major holiday, and each interval associated with a failure in data transmission, are assigned the identical low, positive return, and a dummy variable is 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 these by constrained second-order polynomials over the corresponding intervals, resulting in two regression coefficients for each period. 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 Saving 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 (Summer) 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. We also 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 (1997a), Payne (1996), and Kofman and Martens (1997) for earlier specifications of this form.

We control for four different types of macroeconomic announcements in this section. The most influential is the Employment Report—the "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 incorporated as "Category I" (more important) or "Category II" (less important) releases. The former includes GDP and trade balance figures, and Durable Goods Orders, 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 have a major impact, so this effect is 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 (7a), are recorded in the second column of Table III. All coefficients associ-

ated with the intraday pattern are highly significant, except for the last sine term.²⁰ As mentioned, the volatility slowdown over the latter part of the Summer days compensates for increased activity earlier in the day due to Daylight Saving Time. The strong market opening effect in Tokyo is noteworthy; the pronounced announcement and holiday effects were expected. In contrast, the Monday morning effect is insignificant, once all 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 conventional OLS standard errors, the effect is likely an artifact of specific events that happened to occur on Fridays. When evaluated against the robust standard error the effect is decidedly insignificant.

These results justify estimation without the day-of-the-week dummies, as given in Table III, column three. The only qualitative difference is that the Friday afternoon effect now is insignificant at the 5 percent level. As a last robustness check, we estimate the identical system, imposing the constant daily volatility factor, (7b). 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 Summer Time intraday volatility pattern, as dictated by the estimates in column three, Table III, is 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 predicted and actual average absolute five-minute returns for the corresponding period in the FFF dimension underlying the estimation. The results for the Winter Time regime are similar.

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 (3) through (7) imply

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

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 θ . If we ignore potential correlation between $\hat{\sigma}_t$ and the transformed error term, we simply have to obtain the unconditional expectation, $E[\exp(\hat{u}_{t,n}/2)]$, which in turn may be estimated by averaging the corresponding expression, $\exp(\hat{u}_{t,n}/2)$, over the relevant residuals in the sam

²⁰ The large differences between the heteroskedasticity and autocorrelation consistent standard errors 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.

Table III

**Parameter Estimates for the Regression of Logarithmic Squared
Demeaned Five-Minute Deutsche Mark-Dollar Returns
on Deterministic Regressors Capturing Calendar
and Announcement Effects**

The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the deutsche mark-dollar (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 return observations. The robust standard errors reflect a Newey and West (1987) type correction incorporating 289 lags. The regression equation takes the form

$$2 \log \frac{|R_{t,n} - \bar{R}|}{\hat{\sigma}_{t,n}} = \hat{c} + \mu_0 + \sum_{k=1}^D \lambda_k \cdot I_k(t, n) + \sum_{p=1}^4 \left(\delta_{c,p} \cdot \cos \frac{p2\pi}{N} n + \delta_{s,p} \cdot \sin \frac{p2\pi}{N} n \right) + \hat{u}_{t,n},$$

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, $\hat{\sigma}_{t,n}$ 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 the right-hand side variables represent the deterministic calendar and announcement regressors. The volatility estimates, $\hat{\sigma}_{t,n}$, for interval n on day t , are obtained from an 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 $\hat{\sigma}_t$. During the U.S. Summer Time, the sinusoid regressors 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 is 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, and Category II covers U.S. releases of PPI, retail sales, housing starts, leading indicators, jobless claims, and factory orders, and the German M3 figures. Robust t-statistics are given in brackets and regular OLS t-statistics are in parentheses.

Parameter	Full System	Day-of-Week Effect Excluded	Day-of-Week, Daily Volatility Excluded
$\mu_0 + \hat{c}$	-1.77 [-32.8] (-79.4)	-1.76 [-69.2] (-155.6)	-1.85 [-56.3] (-162.0)
$\delta_{c,1}$	-0.12 [-4.41] (-8.27)	-0.13 [-4.58] (-8.62)	-0.13 [-4.78] (-8.77)
$\delta_{c,2}$	-0.13 [-4.93] (-8.16)	-0.13 [-5.09] (-8.30)	-0.13 [-5.10] (-8.26)
$\delta_{c,3}$	-0.28 [-11.8] (-18.4)	-0.28 [-12.0] (-18.5)	-0.29 [-11.4] (-18.5)

Table III—Continued

Parameter	Full System	Day-of-Week Effect Excluded	Day-of-Week, Daily Volatility Excluded
$\delta_{c,4}$	0.14 [8.10] (10.6)	0.14 [8.01] (10.6)	0.14 [8.10] (10.6)
$\delta_{s,1}$	-0.62 [-24.1] (-38.6)	-0.62 [-23.8] (-38.5)	-0.62 [-23.4] (-38.4)
$\delta_{s,2}$	-0.21 [-10.4] (-14.3)	-0.21 [-10.2] (-14.1)	-0.21 [-10.2] (-14.0)
$\delta_{s,3}$	0.17 [8.64] (11.9)	0.18 [8.76] (12.1)	0.17 [8.55] (11.8)
$\delta_{s,4}$	-0.01 [-0.68] (-0.91)	-0.01 [-0.46] (-0.62)	-0.01 [-0.68] (-0.94)
Summer slowdown	-1.14 [-5.91] (-10.6)	-1.15 [-5.95] (-10.7)	-1.08 [-5.06] (-9.93)
Tokyo opening	0.59 [8.96] (9.81)	0.58 [8.91] (9.76)	0.59 [9.06] (9.76)
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)

Table III—Continued

Parameter	Full System	Day-of-Week Effect Excluded	Day-of-Week, Daily Volatility Excluded
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)	—	—

ple.²¹ 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 for the Summer Time regime 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 (5) are readily interpreted. For example, a coefficient of unity is tantamount to the addition, in 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’s Day, 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

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

²¹ 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.




Figure 7. Flexible Fourier form fit. The figure graphs the fit to the average logarithmic-squared, normalized, and demeaned five-minute deutsche mark–dollar (DM–\$) returns across the 24-hour weekday trading cycle during U.S. Summer Time. 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 interval 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.

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 $\lambda(k, i) = \lambda_k \cdot \gamma(i)$, $i = 0, 1, \dots, N_k$, 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

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




Figure 8. Average intradaily log-volatility fit. The figure graphs the fit to the average logarithmic-squared, normalized, and demeaned five-minute deutsche mark-dollar (DM-\$) returns across the 24-hour weekday trading cycle plotted against the corresponding average 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 Time average is based on 145 weekdays.

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 percent at 9 a.m. 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 percent per interval, so the full impact is only around 0.053 percent. Because the (median) cumulative absolute return is about 9 percent over our sample, this constitutes less than 0.6 percent 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 eco-




Figure 9. Average intraday absolute return fit. The figure graphs the fit to the average absolute five-minute deutsche mark–dollar (DM–\$) returns across the 24-hour weekday trading cycle plotted against the corresponding average 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 Time average is based on 145 weekdays.

nomie significance of the early Monday effect to be negligible. The first few intervals have an estimated 14 percent reduction in volatility, but the total effect amounts to about 0.38 percent of the daily cumulative absolute returns. In contrast, the late Friday slowdown exerts a considerable effect. Due to Daylight Saving Time, separate estimates are obtained for Summer and Winter, but the reduction in volatility over the last interval of the day is 31 percent in both cases, with a cumulative impact of 3 to 4 percent at the daily level.

B. Announcement Effects

This section reports on our estimation of the volatility responses associated with regularly scheduled macroeconomic announcements in the United States, Germany, and Japan. Extensive experimentation reveals the qualitative fea-




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 deutsche mark-dollar (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.

tures of the average volatility impact to be remarkably similar across 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 adopt 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 a few releases, notably the Employment Report, follow elongated versions, so that their response horizons extend 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 are highly significant, with categories I and II consolidating those that have similar response patterns. The estimated average effects take the form displayed in Figure 10. For comparison, the figure also in-

cludes an estimated response pattern for the period following the widening of the ERM band. Because 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 is a major event, and the large and prolonged volatility response is no surprise. The ranking of the regular announcements reflects the fact that they are 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 $i = 0, 1, \dots, 32$ 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 percent. The corresponding cumulative response from equation (9) amounts to 27.17. Since a conservative estimate of the expected absolute return during 8:30 to 10:30 EST, absent announcement effects, is approximately 0.05 percent per five-minute interval, the overall effect is an elevation of volatility by $27.27 \cdot 0.05 = 1.3585$ percent. Therefore, we find about a 15 ($= 1.3585/9.0$) percent average increase in the cumulative 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 percent following Bundesbank meetings, while the cumulative impact represents an increase in the daily cumulative absolute returns of about 3.6, 2.0, and 5.1 percent, 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 two thirds 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 are 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 percent, and to have increased the cumulative absolute return on August 2, 1993, by 36.6 percent. 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.

The strict categorization in Table III is adequate for general characterization of the announcement effects, but clearly 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.

Table IV
Parameter Estimates of Specific Announcements Effects Obtained from Regressions of the Logarithmic Squared Demeaned Five-Minute Deutsche Mark-Dollar Returns on Deterministic Regressors Allowing for Calendar and Other Announcement Effects.

The returns are calculated from interpolated five-minute logarithmic average bid-ask quotes for the deutsche mark-dollar (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 return observations. The regression takes the form

$$2 \log \frac{|R_{t,n} - \bar{R}|}{\hat{\sigma}_{t,n}} = \hat{c} + \mu_0 + \sum_{k=1}^D \lambda_k \cdot I_k(t,n) + \sum_{p=1}^4 \left(\delta_{c,p} \cdot \cos \frac{p2\pi}{N} n + \delta_{s,p} \cdot \sin \frac{p2\pi}{N} n \right) + \hat{u}_{t,n},$$

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, $\hat{\sigma}_{t,n}^2$ 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 the right-hand side variables represent the deterministic calendar and announcement regressors. The volatility estimates, $\hat{\sigma}_{t,n}$, for interval n on day t , are obtained from an 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 $\hat{\sigma}_{t,n} \equiv N^{-1/2} \cdot \hat{\sigma}_t$. During U.S. Summer Time, the sinusoid regressors 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 is restricted to last one hour (13 intervals), except for the U.S. Employment Report pattern which lasts for two hours (25 intervals). All response patterns are approximated by a third-order polynomial 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 covers announcements on the U.S. GDP, trade balance, and durable goods. Category II includes U.S. releases of PPI, Retail Sales, Housing Starts, Leading Indicators, Jobless Claims, and Factory Orders, and German M3 figures. If the specific announcement under investigation belongs to one of these categories, it 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; the estimated total cumulative absolute return induced by the announcement over the assumed response horizon is measured relative to the median cumulative absolute return over the sample of 9.0 percent per day.

Panel A: Important Announcement Effects			
Announcement	Coefficient [robust <i>t</i> -stat]	Instantaneous Jump in Volatility (%)	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

Table IV—Continued

Announcement	Coefficient [robust <i>t</i> -stat]	Instantaneous Jump in Volatility (%)	Impact in Percent of Daily Cum. Abs. Return
Bundesbank meeting	1.46 [9.74]	392	5.11
Merchandise Trade	0.889 [4.24]	164	3.12
Gross Domestic Product (GDP)	0.836 [3.43]	150	2.87
Producer Price Index (PPI)	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

Panel B: Less Important U.S. Announcements

Announcement	Coefficient	Robust <i>t</i> -stat
U.S. Treasury Report	0.338	1.62
Consumer Confidence (Conference Board)	0.273	1.20
Consumer Price Index (CPI)	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

Table IV—Continued

Announcement	Coefficient	Robust <i>t</i> -stat
NAPM Survey	−0.205	−0.777
Consumer Installment Credit	−0.375	−1.02
Wholesale Sales	−0.181	−1.06
Panel C. Less Important German Announcements		
Announcement	Coefficient	Robust <i>t</i> -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

These are 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, Panel 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, Lamont, and Lumsdaine (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.²² 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 and other rates see, e.g., Hardouvelis (1984), Goodhart and Smith (1985), Hakkio and Pearce (1985), Ito and Roley (1987), 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.²³ 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, Panel 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 importance, 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 United States, and a shrinkage of the overall U.S. deficit vis-à-vis Japan. However, the small sample precludes any firm conclusions. For comparison purposes the table also 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 percent 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 Table IV, Panels B and 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 modeling daily exchange rate move-

²² Eddelbüttel and McCurdy (1996) and Chang and Taylor (1996) also find that a simple frequency count of the news headlines on the Reuters screen is positively related to the intraday DM–\$ volatility, but has low overall explanatory power.

²³ For a recent discussion of the Bundesbank monetary policy rules see Clarida and Gertler (1996). In a related context, Peiers (1997) provides an empirical analysis on the role of price leadership by Deutschebank in the DM–\$ interbank market in the two hours surrounding Bundesbank foreign exchange interventions.

ments; see, e.g., McFarland, Pettit, and Sung (1982, 1987), So (1987), Hsieh (1988, 1989), and Baillie and Bollerslev (1989). Upon running the FFF regression, using the volatility estimates from equation (6), and 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 $\exp(\frac{1}{2} \cdot 0.165) - 1 \approx 8.6$ percent, and Thursday of $\exp(\frac{1}{2} \cdot 0.118) - 1 \approx 6.1$ percent. Moreover, the effect is highly significant based on the conventional heteroskedasticity adjusted OLS standard errors, and the Friday effect remains significant at the 5 percent 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, 1 Leading Indicator, 1 Jobless Claims, and 1 Bundesbank meeting releases over the year. In total, this increases the average cumulative absolute returns on Fridays by about 5.4 percent. The unexplained gap of about 3.1 percent 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 percent. The residual 1.0 percent is comparable to the implied variation across the first three weekdays, and is 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 to 8 percent, but 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 five-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 average excess cumulative absolute returns of approximately, or less than, 5 percent of that for a typical trading day. Only the employment report is associated with a substantially higher impact. Section V further explores the significance of the documented effects for explaining overall volatility.

C. Longer-Run Volatility Components

A significant finding to emerge from our study is that the high-frequency returns contain valuable information for measurement of volatility at the daily level. Specifically, the cumulative absolute returns provide a much better ex post measure of the underlying daily latent volatility factor than either absolute or squared daily returns. These results encourage the development of new and improved techniques for the estimation and prediction of daily or lower frequency volatility that explicitly incorporate the information in high-frequency returns. In addition, as we show below, the intraday returns provide new insights that are of critical importance for the understanding of the lower frequency return dynamics.

Given our estimates of the systematic, or deterministic, calendar and announcement effects, we may filter the high-frequency returns to obtain an innovation process that retains only the purely stochastic components of the volatility process. If our modeling strategy is warranted, the properties of this (residual) return series should be largely void of calendar effects and display the type of volatility dependencies usually associated with ARCH-type processes. To investigate this hypothesis, 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 11a, is dominated by the strong periodicity at the daily frequency and does not appear particularly informative.²⁴ Figure 11b, in contrast, features a strictly positive and slowly declining correlogram. Spikes are visible at the daily frequencies, but 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 11b compares favorably to those of similarly filtered absolute returns presented in Andersen and Bollerslev (1997a) and Payne (1996). The excellent fit afforded by the hy-

²⁴ 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.




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 deutsche mark-dollar 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 are also made for extremely low quoting activity during holidays and gaps in the data series. The filtered returns in Figure 11b are obtained by standardizing the raw demeaned absolute returns by the estimated volatility impact of calendar, holiday, and announcement effects.

perbolic decay superimposed in the figures is particularly noteworthy. This rate of decay is inconsistent with ordinary ARCH models, and instead points toward a fractionally integrated, or long-memory, volatility process, as proposed by Baillie, Bollerslev, and Mikkelsen (1996) in the ARCH framework, and Harvey (1994) and Breidt, Crato, and de Lima (1995) within the context of stochastic volatility models.²⁵ The estimate of the 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 semiparametric frequency domain methods in Andersen and Bollerslev (1997b) and Henry and Payne (1996).²⁶ Thus, the long-memory characteristics appear inherent to the return series, as they manifest themselves, even over shorter 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 account for announcement and calendar effects, the high-frequency data provide important evidence on the plausibility of two alternative hypotheses that appear almost observationally equivalent from the perspective of lower frequency returns.

V. The Relative Importance of 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, but standard asset pricing models typically are specified and tested at daily or lower frequencies. Although 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. 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 of each of the three separate components; the daily volatility factor, the calendar effects, and the announcement effects. We let the indicator variable, I_σ , be unity if the daily (ARCH-based) volatility factor from equation (7a) is included in the construction of a given forecast, and

²⁵ The autocorrelations for a fractionally integrated process of order d eventually decay at the hyperbolic rate of j^{2d-1} .

²⁶ The reported estimate of 0.387 is obtained from the regression $\log(\hat{\rho}_j) = \beta_0 + \beta_1 \log(j) + u_j$, $j = 5, 6, \dots, 2670$, where $\hat{\rho}_j$ denotes the sample autocorrelation for the absolute returns, and $\hat{d} = \frac{1}{2}(\hat{\beta}_1 + 1)$; see Andersen and Bollerslev (1997b) for details.

zero if the forecast is based on a constant daily volatility factor, as in equation (7b). Formally, this 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 volatility forecast. The calendar coefficients, f_c , include the FFF sinusoids, the Tokyo open, the Daylight Saving Time, the early Monday, and the late Friday regressors, the announcement effects, f_a , signify the contribution of the four announcement regressors from Table III, and 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 would not interfere with the interpretation of the results. Based on the FFF regressors in equation (5), it is now straightforward to construct a volatility forecast from equation (8), and identify the 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

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

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, $(t = 1, \dots, T)$,

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

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

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

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 percent 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 percent), is only marginally improved by incorporating calendar effects (8.3 percent) and only slightly improved when allowing for announcement effects (11.4 percent). In contrast, the daily GARCH volatility factor alone explains 57.8 percent. 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. Although announcement effects explain a significant proportion of the return variability over shorter intervals, as also documented in Ederington and Lee (1993), this is not the case relative to the full 24-hour trading day, which is characterized by fairly large fluctuations in the overall level of volatility.²⁷ 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 drops to 15.9 percent, but more revealing is the fraction explained by the calendar effects alone (8.1 percent) relative to the announcement effects (4.9 percent) and the daily volatility factor (3.4 percent). In other words, the intraday pattern accounts for the majority of the explained variation, but 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.

VI. Concluding Remarks

The volatility process of the DM–\$ spot exchange rate market is involved, with entirely new phenomena becoming visible as one proceeds from daily returns to high-frequency intraday returns. Nonetheless, it is possible to identify three general sets of characteristics that govern the systematic features of the process. At the high-frequency level, the pronounced intraday volatility pattern is dominant, accounting for an average variation in absolute returns of more than 250 percent across the 24-hour trading cycle (after exclusion of the Tokyo lunch period). At the intraday level, the magnitude of this effect overwhelms the predictable changes in volatility captured by, e.g.,

²⁷ Subsample analysis reveals that the explained variation drops when the overall level of volatility is more stable. However, the ranking of the effects remains the same across all investigated subsamples.

Table V

**Explained Variation (R^2) from Regressions of Deutsche Mark–Dollar
Daily Cumulative Absolute Returns or Five-Minute Absolute
Returns on Alternative Absolute Return Forecasts**

For the daily cumulative absolute returns, the regression 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

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

The five-minute returns are based on interpolated logarithmic average bid–ask quotes for the deutsche mark–dollar (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

$$v(I;t,n) = (N^{-1/2}[\hat{\sigma}_t \cdot I_\sigma + \bar{\sigma} \cdot (1 - I_\sigma)]) \cdot c_0 \cdot \exp\left(\frac{\hat{f}_c(t,n)I_c + \hat{f}_a(t,n)I_a + \hat{f}_h(t,n)}{2}\right),$$

where the first term on the right-hand side 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–\$ 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 accounted for in the construction of a particular forecast. For example, the indicator vector $I \equiv (I_\sigma, I_c, I_a) = (0,1,1)$ corresponds to the model where the daily volatility is constant, and calendar (c) and announcement (a) effects are accounted for. Because the holiday effect, $\hat{f}_h(t,n)$, is included in all of the forecasts, the R^2 's reported here exceed the simple correlations given in Table I. Estimates for the time-varying daily return standard deviations over the one-year sample, $\hat{\sigma}_t$, are obtained from an 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 they are excluded otherwise.

Design	Daily Cumulative Absolute Returns	Five-Minute Absolute Returns
Complete model (I_σ, I_c, I_a) = (1,1,1)	0.606	0.159
No announcements (I_σ, I_c, I_a) = (1,1,0)	0.579	0.113
No calendar effects (I_σ, I_c, I_a) = (1,0,1)	0.603	0.084
Only daily volatility (I_σ, I_c, I_a) = (1,0,0)	0.578	0.034
No daily volatility (I_σ, I_c, I_a) = (0,1,1)	0.119	0.124

Table V—Continued

Design	Daily Cumulative Absolute Returns	Five-Minute Absolute Returns
Only announcements (I_σ, I_c, I_a) = (0,0,1)	0.114	0.049
Only calendar effects (I_σ, I_c, I_a) = (0,1,0)	0.083	0.081
Calendar + Day-of-week (I_σ, I_c, I_a) = (0,1,0)*	0.107	0.083
Only holiday effects (I_σ, I_c, I_a) = (0,0,0)	0.080	0.002

ARCH models, which rarely move by more than 25 percent over any 24-hour period. Additionally, strong but short-lived announcement effects are prevalent at the highest frequencies.

Our results verify that the high-frequency calendar and announcement effects may be estimated efficiently, even without accounting for the broader movements in daily volatility. On the other hand, real-time decision-making and the analysis of one-time events require controlling not only for the intraday pattern and the release of economic or political news, but also for the overall level of volatility. Furthermore, when analyzing the economic implications of the identified factors, it is evident that the daily volatility factor dominates at the daily and lower frequencies. Thus, it might be argued that the intraday pattern and announcement effects are of lesser economic importance, and that high-frequency data are of interest only for the area of market microstructure.

Any such conclusion, dismissing the importance of intraday return series for broader economic issues, is misleading. First, the high-frequency returns contain extremely valuable information for the measurement of volatility at the daily level. Second, the intraday returns reveal that there are significant long-memory features in the return dynamics. These features are critical for portfolio management and derivatives pricing. Moreover, they are relevant for the analysis of information transmission and volatility spillover, both contemporaneously across markets and intertemporally between the geographical regions of the identical market (see, e.g., Engle, Ito, and Lin (1990), Hamao, Masulis, and Ng (1990), and Hogan and Melvin (1994)). In summary, the information provided by high-frequency returns is valuable to a broad range of issues in financial economics, both within and beyond the realm of market microstructure.

REFERENCES

- Admati, Anat R., and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.
- Andersen, Torben G., 1994, Stochastic autoregressive volatility: A framework for volatility modeling, *Mathematical Finance* 4, 75–102.
- Andersen, Torben G., 1996, Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance* 51, 169–204.
- Andersen, Torben G., and Tim Bollerslev, 1997a, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.
- Andersen, Torben G., and Tim Bollerslev, 1997b, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long run in high-frequency returns, *Journal of Finance* 52, 975–1005.
- Baillie, Richard T., and Tim Bollerslev, 1989, The message in daily exchange rates: A conditional variance tale, *Journal of Business and Economic Statistics* 7, 297–305.
- Baillie, Richard T., and Tim Bollerslev, 1991, Intra-day and inter-market volatility in foreign exchange rates, *Review of Economic Studies* 58, 565–585.
- Baillie, Richard T., Tim Bollerslev, and Hans O. Mikkelsen, 1996, Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 74, 3–30.
- Berry, Thomas D., and Keith M. Howe, 1994, Public information arrival, *Journal of Finance* 49, 1331–1346.
- Bollerslev, Tim, and Ian Domowitz, 1993, Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance* 48, 1421–1443.
- Bollerslev, Tim, and Michael Melvin, 1994, Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis, *Journal of International Economics* 36, 355–372.
- Breidt, F. Jay, Nuno Crato, and Pedro J. F. de Lima, 1995, On the detection and estimation of long memory in stochastic volatility, *Journal of Econometrics*, forthcoming.
- Brock, William A., and Allan W. Kleidon, 1992, Periodic market closure and trading volume, *Journal of Economic Dynamics and Control* 16, 451–489.
- Brown, David P., and Robert H. Jennings, 1989, On technical analysis, *Review of Financial Studies* 2, 527–552.
- Canina, Linda, and Stephen Figlewski, 1993, The informational content of implied volatility, *Review of Financial Studies* 6, 659–681.
- Carnes, W. Stansbury, and Stephen D. Slifer, 1991, *The Atlas of Economic Indicators: A Visual Guide to Market Forces and The Federal Reserve* (Harper Collins Publishers, Washington, D.C.).
- Chang, Yuanchen, and Stephen J. Taylor, 1996, Information arrivals and intraday exchange rate volatility, Manuscript, Department of Accounting and Finance, Lancaster University, United Kingdom.
- Clarida, Richard, and Mark Gertler, 1996, How the Bundesbank conducts monetary policy, Manuscript, Department of Finance, New York University, New York.
- Clark, Peter K., 1973, A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41, 135–155.
- Cumby, Robert, Stephen Figlewski, and Joel Hasbrouck, 1993, Forecasting volatility and correlations with EGARCH models, *Journal of Derivatives* 1, 51–63.
- Cutler, David M., Jim M. Poterba, and Lawrence H. Summers, 1989, What moves stock prices? *Journal of Portfolio Management* 15, 4–12.
- Dacorogna, Michel M., Ulrich A. Müller, Robert J. Nagler, Richard B. Olsen, and Olivier V. Pictet, 1993, A geographical model for the daily and weekly seasonal volatility in the foreign exchange market, *Journal of International Money and Finance* 12, 413–438.
- Day, Theodore E., and Craig M. Lewis, 1992, Stock market volatility and the information content of stock index options, *Journal of Econometrics* 52, 267–287.
- DeGennaro, Ramon P., and Ronald E. Shrieves, 1995, Public information releases, private information arrival, and volatility in the foreign exchange market, Working paper, Department of Finance, University of Tennessee, Knoxville.

- Drost, Feike C., and Theo E. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 909–927.
- Drost, Feike C., and Bas J. M. Werker, 1996, Closing the GARCH gap: Continuous time GARCH modeling, *Journal of Econometrics* 74, 31–57.
- Eddelbüttel, Dirk, and Thomas H. McCurdy, 1996, The impact of news on foreign exchange rates: Evidence from very high frequency data, Working paper, Department of Economics, Queen's University, Kingston, Ontario, Canada.
- Ederington, Louis H., and Jae Ha Lee, 1993, How markets process information: News releases and volatility, *Journal of Finance* 48, 1161–1191.
- Ederington, Louis H., and Jae Ha Lee, 1995a, The creation and resolution of market uncertainty: The impact of macroeconomic announcements on implied volatility, Manuscript, Department of Finance, University of Oklahoma, Norman.
- Ederington, Louis H., and Jae Ha Lee, 1995b, Volatility prediction: ARCH, announcement, and seasonality effects, Manuscript, Department of Finance, University of Oklahoma, Norman.
- Engle, Robert F., Takatoshi Ito, and Wen-Ling Lin, 1990, Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market, *Econometrica* 58, 525–542.
- Engle, Robert F., and Jeffrey R. Russell, 1997, Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model, *Journal of Empirical Finance* 4, 187–212.
- Epps, Thomas W., and Mary 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, Stephen, 1995, Forecasting volatility using historical data, Manuscript, Department of Finance, New York University, New York.
- Foster, F., Doug, and S. Viswanathan, 1996, Strategic trading when agents forecast the forecast of others, *Journal of Finance* 51, 1437–1478.
- French, Kenneth R., and Richard Roll, 1986, Stock returns variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.
- Gallant, A. Ronald, 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. Ronald, Peter E. Rossi, and George E. Tauchen, 1992, Stock prices and volume, *Review of Financial Studies* 5, 199–242.
- Garman, Mark B., and Michael J. Klass, 1980, On the estimation of security price volatilities from historical data, *Journal of Business* 53, 67–78.
- Ghysels, Eric, Andrew Harvey, and Eric Renault, 1997, Stochastic volatility, forthcoming, in G. S. Maddala and C. R. Rao, eds.: *Handbook of Statistics*, vol. 14, *Statistical Methods in Finance* (North Holland, Amsterdam).
- Goodhart, Charles 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, Charles 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, Charles A. E., Takatoshi Ito, and Richard Payne, 1996, One day in June 1993: A study of the working of Reuters' dealing 2000-2 electronic foreign exchange trading system, in Jeffrey A. Frankel, Giampaolo Galli, and Alberto Giovannini, eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago).
- Goodhart, Charles A. E., and Maureen O'Hara, 1997, High frequency data in financial markets: Issues and applications, *Journal of Empirical Finance* 4, 73–114.
- Goodhart, Charles A. E., and Richard G. Smith, 1985, The impact of news on financial markets in the United Kingdom, *Journal of Money, Credit and Banking* 17, 507–511.
- Grundy, Bruce D., and Maureen McNichols, 1989, Trade and revelation of information through prices and direct disclosure, *Review of Financial Studies* 2, 495–526.
- Guillaume, Dominique M., Michel M. Dacorogna, R. R. Davé, Ulrich A. Müller, Richard B. Olsen, and Olivier V. Pictet, 1995, From the bird's eye to the microscope: A survey of new

- stylized facts of the intradaily foreign exchange markets, *Finance and Stochastics* 1, 95–129.
- Hakkio, Craig, and Douglas Pearce, 1985, The reaction of exchange rates to economic news, *Economic Inquiry* 23, 621–635.
- Hamao, Yasushi, Ronald W. Masulis, and Victor Ng, 1990, Correlations in price changes and volatility across international stock markets, *Review of Financial Studies* 3, 281–307.
- Hardouvelis, Gilas, 1984, Market perceptions of Federal Reserve policy and the weekly monetary announcements, *Journal of Monetary Economics* 14, 225–240.
- Harris, Lawrence, 1986, A transaction data study of weekly and intradaily patterns in stock returns, *Journal of Financial Economics* 16, 99–117.
- Harvey, Andrew C., 1994, Long-memory in stochastic volatility, Manuscript, London School of Economics, U.K.
- Harvey, Campbell R., and Roger D. Huang, 1991, Volatility in the foreign currency futures market, *Review of Financial Studies* 4, 543–569.
- Henry, Marc, and Richard Payne, 1996, An investigation of long range dependence in intra-day foreign exchange rate volatility, Manuscript, Financial Markets Group, London School of Economics, U.K.
- Hogan, Kedreth C., Jr., and Michael T. Melvin, 1994, Sources of heat waves and meteor showers in the foreign exchange markets, *Journal of International Economics* 37, 239–247.
- Hong, Harrison, and Jiang Wang, 1995, Trading and returns under periodic market closures, Manuscript, Sloan School of Management, MIT, Cambridge, Massachusetts.
- Hsieh, David A., 1988, The statistical properties of daily foreign exchange rates: 1974–1983, *Journal of International Economics* 24, 129–145.
- Hsieh, David A., 1989, Modeling heteroskedasticity in daily foreign exchange rates, *Journal of Business and Economic Statistics* 7, 307–317.
- Hsieh, David A., 1991, Chaos and nonlinear dynamics: Application to financial markets, *Journal of Finance* 46, 1839–1877.
- Hsieh, David A., and Allen W. Kleidon, 1996, Bid-ask spreads in foreign exchange markets: Implications for models of asymmetric information, in Jeffrey A. Frankel, Giampaolo Galli, and Alberto Giovannini, eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago).
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1996, Is there private information in the fx market? The Tokyo experiment, Manuscript, University of California at Berkeley.
- Ito, Takatoshi, and V. Vance Roley, 1987, News from U.S. and Japan: Which moves the yen/dollar exchange rate?, *Journal of Monetary Economics* 19, 255–277.
- Jacklin, Charles J., Allan W. Kleidon, and Paul Pfleiderer, 1992, Underestimation of portfolio insurance and the crash of October 1987, *Review of Financial Studies* 5, 35–63.
- Jones, Charles, Gautam Kaul, and Marc L. Lipson, 1994, Transactions, volume and volatility, *Review of Financial Studies* 7, 631–651.
- Jones, Charles, Owen Lamont, and Robin Lumsdaine, 1995, Public information and the persistence of bond market volatility, Manuscript, Department of Economics, Princeton University, New Jersey.
- Jorion, Phillippe, 1995, Predicting volatility in the foreign exchange market, *Journal of Finance* 50, 507–528.
- Kim, Oliver, and Robert E. Verrecchia, 1991, Market reaction to anticipated announcements, *Journal of Financial Economics* 30, 273–309.
- Kofman, Paul, and Martin Martens, 1997, Interaction between the London and New York stock markets during common trading hours, *Journal of International Money and Finance* 16, 387–414.
- Lamoureux, Christopher G., and William D. Lastrapes, 1990a, Heteroskedasticity in stock return data: Volume versus GARCH effects, *Journal of Finance* 45, 221–229.
- Lamoureux, Christopher G., and William D. Lastrapes, 1990b, Persistence in variance, structural change and the GARCH model, *Journal of Business and Economic Statistics* 8, 225–234.
- Latané, Henry A., and Richard J. Rendleman, Jr., 1976, Standard deviations of stock price ratios implied in option prices, *Journal of Finance* 31, 359–381.

- Lyons, Richard, 1995, Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321–351.
- Lyons, Richard, 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275–298.
- McFarland, James W., R. Richardson Pettit, and Sam K. Sung, 1982, The distribution of foreign exchange price changes: Trading day effects and risk measurement, *Journal of Finance* 37, 693–715.
- McFarland, James W., R. Richardson Pettit, and Sam K. Sung, 1987, The distribution of foreign exchange price changes: Trading day effects and risk measurement—a reply, *Journal of Finance* 42, 189–194.
- Mitchell, Mark L., and J. Harold Mulherin, 1994, The impact of public information on the stock market, *Journal of Finance* 49, 923–950.
- Müller, Ulrich A., Michel M. Dacorogna, Richard B. Olsen, Olivier V. Pictet, Matthias Schwarz, and Claude Morgenegg, 1990, Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis, *Journal of Banking and Finance* 14, 1189–1208.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity consistent covariance matrix, *Econometrica* 55, 703–708.
- Pagan, Adrian R., 1984, Econometric issues in the analysis of regressions with generated regressors, *International Economic Review* 25, 221–247.
- Parkinson, Michael, 1980, Extreme value method for estimating the variance of the rate of return, *Journal of Business* 53, 61–65.
- Payne, Richard, 1996, Announcement effects and seasonality in the intraday foreign exchange market, Manuscript, Financial Markets Group, London School of Economics, United Kingdom.
- Peiers, Bettina, 1997, Informed traders, intervention and price leadership: A deeper view of the microstructure of the foreign exchange market, *Journal of Finance* 52, 1589–1614.
- Rogers, L. Chris G., and Stephen E. Satchell, 1991, Estimating variance from high, low and closing prices, *Annals of Applied Probability* 1, 504–512.
- Roll, Richard, 1984, Orange juice and weather, *American Economic Review* 74, 861–880.
- Roll, Richard, 1988, R^2 , *Journal of Finance* 43, 541–566.
- Romer, David, 1993, Rational asset-price movements without news, *American Economic Review* 83, 1112–1130.
- So, Jacky C., 1987, The distribution of foreign exchange price changes: Trading day effects and risk measurement—A comment, *Journal of Finance* 42, 181–188.
- Tauchen, George E., and Mark Pitts, 1983, The price variability–volume relationship on speculative markets, *Econometrica* 51, 485–505.
- Thornton, Daniel, 1989, The effect of unanticipated money on the money and foreign exchange markets, *Journal of International Money and Finance* 8, 573–587.
- West, Kenneth D., and Dongchul Cho, 1995, The predictive ability of several models of exchange rate volatility, *Journal of Econometrics* 69, 367–391.
- Wood, Robert A., Thomas H. McInish, and J. Keith Ord, 1985, An investigation of transaction data for NYSE stocks, *Journal of Finance* 25, 723–739.
- Xu, Xinzhong, and Stephen J. Taylor, 1995, Conditional volatility and the informational efficiency of the PHLX currency options market, *Journal of Banking and Finance* 19, 803–821.
- Zhou, Bin, 1996, High frequency data and volatility in foreign exchange rates, *Journal of Business and Economic Statistics* 14, 45–52.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.