Journal of International Financial Markets, Institutions and Money 10 (2000) 107–130

Journal of INTERNATIONAL FINANCIAL MARKETS, INSTITUTIONS & MONEY

www.elsevier.com/locate/econbase

Intraday and interday volatility in the Japanese stock market

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

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

[b] *Department of Economics, and National Bureau of Economic Research, Duke University, Durham, NC 27708, USA*

[c] *Department of Economics and Finance, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong*

Received 1 November 1999; accepted 6 December 1999

Abstract

This paper characterizes the volatility in the Japanese stock market based on a 4-year sample of 5-min Nikkei 225 returns from 1994 through 1997. The intradaily volatility exhibits a doubly U-shaped pattern associated with the opening and closing of the separate morning and afternoon trading sessions on the Tokyo Stock Exchange. This feature is consistent with market microstructure theories that emphasize the role of private and asymmetric information in the price formation process. Meanwhile, readily identifiable Japanese macroeconomic news announcements explain little of the day-to-day variation in the volatility, confirming previous findings for US equity markets. Furthermore, by appropriately filtering out the strong intradaily periodic pattern, the high-frequency returns reveal the existence of important long-memory interdaily volatility dependencies. This supports recent results stressing the importance of exploiting high-frequency intraday asset prices in the study of long-run volatility properties of asset returns. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Nikkei 225; High-frequency data; Intraday periodicity; ARCH; Long-memory; Announcement effects

JEL classification: C14; C22; G14; G15

* Corresponding author. Tel.: + 1-847-4671285; fax: + 1-847-4915719.

E-mail addresses: t-andersen@nwu.edu (T.G. Andersen), boller@econ.duke.edu (T. Bollerslev), efjuncai@cityu.edu.hk (J. Cai)

PII: S1042-4431(99)00029-3

1. Introduction

An extensive literature has characterized the systematic features of return volatility at the daily and lower frequencies for all major primary financial assets, including foreign exchange, equity, fixed-income, and commodity related securities. Although there are some important differences across this diverse set of assets, the qualitative findings are remarkably uniform, enabling a variant of the nonlinear time-series models within the stochastic volatility and ARCH class to provide a good first-order approximation to the underlying volatility dynamics in practically all cases, see, e.g. Bollerslev et al. (1992) and Ghysels et al. (1996).

Meanwhile, research based on high-frequency, intraday financial asset prices has documented a striking diversity of the intraday return volatility process across the various asset categories and market structures. For example, bond markets appear to have a relatively constant volatility pattern across the trading day, but react quite dramatically to the release of macroeconomic news, see, e.g. Ederington and Lee (1993), Fleming and Remolona (1999). In contrast, foreign exchange markets display pronounced intraday patterns, seemingly driven by the daily activity cycle of the world's financial centers, but have a smaller, yet still highly significant, response to macroeconomic announcements, cf. Dacorogna et al. (1993), Andersen and Bollerslev (1998a). Finally, the findings of a U-shaped volatility pattern and a relatively low sensitivity to macroeconomic releases seem ubiquitous across equity markets, although the studies largely rely on US data, see Wood et al. (1985), Harris (1986) for early evidence.[1] Consequently, the similarity of the low-frequency return volatility features masks the fundamentally different real-time asset price dynamics that generate the long-run price variability across markets.

The time-series literature on volatility dynamics, the empirical market microstructure research on the shape of the intraday activity patterns, and the event studies on reactions to news releases have largely evolved independently. However, much is to be gained through an integration of these approaches. A comprehensive understanding of the real-time price discovery process is possible only when we capture the simultaneous impact of the various factors at different return horizons. Moreover, a multitude of real-time applications, including short-term trading models, option pricing and risk measurement, require joint modeling along these dimensions. Andersen and Bollerslev (1998a) take a first step in that direction by characterizing the relative explanatory power of intraday patterns, macroeconomic announcements and long-term volatility factors for intraday and daily return volatility in the foreign exchange market. In the process, they establish a number of critical empirical facts in this context. First, tests for significant volatility responses to specific events are reliable only if one controls appropriately for the persistent interdaily volatility process and the intraday pattern. Second, the intraday pattern

[1] A number of studies have sought to rationalize the U-shaped pattern in intradaily US equity volatility by the strategic interaction of asymmetrically informed agents, see, e.g. Admati and Pfleiderer (1988), Foster and Viswanathan (1990), Brock and Kleidon (1992), Slezak (1994), and Spiegel and Subrahmanyan (1995).

may only be extracted with precision if one accounts for the scheduled macroeconomic announcements. Third, there is a lot of information in the high-frequency data regarding the current state of the daily volatility factor, but both the intraday pattern and the news announcements affect the interpretation. Unfortunately, this type of comprehensive characterization has only been undertaken for a few foreign exchange rates, so the generality of the conclusions remain unclear. Among the primary obstacles are the data requirements — access to high-quality high-frequency asset prices over a relatively long calendar time period along with the associated macroeconomic news releases is indispensable.

The current paper seeks to fill this void by examining a unique dataset of 5-min returns for the Nikkei 225 equity index spanning the 4-year period from 1994 through 1997. The Japanese equity market provides a particularly attractive object of study for several reasons. First, the institutional arrangements of the Tokyo Stock Exchange (TSE) are distinctly different from the commonly analyzed US equity exchanges as well as the foreign exchange and fixed-income dealer markets. The TSE opens at 09:00 h, closes for lunch from 11:00 to 12:30 h, and closes for the day at 15:00 h. At the opening of the morning and afternoon sessions a batched trading process — *Itayose* — is used to clear the orders. Following this opening procedure, a continuous auction — *Zaraba* — is employed for the remainder of the session. The actual trading on the exchange is done by specialized security houses — *Saitori* members — who are in charge of matching the orders without taking positions themselves; see Amihud and Mendelson (1991), Lehmann and Modest (1994) and Hamao and Hasbrouck (1995) for a detailed description. These alternative market arrangements lead us to ask whether the typical U-shape for US equities applies to the Japanese market as well. Among other things, we find that the Nikkei 225 index volatility is significantly higher at the opening of the morning and the close of the afternoon sessions than during the mid-morning and mid-afternoon sessions. These features, combined with an increase in volatility immediately before and after the lunch break, result in two distinct U-shapes; one in the morning and one in the afternoon. These findings refine and extend the evidence in Lehmann and Modest (1994) who explore a courser set of half-hourly transaction and quotation returns on size-sorted portfolios. Although they do not detect an afternoon U-shape, the conclusions are consistent since our index returns are computed at higher frequencies and also, in part, reflect a staggered completion of the Itayose across individual stocks following the lunch closure. The results are also in line with the pertinent theoretical market microstructure literature and the notion that private information and asymmetrically informed agents play an important role in the price formation process. In sum, our results confirm that the unique market structure of the TSE does shape the intraday volatility pattern of Japanese equity-index returns in significant ways.

Second, recent studies suggest that intraday data contain valuable information for inference regarding the persistence of the interdaily return volatility process. Specifically, high-frequency data allow for improved volatility estimation at both daily and monthly horizons and provide a direct gauge on the strength and reliability of the so-called long-memory features of the volatility process, e.g.

Andersen and Bollerslev (1997a) and Andersen et al. (1999). The latter observation is of particular interest because samples collected over a relatively short calendar time span may help resolve the controversy regarding whether the long-memory features reflect a genuine strongly persistent volatility process rather than infrequent structural breaks in the level of volatility, as suggested by Lamoureux and Lastrapes (1990) among others. So far, comprehensive studies of the joint interaction between the various high-frequency volatility components have been limited to foreign exchange data, except for a preliminary exploration of S&P500 futures equity-index returns in Andersen and Bollerslev (1997b). Hence, the Japanese equity-index returns offer an opportunity to assess the generality of these intriguing findings. Our results provide a strong confirmation of the earlier evidence. Upon controlling for the daily periodicity in the absolute 5-min return autocorrelogram induced by the pronounced intraday activity pattern, the evidence of important long-memory volatility dependencies for returns over even long interdaily horizons is compelling. The presence of long-memory, or fractionally integrated, volatility in the Nikkei 225 returns corroborates recent results for other markets and longer time series of daily returns in Ding et al. (1993), Bollerslev and Mikkelsen (1996), Andersen and Bollerslev (1997a) and among others.

Lastly, we investigate the impact of Japanese macroeconomic announcements on the overall volatility of the market. Earlier studies of public news announcements related to the US equity market, including Cutler et al. (1989), Berry and Howe (1994), and Mitchell and Mulherin (1994), find the proportion of return variability explained by public information to be very low. These studies may, however, suffer from a lack of power as they rely on daily returns. This is not readily remedied as the most important US macroeconomic announcements are released early in the morning prior to the opening of the US stock exchanges. Such problems are circumvented with the Nikkei 225 data, since the majority of the announcements occur while the TSE is actually operating.[2] By simultaneously accounting for the intraday pattern and the interday volatility persistence, we provide new supporting evidence for the limited importance of public news releases for the overall equity return volatility process. Even so, some of the individual announcements do have a significant short-term impact on the volatility. Moreover, the inference regarding these effects is shown to be sensitive to the control for the other components of the volatility process.

The remainder of the paper is organized as follows. Section 2 provides a description of the high-frequency Nikkei 225 returns. The intraday return volatility pattern and corresponding return volatility decomposition and Flexible Fourier Form (FFF) estimation are discussed in Section 3. Section 4 presents estimates for the degree of long-memory volatility dynamics implicit in the high-frequency returns. The impact of the macroeconomic announcements is detailed in Section 5, while the overall importance of the distinct volatility components for explaining the

[2] This also permits a more direct comparison to the existing studies based on high-frequency data for the foreign exchange market, which generally report significant announcement effects; e.g. Ito and Roley (1987), Ederington and Lee (1993), Payne (1996), and Andersen and Bollerslev (1998a).

intraday and interdaily variation in the Japanese stock market volatility is assessed in Section 6. Section 7 concludes.

2. Data description and notation

The intraday Nikkei 225 index data were provided by Nihon Keizei Shimbun Incorporated and covers the 4-year period from January 2, 1994 to December 31, 1997.[3] Each data record specifies the time, to the nearest minute, and the cash index level. We partitioned the series into 5-min intervals and computed continuously compounded returns based on the last quotation within each interval.[4] The daily interval covers the period from 09:00 through 15:00 h local time, corresponding to the trading hours of the Nikkei 225 futures contract on the Osaka Stock Exchange. The initial return of the trading day, 09:00–09:05 h, largely reflects the adjustment to information accumulated overnight, and consequently displays a much higher average return variability than any other 5-min interval.[5] In order not to confound our statistical inference, we deleted this 'unusual' return interval. This leaves us with a sample of 984 trading days, each consisting of 53 intraday 5-min returns, for a total of 52 152 observations; i.e. r_τ where $\tau = 1, 2, \ldots, 52\,152$.[6]

The sample mean of this 5-min return series equals $\bar{r} = 0.00075\%$, corresponding to an annual rate of return of about 10%. However, given the sample standard deviation of 0.128, the high-frequency mean return is effectively zero from a statistical perspective. On the other hand, the first order autocorrelation of the 5-min returns equals $\rho_1 = 0.144$. This highly significant autocorrelation may be explained by a non-synchronous trading effect in the construction of the index returns; see, e.g. Chan et al. (1991) and Lo and MacKinlay (1990). This particular feature of the returns is outside the scope of the present analysis, so we adjusted the 5-min returns to eliminate the mean and first order serial-correlation effect, thus producing the residual series $R_\tau \equiv (r_\tau - \bar{r}) - \rho_1(r_{\tau-1} - \bar{r})$, where $\tau = 1, 2, \ldots, 52\,152$. This series still displays weak serial dependence at lower lags, but the effect is economically negligible, cf. Section 3 below. In the sequel we refer to this adjusted

[3] The Nikkei 225 is a price-weighted index consisting of 225 firms in the first section of the TSE selected subject to certain industry-balance considerations. It represents about 50% of the total market capitalization of the TSE.

[4] Trading suspensions and/or missing data are rare, and amount to less than two percent of the observations. In these cases, the missing quotes are determined by linear interpolation, leading to identical returns over the intermediate intervals.

[5] Lee and Mathur (1999) also document unusually high volatility in the first 5-min return of the day in their analysis of MMI, S&P500, and NYSE US equity index futures.

[6] We also examined the corresponding Nikkei 225 futures returns. Since many of these 5-min returns are zero, this creates an 'inlier' problem in the flexible Fourier form regressions discussed in Section 3 below. However, the main empirical findings remain intact. Details of these results are available upon request.

return series as the 'raw' 5-min returns.[7] For notational simplicity, we also denote these returns by $R_{t,n}$ where the subscripts refer to the nth 5-min return interval on day t, so that $n = 1, 2, \ldots, 53$, and $t = 1, 2, \ldots, 984$, respectively.

3. Intraday return periodicity and flexible Fourier form filtering

This section characterizes the intraday pattern in the Nikkei 225 5-min returns. Fig. 1 depicts the raw and absolute returns averaged across the 984 trading days. The raw returns in panel A appear to be systematically lower at the opening and closing of the morning session, and increase towards the end of the trading day. A similar tendency for low returns at the open and high returns at the close have been documented for US equity-index returns by Harris (1986). However, the orders of magnitude are extremely small and statistically insignificant.

In contrast, the average absolute returns in Fig. 1B display a strikingly regular pattern. Starting out at about 0.11% in the morning, they drop to almost half during the afternoon session, only to rise to around 0.14% towards the end. Interestingly, there is a pronounced bulge in this otherwise U-shaped curve associated with the lunch break from 11:00 to 12:30 h, along with an increase in volatility just before and after lunch. In effect, this produces the outlines of a doubly U-shaped pattern across the trading day. This basic pattern is remarkably stable across subsamples and days-of-the-week.[8] A couple of additional features are noteworthy. First, the limited increase in volatility following the lunch break reflects the fact that volatility per unit time is much lower during the lunch period than over the rest of the day.[9] Second, the increase in Nikkei 225 volatility towards the end of the day is more pronounced than observed in US equity markets. This may reflect the fact that market-on-close orders on the TSE may fail to execute given the restrictions on price fluctuations — agents eager to execute transactions within the day therefore have an added incentive to submit their orders during the continuous auction prior to the market closure.[10]

Turning to the time-series dependencies, Fig. 2 plots the autocorrelograms for the raw and absolute returns out to a lag of 530 5-min intervals, or 10 trading days.

[7] The sample skewness and kurtosis of this 5-min return series equal 0.108 and 10.3, respectively, both of which indicate significant departures from normality.

[8] The corresponding plots for different subsamples and days-of-the-week are available upon request.

[9] In a related context, Ito and Lin (1992) find that the hourly return volatility for the lunch break vis-a-vis regular trading hours on the TSE is much smaller than the comparable variance-ratio statistics for the New York Stock Exchange (NYSE). These findings are in line with earlier evidence in French and Roll (1986) related to the suspension of Wednesday trading on the NYSE in 1968, and the analysis of Saturday trading on the TSE in Barclay et al. (1990).

[10] In support of this conjecture, Muranaga (1998) reports that the settled market-on-close orders comprise less than 2% of the daily volume on the TSE — a figure consistent with earlier evidence in Lehmann and Modest (1994).

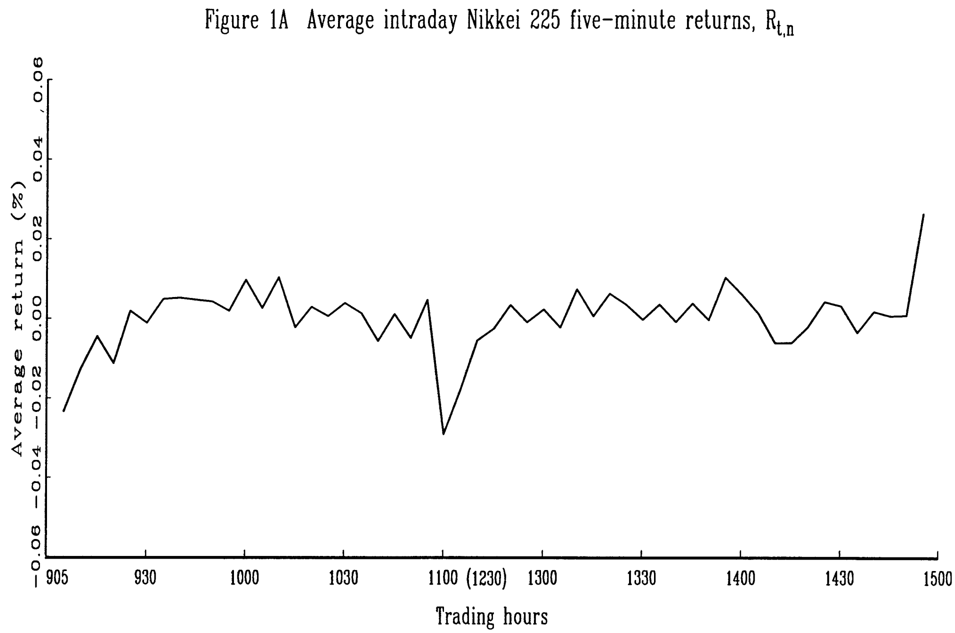

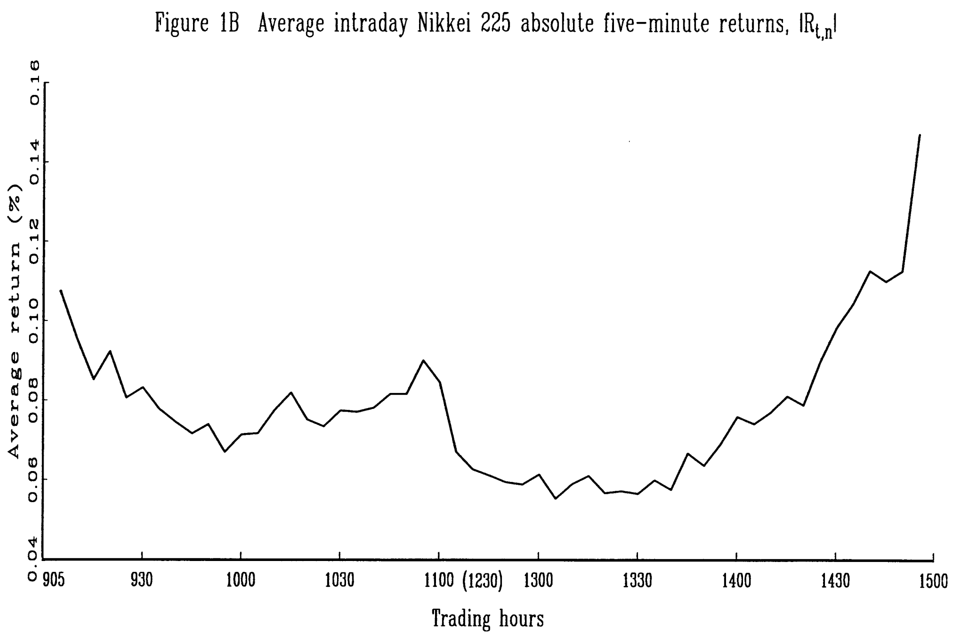

Fig. 1. Panel A graphs the average intradaily 5-min returns for the Nikkei 225 index. Panel B graphs the corresponding average absolute 5-min returns. The sample period extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

For the raw returns all of the autocorrelations in Fig. 2A are small and the overall pattern is consistent with the realization of a white noise process. In contrast, the 10-day autocorrelogram for the absolute returns in Fig. 2B displays a pronounced repetitive pattern. The strong intraday periodicity in the average absolute returns induces a distorted U-shape in the sample autocorrelations, each occupying exactly 1 day. These results are in line with the earlier evidence in Andersen and Bollerslev (1997a,b, 1998a) pertaining to US equity and foreign exchange markets. Because of this strong periodicity, standard ARCH and stochastic volatility models, which generally imply a monotone geometric decay in the autocorrelogram of the absolute returns, are not directly applicable.

In order to address this issue, we rely on the general framework developed by Andersen and Bollerslev (1997b, 1998a) for analyzing high-frequency return series. The basic idea involves the following decomposition,

$$R_{t,n} = N^{-1/2} \cdot \sigma_t \cdot s_n \cdot Z_{t,n} \tag{1}$$

where N denotes the number of high-frequency returns per day, σ_t is intended to capture the overall volatility level on day t, s_n refers to the periodic intraday volatility component, and $Z_{t,n}$ is an i.i.d. mean zero, unit variance error term.

A number of alternative parametric formulations could be used to approximate s_n and the smooth intradaily periodic pattern evident in Fig. 2B. However, the flexible Fourier form (FFF), originally proposed by Gallant (1981, 1982), is particularly convenient. Specifically, on defining,

$$x_{t,n} \equiv \log(R_{t,n}^2) + \log(N) - \log(\hat{\sigma}_t^2) \tag{2}$$

where $\hat{\sigma}_t^2$ denotes an a priori estimate of the daily volatility factor, the logarithmic periodic component, $\log(s_n^2)$, may be estimated from the linear FFF-regression,

$$x_{t,n} = \textstyle\sum_{q=0}^{Q} \delta_{0,q} \cdot n^q + \sum_{p=1}^{P} [\delta_{c,p} \cdot \cos(2\pi pn/N) + \delta_{s,p} \cdot \sin(2\pi pn/N)] \tag{3}$$

where the tuning parameters, Q and P, determine the order of the Fourier expansion. By assumption this estimate for the periodic component is constant across days. In a more general setup s_n in Eq. (1) and the corresponding FFF-regression in Eq. (3) could easily be allowed to depend upon day t information also. In particular, it is straightforward to incorporate the effects of readily identifiably exogenous events through the inclusion of simple dummy variables, say $\lambda_k \cdot I_k(t,n)$, for event k during time interval n on day t. Our analysis of day-of-the-week and macroeconomic announcement effects involves such a generalization. However, for notational convenience we continue to refer to the periodic component, s_n, by the single subscript, n.

The actual implementation of the regression in Eq. (3) requires a first-stage estimate of the daily volatility factor, $\hat{\sigma}_t$, entering the definition of $x_{t,n}$ in Eq. (2).[11]

[11] Still, the regression estimates for δ based on Eq. (3) remain consistent, albeit inefficient, under quite general conditions, including misspecification of the daily volatility model; see Andersen and Bollerslev (1998a) for further discussion.

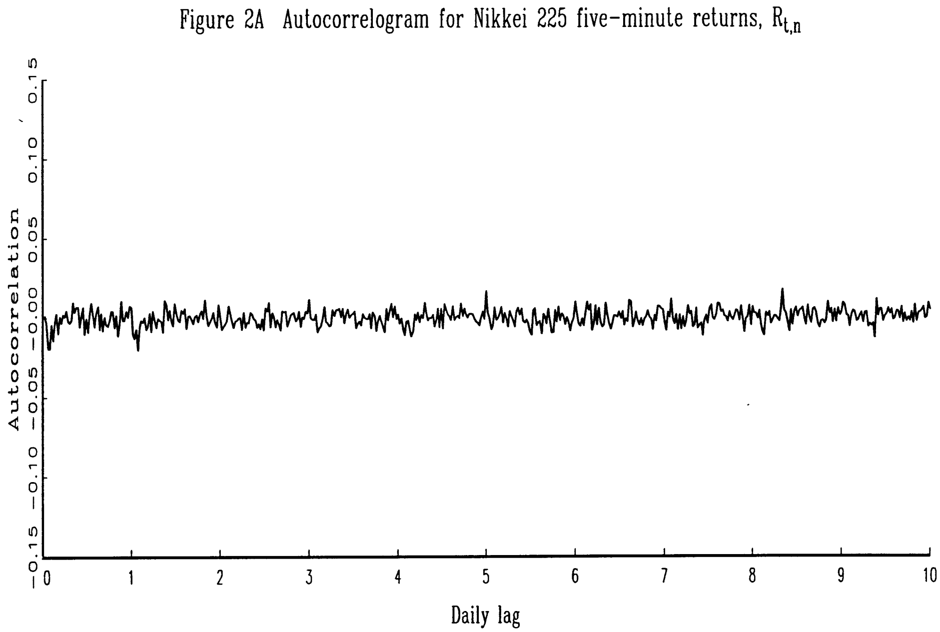

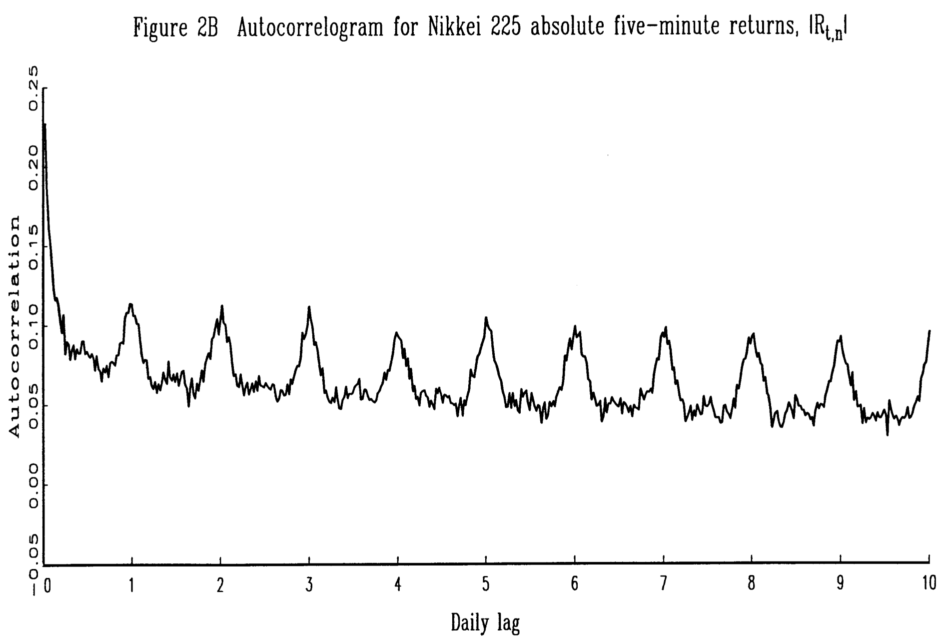

Fig. 2. Panel A gives the autocorrelogram for the 5-min Nikkei 225 returns out to a lag length of 530, or 10 trading days. Panel B graphs the corresponding autocorrelogram for the absolute 5-min returns. The sample period extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

Numerous ARCH and stochastic volatility models have been proposed for capturing the day-to-day variation in financial market volatility. Motivated by the results in the next section, our estimate for σ_t will be based on the Fractionally Integrated Exponential GARCH, or FIEGARCH, model in Bollerslev and Mikkelsen (1996). This model explicitly incorporates asymmetries, long-memory features, and the effects of market closures within a coherent modeling framework.[12]

The choice of the tuning parameters in the second-stage regression is also an empirical issue. In the present context, expanding the Fourier terms in Eq. (3) beyond $Q=2$ and $P=6$ produces insignificant estimates for any additional $\delta_{0,q}$, $\delta_{c,p}$, and $\delta_{s,p}$ coefficients. Since the approach provides a non-parametric fit to the overall intraday pattern, individual point estimates are generally void of direct economic interpretation.[13] Nonetheless, Fig. 3 indicates that the FFF-regression and the associated choice of tuning parameters provides a good approximation to the average intraday pattern.[14] Given the robustness of the FFF-approach detailed in Andersen and Bollerslev (1998a), this underscores the reliability of the features of the intraday volatility pattern discussed above.

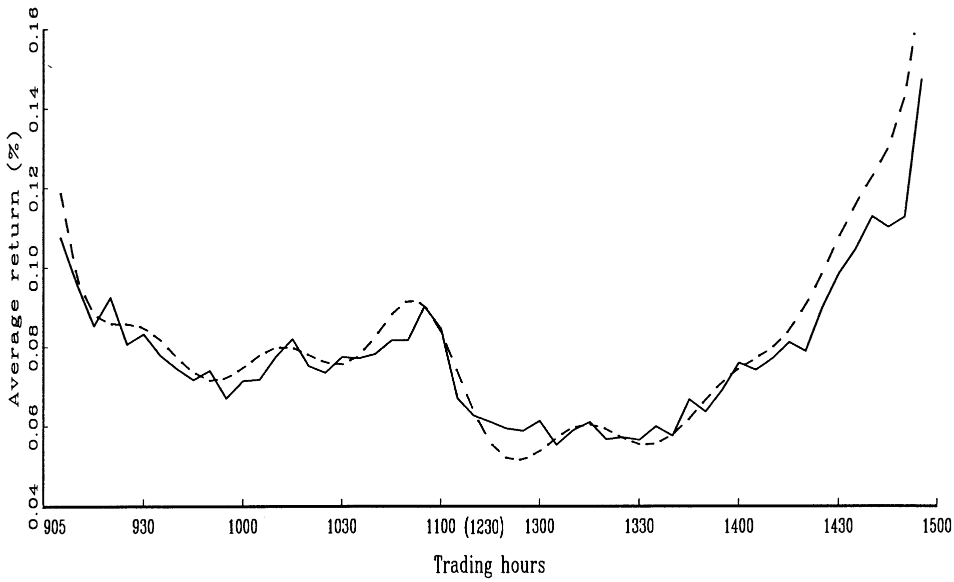

Fig. 3. The figure graphs the intradaily average absolute 5-min Nikkei 225 returns (solid line) along with the average time series estimate from the FFF-regression in Eq. (3) (dashed line). The sample period extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

[12] The actual estimation results for the specific MA(1)-FIEGARCH(1,d,1) model applied to the long time series of daily Nikkei 225 returns from September 5, 1988 through December 30, 1997, are available upon request.

[13] Details of these parameter estimates are available upon request.

[14] Let $\hat{x}_{t,n}$ denote the estimate for the right-hand-side of Eq. (3). The normalized estimate for s_n depicted in Figure 3 is then based on the standardization $982\cdot 53\cdot \exp(\hat{x}_{t,n}/2)/[\Sigma_{t=1,\ldots,982}\Sigma_{n=1,\ldots,53}\exp(\hat{x}_{t,n}/2)]$.

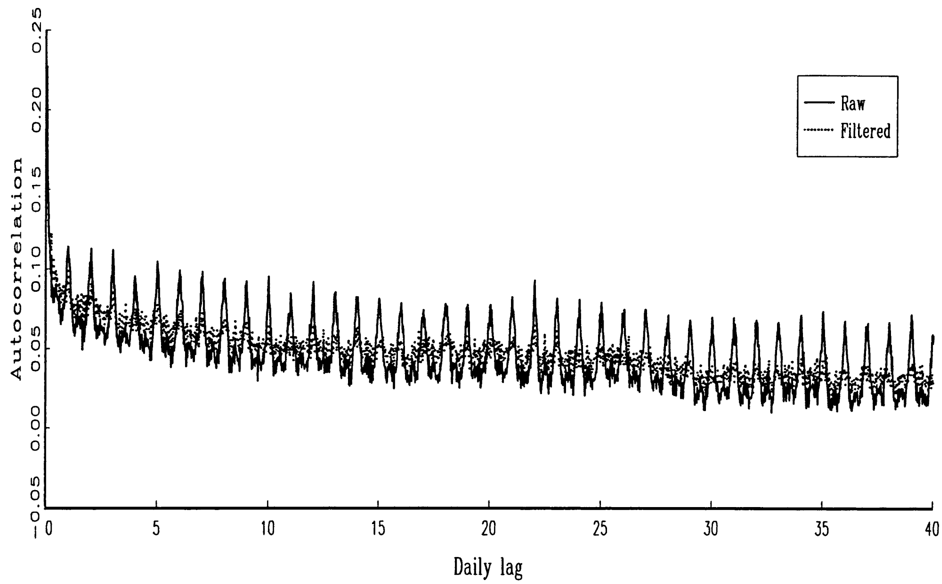

Fig. 4. The figure graphs the autocorrelogram for the absolute 5-min Nikkei 225 returns (solid line), along with autocorrelogram for the FFF-filtered absolute returns (dashed line), out to a lag length 2120, or 40 trading days. The sample period extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

In spite of the seemingly satisfactory fit from the regression approach above, the usefulness of the FFF-procedure as a method for filtering out the intraday market microstructure effects ultimately depends upon the ability of the approach to also account for the repetitive pattern of the autocorrelations in Fig. 2B. To gauge the performance along this dimension, Fig. 4 plots the 40-day autocorrelogram for the 5-min raw and filtered absolute 5-min returns, $|R_{t,n}|$ and $|R_{t,n}/\hat{s}_n|$, respectively. Clearly, the periodic dependencies have been significantly reduced, albeit not entirely eliminated, for the filtered compared to the raw absolute returns. Another striking feature of the filtered correlogram, which is hard to identify from direct analysis of the raw returns, is the rapid initial decay in the autocorrelations followed by an extremely slow rate of decay thereafter. This pattern is at odds with the exponential decay implied by most ARCH and stochastic volatility models. However, the identical hyperbolic shape of the autocorrelogram for long time series of daily absolute returns has recently been documented for a range of different assets by Ding et al. (1993), Ding and Granger (1996), Breidt et al. (1998) among others. In particular, on calculating the correlogram for the daily Nikkei 225 absolute return from 1970 to 1992, Ding and Granger (1996) report that the autocorrelations remain as high as 0.077 out to a lag of 200 days. Of course, this long-range dependence is consistent with our FIEGARCH model for the daily volatility factor, σ_t. We now turn to a more formal characterization of the long-memory dependencies implicit in the 5-min absolute returns.

4. Long-memory volatility dynamics and high-frequency returns

The strong intraday periodicity complicates the application of traditional correlation-based measures for characterizing the dynamic dependencies in the high-frequency return volatility. The semi-parametric FFF-regression provides a relatively simple procedure for filtering out the intradaily volatility patterns. Doing so, the autocorrelations in Fig. 4 suggest that the long-run dependencies in the 5-min Nikkei 225 absolute returns are extremely persistent. At the same time, formal statistical inference regarding the degree of fractional integration implied by the correlogram in Fig. 4 is complicated by a generated regressors problem. Frequency-domain procedures offer an alternative non-parametric approach for eliminating the distorting intradaily volatility dependencies. The application of such spectral-based methods in the analysis of high-frequency returns have recently been advocated by Andersen and Bollerslev (1997a). Furthermore, Bollerslev and Wright (2000) demonstrate that, in contrast to estimates for long-run dependencies in the mean where the time-span of the data is of the utmost importance in determining the accuracy of the estimates, high-frequency data allow for more accurate volatility measurements which in turn can lead to important efficiency gains when characterizing the interdaily dynamic dependencies.

A number of alternative frequency-domain methods have been proposed for estimating the degree of fractional integration, or d. We use the log-periodogram regression estimator proposed by Geweke and Porter-Hudak (1983) and later formalized by Robinson (1994a,b, 1995) and Hurvich et al. (1998). The estimator exploits that, if $|R_{t,n}|$ is long-memory, the spectrum for the process should be linear for frequencies 'close' to zero. The plot of the estimated logarithmic spectrum for the absolute 5-min Nikkei 225 returns in Fig. 5 is obviously close to being linear in the logarithm of the longer-run interdaily frequencies, or $\omega_j \geq 2\pi \cdot 53/52{,}152 \approx 0.00639$. To formalize the corresponding estimator for d, let $I(\omega_j)$ denote the sample periodogram at the jth Fourier frequency. The estimate is then obtained from the least squares regression,

$$\log[I(\omega_j)] = \beta_0 + \beta_1 \cdot \log(\omega_j) + u_j \tag{4}$$

where $j = 1, 2, \ldots, m$, $\hat{d} \equiv -\frac{1}{2} \cdot \hat{\beta}_1$, and the asymptotic standard error for $\hat{d}$ depends only on m, as in $m^{1/2}(\hat{d} - d) \sim N(0, \pi^2/24)$. Letting $m = (52\,152)^{1/2} \approx 228$, the estimate for the 5-min absolute returns is $\hat{d} = 0.429$, which is significantly different from both zero and one given the standard error of 0.042.

It is well known that the class of fractionally integrated time series models is self-similar, implying that the degree of fractional integration should be identical across different sampling frequencies; see, e.g. Beran (1994). Of course, the absolute value of the temporally aggregated 5-min returns does not equal the temporally aggregated absolute 5-min returns. Nonetheless, as shown by Andersen and Bollerslev (1997a) and Bollerslev and Wright (2000), the degree of fractional integration should remain the same across the temporally aggregated absolute returns under quite general distributional assumptions. In order to test this strong implication of the long-memory feature, we therefore also estimated the degree of fractional integration for the 10-, 15-and 30-min absolute Nikkei 225 returns. The results are striking. The estimates based on the three longer intraday sampling frequencies, with asymptotic standard errors in parentheses, are statistically indistinguishable at 0.404 (0.051), 0.482 (0.056), and 0.485 (0.068), respectively.[15] Although these estimates are not independent, the commonality across the intraday sampling frequencies and the longer daily sample suggest that the long-memory feature is an inherent property of the Nikkei 225 volatility. Further evidence is provided by the time-domain correlogram for the fractionally differenced absolute 5-min-returns, $(1-L)^{0.429}|R_{t,n}|$, in Fig. 6.[16] In contrast to the overall slow hyperbolic rate of decay in the correlograms for the raw and filtered absolute returns in Fig. 4, the autocorrelations for the fractionally differenced absolute returns are concentrated

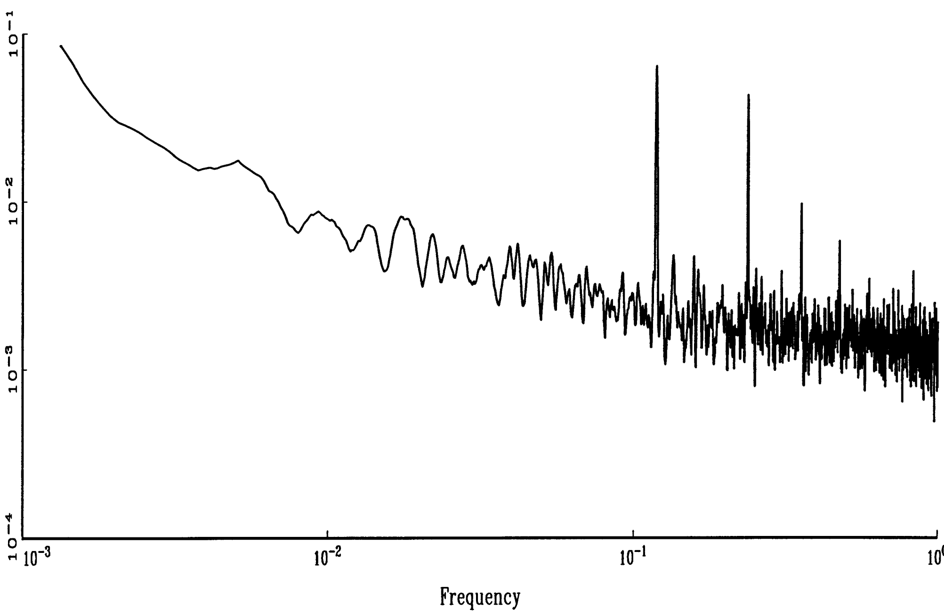

Fig. 5. The figure graphs the spectrum for the absolute 5-min Nikkei 225 returns. The estimate is based on the smoothed sample periodogram and a triangular kernel with a bandwidth of ten. The sample period for the returns extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

[15] Even more strikingly, the log-periodogram estimate based on the longer time series of 2308 daily Nikkei 225 absolute returns covering September 5, 1988 to December 30, 1997, equals 0.476 (0.093).

[16] The binomial expansion for $(1-L)^{0.429}$ is truncated at a lag length of 265, corresponding to 1 week. Also, the first order autocorrelation for $(1-L)^{0.429}|R_{t,n}|$ equals -0.218, and is omitted from Fig. 6.

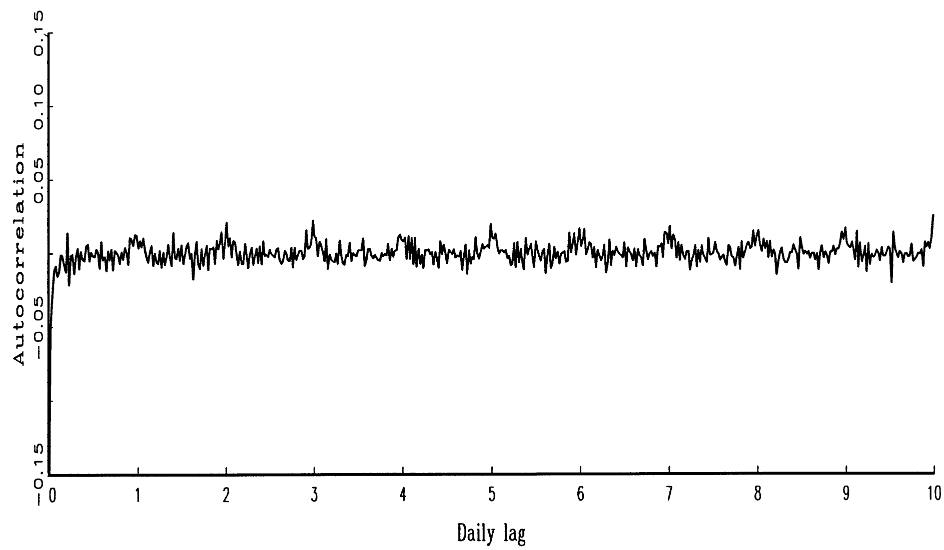

Fig. 6. The figure graphs the autocorrelogram for the fractionally differenced absolute 5-min Nikkei 225 returns, $(1-L)^{0.429}|R_{t,n}|$, out to a lag length 530, or 10 trading days. The sample period for the returns extends from January 2, 1994 through December 31, 1997, for a total of 52 152 observations.

around zero. The fractional differencing operator evidently eliminates the long-run interdaily dependencies.

It is noteworthy that the point estimates for Japanese high-frequency equity returns reported above are fully consistent with those for the US Dollar–Deutschemark exchange rate in Andersen and Bollerslev (1997a). The documentation of fractionally integrated volatility processes across daily financial return series for numerous distinct assets traded in a variety of different market structures, and the apparent stability of the degree of integration over intraday frequencies and subsamples suggest that this is a universal characteristic of speculative returns. It remains an open issue to identify the specific economic forces that may generate such long-run persistence patterns. At an abstract level, one possibility is that it may arise from the interaction of a large number of heterogeneous information arrival processes.[17]

5. Macroeconomic news announcement effects

Although it is evident from both theory, e.g. Ross (1989), and stylized facts that the intensity of information flow impacts the level of market volatility, it is generally difficult, if not impossible, to explicitly identify the information arrival

[17] Building on the mixture-of-distributions hypothesis advocated by Clark (1973), Epps and Epps (1976) and Tauchen and Pitts (1983), Andersen and Bollerslev (1997a) propose one such statistically oriented model with long-memory volatility.

processes underlying the systematic intraday patterns and dynamic dependencies documented above. However, scheduled macroeconomic news announcements provide a possible exception. There is already a sizeable literature on the effects of public information releases for US financial markets. For instance, Ederington and Lee (1993) find that regularly scheduled US macroeconomic announcements lead to significant time-of-day and day-of-week patterns in the volatility of US Treasury and foreign exchange futures. Similarly, Goodhart et al. (1993), DeGennaro and Shrieves (1997), and Payne (1996) report significant news effects in high-frequency foreign exchange spot rates. Also, Ito and Roley (1987) find that US money supply figures help to explain the movements in the Japanese Yen–US Dollar exchange rate. In contrast, Cutler et al. (1989) find little association between the largest daily price changes in the US stock market and readily identifiable economic news. Likewise, Mitchell and Mulherin (1994) show that the number of Dow Jones announcements is only weakly related to US stock market volume and volatility, while Berry and Howe (1994) report no significant relationship between US equity volatility and the total number of news releases by the Reuters News Service.

Most of these studies treat the three distinct volatility components — the intraday volatility pattern, the persistent interday volatility factor, and the announcement effects — in isolation. This is problematic, as the pronounced intraday periodicity may obscure the relationship between volatility and information releases. In order to avoid such potential pitfalls, we continue to rely on the decomposition and FFF modeling approach outlined in Section 3. In addition to the intraday patterns and the daily FIEGARCH factor, we also include a set of dummy variables to control for day-of-week effects and a systematic impact of the bi-weekly Japanese cabinet meetings.[18]

Our analysis of the announcement effects for the Nikkei 225 index is based on the last 2 years of our 4-year sample only; i.e. January 2, 1996 through December 31, 1997, for a total of 26 076 5-min return observations.[19] We analyze 19 recurring announcements.[20] These include figures on household spending, consumer sentiment, housing, industrial production and orders, vehicle production and sales, retail sales, unemployment, inflation and bankruptcies. In contrast to most US macroeco-

[18] The Japanese cabinet meets at 09:00 every Tuesday and Friday. Afterwards the Finance Minister holds a press conference. Our cabinet meeting dummies are set to unity between 09:45 and 10:45 on those days.

[19] Bloomberg Information Services have provided information on Japanese macroeconomic announcements since January 1996, but otherwise this information is not readily available. The Nikkei Telecom/Japan News & Retrieval Services provide the release times for the full 4-year sample from 1994 through 1997, but cover only a small fraction of the relevant news announcements.

[20] Several other announcements, including the current account, merchandise trade, GDP, and money supply figures, the whole-sale and corporate service price indices, and the Tankan (business survey) report, were always announced outside of the TSE trading hours. In addition, the timing of a number of announcements coincide. For instance, the unemployment rate and the job-to-applicant ratio are announced together, as are the housing starts and construction orders, as well as the nationwide consumer price index and the separate index for the Tokyo area.

nomic announcements, the Japanese news are usually not released on a punctual schedule, but varies both in terms of frequency and time-of-day. The actual number of releases for each category over the sample period is given in Table 1.

Given the limited number of occurrences of each type of news and the inherent noise in the return process, it is not feasible to accurately estimate separate coefficients for each 5-min interval following each news release. Instead, we impose

Table 1
Macroeconomic announcement effects[a]

Announcement (# observations)	FFF-coefficient (Robust t-stat.)	Instantaneous volatility (%)	Daily volatility (%)
All News (192)	0.39 (1.59)	10	0.7
Household Spending (23)	2.37 (3.68)	82	6.2
Vehicle exports (24)	1.45 (1.86)	44	2.0
Vehicle production (23)	1.45 (1.78)	44	2.0
Crude oil imports (13)	1.72 (1.65)	55	2.9
Housing starts/Construction orders (24)	1.16 (1.63)	40	2.8
Vehicle sales (23)	1.33 (1.41)	44	3.2
Consumer sentiment index (6)	2.77 (1.36)	102	7.5
Industrial production (21)	1.01 (1.13)	29	1.6
Tokyo new condominium sales (19)	0.64 (0.76)	18	0.8
Machinery Orders (22)	0.33 (0.48)	9	0.7
CPI-Tokyo and nationwide (17)	0.13 (0.10)	3	0.2
Unemployment (16)	0.10 (0.09)	3	0.2
Large scale retail sales (9)	−0.05 (−0.03)	−1	−0.1
Tokyo department store sales (21)	−0.23 (−0.24)	−6	−0.6
International securities investment (19)	−0.33 (−0.40)	−8	−0.5
Steel production (24)	−0.44 (−0.59)	−11	−0.9
Diffusion index (24)	−0.58 (−0.75)	−14	−2.1
Nationwide department store sales (20)	−1.29 (−1.16)	−28	−3.1
Bankruptcies (19)	−1.87 (−2.22)	−38	−2.5

[a] The table reports the estimated news announcement effects for the Nikkei 225 index return volatility based on the FFF regression in Eq. (3). The 5-min Nikkei 225 returns, $R_{t,n}$, covers the period from January 2, 1996 to December 31, 1997, for a total of 26 076 observations. The estimate for the daily volatility factor, σ_t, is based on a MA(1)-FIEGARCH(1,1) model and daily Nikkei 225 returns from September 5, 1988 through December 31, 1997. All of the FFF-regressions include day-of-the-week dummies, along with 09:45–10:45 h dummies for the Tuesday and Friday Japanese cabinet meetings. The response pattern following each of the announcements is approximated by a third-order polynomial restricted to reach zero at the end of the 1-h response horizon, or the 13th 5-min interval. The reported FFF-coefficients, λ_k, measure the extent to which the FFF-regressions load into this pattern. The instantaneous jump in the volatility gives the increase in the 5-min absolute return for the actual announcement interval, or $\exp(\lambda_k \cdot \gamma(0)/2) - 1$. The daily volatility measure gives the estimated impact on the cumulative absolute returns over the 1-h announcement horizon, $M(k) \equiv \Sigma_{i=0,\ldots,12}[\exp(\lambda_k \cdot \gamma(i)/2) - 1]$, relative to the sample median daily cumulative absolute 5-min returns of 4.00%.

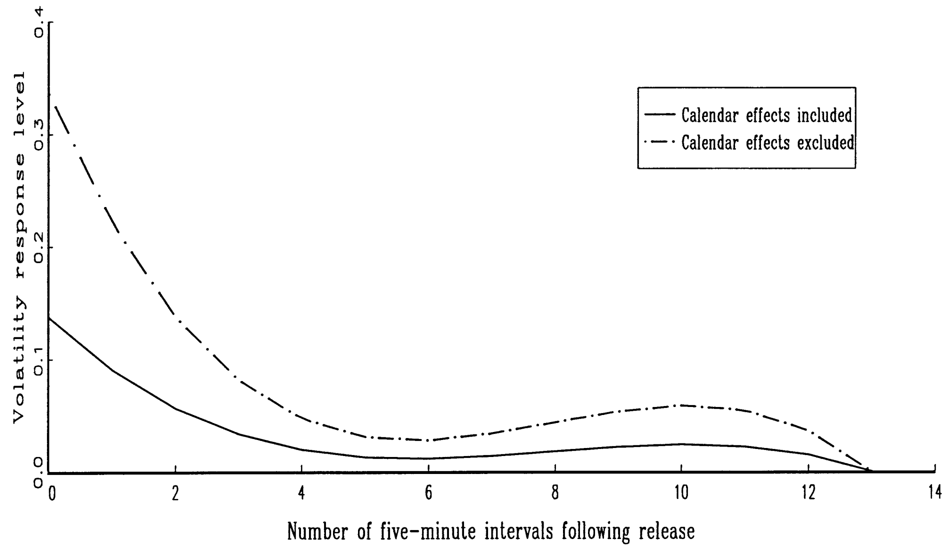

Fig. 7. The figure gives the estimated average response pattern for the Japanese macroeconomic news announcements. The estimates are based on the FFF-regression in Eq. (3) and the ten announcements listed in Table 2, along with a third-order polynomial restricted to reach zero at the 1-h horizon. The sample period extends from January 2, 1996 to December 31, 1997, for a total of 26 076 observations.

a reasonable decay structure based on the average volatility response pattern across all of the announcements and estimate the degree to which the specific announcement of type k 'loads onto' this pattern via a loading coefficient, λ_k. After some experimentation, the response horizon of this pattern was fixed at 1 h. Also, as in Andersen and Bollerslev (1998a), the shape of the average response pattern is well approximated by a third degree polynomial; say $\lambda_k \cdot \gamma(i)$, where $\gamma(i)$ dictates the response at lag $i = 0, 1, 2, \ldots, 12$. This estimated response pattern is depicted in Fig. 7.[21] According to Eq. (3), the instantaneous impact on the absolute 5-min returns is therefore $\exp(\lambda_k \cdot \gamma(0)/2) - 1$, while the response at the ith 5-min lag is $\exp(\lambda_k \cdot \gamma(i)/2) - 1$. We define the corresponding cumulative response measure by $M(k) \equiv \Sigma_{i=0,\ldots,12}[\exp(\lambda_k \cdot \gamma(i)/2) - 1]$. This measure reflects the impact over the entire 1-h response horizon expressed as a multiple of the typical volatility for the associated 5-min intervals.

Turning to the empirical results, the parameter estimate for the λ_k coefficient for all announcements combined equals 0.39 with a t-statistic of 1.59. Not only is the effect statistically insignificant, but the point estimate implies an instantaneous increase in the absolute 5-min return of only $\exp(\lambda_k \cdot \gamma(0)/2) - 1 = \exp(0.39 \cdot 0.507/2) - 1 = 0.104\%$. The associated cumulative response measure $M(k)$, judged relative to the median cumulative absolute 5-min returns, indicates an elevation of the

[21] The actual estimates for the third-order polynomial, when including calendar effects in the FFF-regression, is given by $\gamma(i) = 0.508 \cdot [1 - (i/13)^3] - 0.200 \cdot [1 - (i/13)^2] \cdot i + 0.027 \cdot [1 - (i/13)] \cdot i^2$, $i = 0, 1, 2, \ldots, 12$.

average volatility for trading days that contain a single macroeconomic announcement of only 0.7% relative to days without any news releases. Moreover, these estimates are virtually unaffected by excluding the dummies for the regular Japanese cabinet meetings, or by assuming the daily volatility factor, σ_t, to be constant.

Of course, the insignificant average announcement effects may mask the influence of important individual announcements. Table 1 reports estimates of separate λ_k coefficients for each type of announcement. Only the monthly report on household spending stands out as significant. The instantaneous jump in volatility equals $\exp(2.37 \cdot 0.507/2) - 1 = 1.82 - 1$, or 82%. The figures for household spending are typically announced around 14:00 h, and the average 5-min absolute return between 14:00 and 15:00 h equals 0.097%. The corresponding cumulative response estimate of $M(k) = 2.536$ thus implies an increase of about $2.536 \cdot 0.097\% = 0.246\%$, tantamount to an overall increase of just 0.246%/(53·0.075%), or 6.2%, for trading days on which the household spending figures are announced. Of the remaining 18 news releases, eleven result in an average increase in the volatility, although none of the estimates are statistically significant. Interestingly, the report on bankruptcies is followed by a significant 'decrease' in the volatility in the 1 h following the announcement.[22]

It is difficult to draw any general conclusion from these results. First, there are some important announcements that invariably are released outside of the TSE trading hours. This may reflect the expectation that these particular releases are likely to have a disruptive impact on the market, and hence purposely are released outside of the trading hours. Second, there are only a few announcements of each kind, so the statistical power is relatively low. Nonetheless, a tentative conclusion is that the releases with relatively large estimated coefficients speak to the state of consumer or business confidence through either sentiment indices, household purchases, production numbers or housing starts. In contrast, the releases on price development or unemployment are inconsequential. This may well reflect the perception that the weakness in the Japanese economy over this period was related to the lack of consumer spending, business confidence and activity in the housing markets. Conversely, the drop in volatility following the bankruptcy report may suggest that this figure is quite readily interpreted by the market participants, and consequently entail a resolution of uncertainty (although emphasizing one significant negative coefficient among 19 also may entail a degree of data snooping). Note also that the size of the λ_k coefficients does not uniquely rank the announcements in terms of their impact on the cumulative daily volatility, because the news are released at a different time-of-day. For example, the vehicle sales have a lower

[22] It is also noteworthy, that excluding the intradaily δ-coefficients from the FFF-regression, the estimates for λ_k, with t-statistics in parentheses, for the reports on housing starts and Tokyo department store sales are 1.83 (2.55) and 2.21 (2.23), respectively, both of which are now significant. Furthermore, the estimate for nationwide department store sales changes sign to 1.24 (1.13). These results illustrate the importance of accounting for the strong intradaily periodicities when assessing the quantitative and qualitative impact of the announcement effects.

coefficient than vehicle production, but a larger impact on overall volatility because the (percentage) impact of the former is initiated at 14:00 h when the market is more volatile compared to the typical 13:00 h release time for the latter figures.

These relatively weak macroeconomic announcement effects stand in sharp contrast to the results for the Deutschemark–US Dollar foreign exchange rate in Andersen and Bollerslev (1998a), who report that the US employment report, the 'king-of-kings' among announcements, induces an average instantaneous jump in exchange rate volatility of almost 600%, along with an increase in the cumulative absolute return for days containing a scheduled employment report of about 15%.[23] Meanwhile, our findings are fully consistent with the prior evidence of a weak linkage between public information and US stock market volatility documented by Cutler et al. (1989), Berry and Howe (1994) and Mitchell and Mulherin (1994), and among others.

6. Relative importance of the different volatility components

The intraday volatility pattern, the interday volatility factor, and the macroeconomic announcements all contribute to the volatility in the market. However, the relative importance of the different components is not obvious. This section provides a direct assessment of the joint and marginal predictive power of the components for forecasting the variation in the Nikkei 225 volatility.

The basic idea again builds on the decomposition in Eq. (1). In particular, given the model estimates of the three components discussed in the previous sections, a 1-day-ahead forecast for the absolute 5-min return is readily constructed from,

$$v_{t,n} = N^{-1/2} \cdot [\hat{\sigma}_t \cdot I_\sigma + \bar{\sigma}(1 - I_\sigma)] \cdot \exp\{[f_c(t,n) \cdot I_c + f_a(t,n) \cdot I_a + f_w(t,n)]/2\} \tag{5}$$

where $\hat{\sigma}_t$ refers to the 1-day-ahead volatility forecasts from the FIEGARCH model, and $f_c(t,n)$, $f_a(t,n)$ and $f_w(t,n)$ denote the estimated calendar, announcement, and day-of-week effects, respectively. The indicator variables I_σ, I_c, and I_a, signify whether the corresponding feature is accounted for in the particular forecast. For example, $(I_\sigma, I_c, I_a) = (0, 1, 1)$ corresponds to a model where the daily volatility is assumed to be constant, but the calendar and announcement effects are both included. For simplicity, the day-of-week effects, $f_w(t,n)$, are incorporated in all of the forecasts.

Our assessment of the relative importance of the volatility components is based on a direct in-sample comparison of the deviation between the realized and forecasted absolute returns. Of course, the relative importance of the different components depends on the forecast horizon. For 'instantaneous volatility' our

[23] Andersen and Bollerslev (1998a) also find that the day-of-the-week effects in the foreign exchange market become insignificant when the systematic influences of the US and German macroeconomic announcements are accounted for. Of the 192 Japanese announcements included in our sample, Monday through Friday account for 30, 52, 12, 31 and 67, respectively. However, this clustering does not explain the significant day-of-the-week effects in the FFF-regression for the Nikkei 225 volatility.

measure of forecast performance is simply given by the coefficient of multiple correlation, or R^2, from the regression of the realized 5-min absolute returns on the 1-day-ahead forecast, $v_{t,n}$, from Eq. (5). Similarly, the forecast performance for 'daily volatility' is measured by the R^2 from the regression of the daily sum of the realized absolute 5-min returns on the corresponding cumulated 5-min absolute return forecasts; i.e. $\Sigma_{n=1,\ldots,53}\, v_{t,n}$. Although these criteria do not correspond directly to the loss functions that are relevant in most practical applications involving the pricing of options or the implementation of dynamic portfolio and hedging strategies, the R^2s allow for a simple statistical model-free comparison.[24]

Table 2 summarizes the results for the eight possible model configurations as indicated by the triplet (I_σ, I_c, I_a). The first data column provides the explained variation for the 5-min absolute return forecasts. The explained variation for the full model incorporating all the components is 8.6%. The daily volatility factor and the calendar effects are the main features, explaining about 6.0 and 3.5% in

Table 2
Explained variation for different volatility components[a]

Model (I_σ, I_c, I_a)	Instantaneous volatility	Daily volatility
Full model (1, 1, 1)	0.086	0.450
No announcements (1, 1, 0)	0.086	0.449
No calendar effects (1, 0, 1)	0.059	0.448
No daily volatility (0, 1, 1)	0.035	0.009
Announcements only (0, 0, 1)	0.001	0.009
Calendar effects only (0, 1, 0)	0.035	0.010
Daily volatility only (1, 0, 0)	0.060	0.450
Day-of-the-week effects only (0, 0, 0)	0.001	0.010

[a] The table reports the explained variation for forecasting the Nikkei 225 index volatility. The 5-min Nikkei 225 returns, $R_{t,n}$, covers the period from January 2, 1996 to December 31, 1997, for a total of 26 076 observations. The daily volatility factor, σ_t, is based on a MA(1)-FIEGARCH(1,1) model and daily Nikkei 225 returns from September 5, 1988 through December 31, 1997. The 1-day-ahead forecast for the 5-min absolute returns are obtained as: $v_{t,n} = N^{-1/2}\cdot[\sigma_t\cdot I_\sigma + \bar{\sigma}(1 - I_\sigma)]\cdot\exp\{[f_c(t,n)\cdot I_c + f_a(t,n)\cdot I_a + f_w(t,n)]/2\}$; where $f_c(t,n)$ and $f_a(t,n)$ refer to the estimated calendar and announcement effects, respectively. The day-of-the-week effects, $f_w(t,n)$, is included in all of the forecasts. The indicator variables I_σ, I_c, and I_a, signify whether the corresponding feature is accounted for in the particular model forecast. The forecasting performance for the instantaneous volatility is assessed by the coefficient of multiple correlation, or R^2, from the regression of the realized 5-min absolute Nikkei 225 returns on $v_{t,n}$. The forecasting performance for the daily volatility is measured by the R^2 from the regression of the daily sums of the realized absolute 5-min returns on the corresponding cumulated 5-min absolute return forecasts, $\Sigma_{n=1,\ldots,53}\, v_{t,n}$.

[24] Note that, since the 1-day-ahead absolute return forecasts entail a nonlinear transformation of the daily volatility factor and the pertinent FFF components, it is possible that the explanatory power may actually increase (slightly) by deleting particular components.

isolation, respectively. The negligible impact of the macroeconomic announcements, accounting for only 0.1% of the variation in intraday volatility, is consistent with the results in the previous section.

Turning to the results for the daily cumulative absolute returns, the full model explains an impressive 45.0% of the variance of the daily Nikkei 225 volatility. The figure is virtually unchanged when either the announcement and/or calendar effects are excluded.[25] In contrast, the R^2 drops to a low 0.9% when the daily FIEGARCH volatility factor is omitted. The marginal impact of the news announcements is, of course, not surprising. More importantly, the results in Table 2 suggests that calendar effects associated with the intraday pattern and the day-of-the-week dummies are effectively annihilated when forecasting over longer-run interdaily horizons.

7. Concluding remarks

This paper provides a detailed characterization of the volatility in the Japanese stock market based on a long sample of high-frequency 5-min Nikkei 225 index returns. Our major findings are as follows. The intraday volatility is dominated by two separate U-shaped patterns; one for the morning and one in the afternoon. This heightened volatility around the open and close of the two separate trading sessions on the TSE is broadly consistent with the predictions from theoretical market microstructure models based on the strategic interaction of asymmetrically informed agents. Meanwhile, there is scant evidence for any systematic increase in the volatility in connection with the release of regular Japanese macroeconomic news reports. This conclusion, however, is critically dependent upon appropriately adjusting for the double U-shaped pattern in the intradaily volatility. Our results further show that the high-frequency data contain very useful information about the salient long-lived interdaily volatility dependencies. This long-memory component in turn explains almost 50% of the day-to-day variation in the Nikkei 225 volatility.

The decomposition and estimation procedures employed in the present analysis can readily be applied to other high-frequency return series. It would also be interesting to use the approach for disentangling the complex link between the short- and long-run components when analyzing the dynamic dependencies and announcement effects in intraday bid-ask spreads and trading volume. Future research along these lines holds the promise of delivering a deeper understanding of the role of different information sources and the functioning of alternative market mechanisms.

[25] This relatively high R^2 also highlight the fallacy of conventional volatility forecast evaluation criteria based on the realized daily absolute, or squared, returns, which typically suggest very poor interdaily forecast performance; see Andersen and Bollerslev (1998b) for further discussion along these lines.

Acknowledgements

Andersen's and Bollerslev's research was supported by a grant from the NSF to the NBER. J. Cai gratefully acknowledges the financial support from RGC Competitive Earmarked Research Grants 1994–1996 and 1996–1998.

References

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

Amihud, Y., Mendelson, H., 1991. Volatility, efficiency, and trading: evidence from the Japanese stock market. J. Finance 46, 1765–1789.

Andersen, T.G., Bollerslev, T., 1997a. Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns. J. Finance 52, 975–1005.

Andersen, T.G., Bollerslev, T., 1997b. Intraday periodicity and volatility persistence in financial markets. J. Empirical Finance 4, 115–158.

Andersen, T.G., Bollerslev, T., 1998a. Deutschemark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer-run dependencies. J. Finance 53, 219–265.

Andersen, T.G., Bollerslev, T., 1998b. Answering the skeptics: yes standard volatility models do provide accurate forecasts. Int. Econ. Rev. 39, 885–906.

Andersen, T.G., Bollerslev, T., Lange, S., 2000. Forecasting financial market volatility: sample frequency vis-à-vis forecast horizon. J. Empirical Finance (in press).

Barclay, M., Litzenberger, R., Warner, J., 1990. Private information, trading volume, and stock-return variances. Rev. Financial Stud. 3, 233–253.

Beran, J., 1994. Statistics for Long-Memory Processes. Chapman and Hall, New York.

Berry, T.D., Howe, K.M., 1994. Public information arrival. J. Finance 49, 1331–1346.

Bollerslev, T., Mikkelsen, H.O., 1996. Modeling and pricing long-memory in stock market volatility. J. Econometrics 73, 151–184.

Bollerslev, T., Wright, J.H., 2000. Estimating long-memory volatility dependencies: the role of high-frequency data. J. Econometrics, forthcoming.

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

Breidt, F.J., Crato, N., deLima, P., 1998. On the detection and estimation of long-memory in stochastic volatility. J. Econometrics 83, 325–348.

Brock, W.A., Kleidon, A.W., 1992. Periodic market closure and trading volume: a model of intraday bids and asks. J. Econ. Dynamics Control 16, 451–489.

Chan, K., Chan, K.C., Karolyi, G.A., 1991. Intraday volatility in the stock index and stock index futures markets. Rev. Financial Stud. 4, 1161–1187.

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

Cutler, D.M., Poterba, J.M., Summers, L.H., 1989. What moves stock prices? J. Portfolio Manage. 15, 4–12.

Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V., 1993. A geographical model for the daily and weekly seasoned volatility in the FX market. J. Int. Money Finance 12, 413–438.

DeGennaro, R.P., Shrieves, R.E., 1997. Public information releases, private information arrival, and volatility in the foreign exchange market. J. Empirical Finance 4, 295–316.

Ding, Z., Granger, C.W.J., 1996. Modeling volatility persistence of speculative returns: a new approach. J. Econometrics 73, 185–215.

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

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

Epps, T.W., Epps, M.L., 1976. The stochastic dependence of security price changes and transaction volumes: implications for the mixture-of-distribution hypothesis. Econometrica 44, 305–321.

Fleming, M.J., Remolona, E.M., 1999. Price formation and liquidity in the US treasury market: the response to public information. J. Finance 54, 1901–1915.

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

French, K.R., Roll, R., 1986. Stock return variances: the arrival of information and the reaction of traders. J. Financial Econ. 17, 5–26.

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

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

Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long-memory time series models. J. Time Series Anal. 4, 221–238.

Ghysels, E., Harvey, A., Renault, E., 1996. Stochastic volatility. In: Maddala, G.S., Rao, C.R. (Eds.), Handbook of Statistics, vol. 14. North-Holland, Amsterdam, pp. 119–191.

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

Hamao, Y., Hasbrouck, J., 1995. Securities trading in the absence of dealers: trades and quotes on the Tokyo stock exchange. Rev. Financial Stud. 8, 849–878.

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

Hurvich, C.M., Deo, R., Brodsky, J., 1998. The mean squared error of Geweke and Porter Hudak's estimator of the memory parameter of a long-memory time series. J. Time Series Anal. 19, 19–46.

Ito, T., Lin, W.L., 1992. Lunch break and intraday volatility of stock returns. Econ. Lett. 39, 85–90.

Ito, T., Roley, V.V., 1987. News from US and Japan: which moves the yen/dollar exchange rate? J. Monetary Econ. 19, 255–277.

Lamoureux, C.G., Lastrapes, W.D., 1990. Persistence in variance, structural change and the GARCH model. J. Business and Economic Statistics 8, 225–234.

Lee, C.I., Mathur, I., 1999. The influence of information arrival on market microstructure: evidence from three related markets. Financial Rev. 34, 1–26.

Lehmann, B.N., Modest, D.M., 1994. Trading and liquidity on the Tokyo stock exchange: a bird's eye view. J. Finance 49, 951–984.

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

Mitchell, M.L., Mulherin, J.H., 1994. The impact of public information on the stock market. J. Finance 49, 923–950.

Muranaga, J., 1998. Dynamics of market liquidity of Japanese stocks: an analysis of tick-by-tick data of the Tokyo stock exchange. Working paper, Institute for Monetary and Economics Studies, Bank of Japan, Tokyo.

Payne, R., 1996. Announcement effects and seasonality in the intraday foreign exchange market. Working paper, London School of Economics, United Kingdom.

Robinson, P.M., 1994a. Semiparametric analysis of long-memory time series. Ann. Statistics 22, 515–539.

Robinson, P.M., 1994b. Time series with strong dependence. In: Sims, C. (Ed.), Advances in Econometrics: Sixth World Congress. Cambridge University Press, Cambridge, pp. 47–95.

Robinson, P.M., 1995. Log-periodogram regression of time series with long-range dependence. Ann. Statistics 23, 1048–1072.

Ross, S.A., 1989. Information and volatility: the no-arbitrage martingale approach to timing and resolution irrelevancy. J. Finance 44, 1–18.

Slezak, S., 1994. A theory of the dynamics of security returns around market closures. J. Finance 49, 1163–1211.

Spiegel, M., Subrahmanyan, A., 1995. On intraday risk premia. J. Finance 50, 319–339.

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

Wood, R.A., McInish, T.H., Ord, J.K., 1985. An investigation of transactions data for NYSE stocks. J. Finance 40, 723–739.