NBER WORKING PAPER SERIES

TESTING FOR MARKET MICROSTRUCTURE EFFECTS IN INTRADAY VOLATILITY: A REASSESSMENT OF THE TOKYO FX EXPERIMENT

Torben G. Anderson
Tim Bollerslev
Ashish Das

Working Paper 6666
http://www.nber.org/papers/w6666

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1998

This work was supported by a grant from the National Science Foundation to the NBER. We are also grateful to Olsen and Associates for making the intraday exchange rate quotations available. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Testing for Market Microstructure Effects in
Intraday Volatility: A Reassessment of the
Tokyo FX Experiment
Torben G. Anderson, Tim Bollerslev,
and Ashish Das
NBER Working Paper No. 6666
July 1998
JEL No. C12, C22, F33, G14, G15

ABSTRACT

This paper develops mew robust inference procedures for analyzing the intraday return volatility patterns that constitute a focal point of much market microstructure theory. Our empirical analysis is motivated by the recent lifting of trading restrictions in the interbank foreign exchange (FX) market for Japanese banks during the Tokyo lunch period. Ito, Lyons, and Melvin (1998) (ILM) argue that this deregulation resulted in a highly significant shift in the volatility pattern across the entire Japanese trading day, indicating that private information is an important component of the price formation process in the FX market. In contrast, our robust analysis finds no evidence for any discernible change in the pattern outside of the Tokyo lunch period. Moreover, we document that the standard variance-ratio methodology, employed by ILM, provides very misleading inference in this high-frequency data context.

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

Tim Bollerslev
Department of Economics
Duke University
Durham, NC 27708
and NBER
bollerslev@virginia.edu

Ashish Das
Department of Finance
J.L. Kellogg Gradaute School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
adas@nwu.edu

A number of important financial market variables, including bid-ask spreads and return volatility, display pronounced intraday patterns.[1] For equity markets with daily openings and closures they typically follow a distorted U-shape over the trading day, whereas markets with round-the-clock trading in partially overlapping regional segments, such as the foreign exchange (FX) interbank market, produce more complex patterns, seemingly related to the daily ebb and flow of activity across the global financial centers. These striking regularities have inspired a large theoretical literature, seeking to explain how such patterns may arise from the interaction of distinct customer groups, trading professionals, and market makers.[2] Issues related to the notion of private (asymmetric) versus public information, the risks and inventory costs borne by market makers or dealers, and the regulatory and institutional arrangements surrounding trade in the specific security all play an important role in these theories.

Meanwhile, empirical work on the relative importance of these factors must confront several obstacles. First, a number of the key theoretical concepts are inherently unobservable, such as the prevalence of private or asymmetric information in the market or the strength of, say, discretionary liquidity trading. Second, market microstructure theories are typically not designed to provide quantitative predictions, but merely a qualitative characterization of the pattern that is likely to arise in some market variables. Third, meaningful comparisons across different market structures are complicated by the fact that these generally differ along a number of dimensions simultaneously. Consequently, a number of empirical investigations have resorted to event type studies, where the impact of exogenous or predetermined events may be gauged within a given market setting, thus providing an opportunity to discriminate among theories that offer different predictions regarding the impact of the specific event. Examples of such event studies include the analysis of exogenous modifications in exchange trading hours to test for changes in overall return volatility, and the short-run impact of firm-specific news releases or regularly scheduled macroeconomic announcements.[3]

---

[1] See, e.g., Wood et al. (1985), Harris (1986), Dacorogna et al. (1993), Hsieh and Kleidon (1996), and Andersen and Bollerslev (1998a).

[2] Early contributors include Admati and Pfleiderer (1988) and Foster and Viswanathan (1990, 1993); for a recent survey of this literature see O'Hara (1995).

[3] The event study methodology is largely due to Ball and Brown (1968) and Fama et al. (1969), but the applications to market microstructure issues is much more recent. For example, Barclay et al. (1990) follow French and Roll (1986) in exploring the discrepancy in return volatility over closed versus open trading periods, focusing on the phasing out of Saturday trading on the Tokyo Stock Exchange. Likewise, announcement effects were among the first to be studied via event study techniques, but the use of high-frequency data is fairly new, see, e.g., the discussions in Ederington and Lee (1993) and Andersen and Bollerslev (1998a).

The recent analysis of the interbank FX market by Ito, Lyons and Melvin (1998) (henceforth ILM) provides an intriguing extension of this event study approach to a high-frequency data setting.[4] Before December 22, 1994, banks operating from Japan were not allowed to trade outside of the hours 9:00-12:00 and 1:30-3:30 local time, with 12:00-1:30 constituting the Tokyo lunch period. ILM argue that if private information is important, not only should the yen-dollar (¥-$) return volatility increase over the Tokyo lunch period after the trading restrictions were lifted, but the entire volatility pattern over the Japanese trading segment should shift in response to the endogenous changes in the trading activities of privately informed agents. As such, the ILM study transcends prior empirical work by specifying a set of joint hypotheses that are expressed, not simply as a change in the level of return volatility at a given point in time or over the full trading day, but rather as a string of distinct qualitative changes in the pattern across the trading day. Having more formally specified these hypotheses, ILM go on to show that the intraday pattern does indeed change in all of the predicted directions, thus lending support to the notion that private information is a critical ingredient to the price formation process in the FX market -- an important finding with a range of interesting implications.[5]

The ILM study relies on fairly standard econometric methodology, and as such the problem of inference concerning high-frequency volatility patterns may appear trivial. Almost by definition, the number of intraday return observations is enormous. Standard asymptotic theory therefore suggests that the power of most statistical tests would be very high. Indeed, ILM conclude that virtually all of their null hypotheses of no change in the shape of the intraday volatility pattern may be rejected at extreme significance levels. This is deceptive. In addition to the pronounced intraday volatility pattern emphasized by ILM, high-frequency returns also exhibit a highly persistent conditionally heteroskedastic component, along with discrete information arrival effects associated with the release of public news. In combination, these factors generate strong serial correlation and jump like behavior in the intraday absolute, or squared, return series, to the point of invalidating standard statistical procedures.

Recognizing these complications, the present paper develops a new set of tools for inference

---

[4] The ILM study also prompted a writeup in the May 31, 1997 Economist entitled "Forex with Rice."

[5] For instance, this suggests that traditional models of exchange rate determination, based solely on readily observable macroeconomic variables, could usefully be extended to allow for asymmetrically informed optimizing agents.

in the high-frequency data setting that are valid under quite general assumptions. In contrast to the findings reported by ILM, these new robust procedures indicate that, apart from an increase in the volatility over the Tokyo lunch period, the intraday return volatility pattern in the ¥-$ FX market has remained remarkably stable following the lifting of the trading restriction.[6] While these findings do not disprove the relevance of private information in the FX market per se -- it may simply be the case that the lifting of the restrictions has no discernable impact on the optimal trading strategy of informed traders -- the change in the intraday pattern outside the lunch period in response to the "Tokyo Experiment" apparently has little to contribute to this general debate.

The rest of the paper is structured as follows. Section I takes a first look at the intraday ¥-$ volatility pattern in question. We find that the patterns alluded to in ILM are very hard to pin down, especially when the high-frequency return series only span one month before and after the change in regime. Section II explores the properties of the standard variance ratio procedures employed by ILM. Based on simulations and a bootstrap involving an eight-year sample of high-frequency ¥-$ returns, the evidence is unambiguous: the ILM sample is effectively very small, and standard inference procedures can be grossly misleading in the context of high-frequency data. Section III reviews the flexible and parsimonious Fourier form (FFF) regression approach for estimating intraday volatility patterns originally proposed by Andersen and Bollerslev (1997a, 1998a). Building on this approach, the section develops new inference procedures that may be used more generally in testing for changes in the patterns expressed in terms of functionals of the underlying estimated FFF parameters. Section IV applies these tools to the ILM hypotheses. Our findings support the notion of enhanced volatility during the Tokyo lunch period following the lifting of the trade restrictions, although the true statistical significance is much lower than stated by ILM. However, based on our new robust procedures and a much larger sample of high-frequency returns covering two years before and after the deregulation, we find no evidence of any significant changes in the volatility pattern outside of the Tokyo lunch period. Section V concludes.

---

[6] This is also consistent with the recent findings by Wei and Kim (1997), who argue that FX positions by large market participants do not predict future exchange rate movements, and thus do not indicate the presence of superior private information. On the other hand, the evidence in Peiers (1997) does suggest that Deutschebank may have access to relevant private order flow information immediately preceding central bank interventions by the German Bundesbank.

I. A Preliminary Look at the Volatility Pattern around the Tokyo Experiment

This section presents a preliminary and informal graphical investigation of the shift in the intraday ¥-$ return volatility pattern over the Japanese trading segment around the lifting of the trading restrictions on December 22, 1994. We initially follow the ILM strategy of constructing the intraday patterns from samples of high-frequency returns collected over equally sized windows spanning respectively twenty and sixty weekdays before December 22, 1994, and after January 4, 1995. However, whereas the before-versus-after calculations reported in ILM rely on one-minute ¥-$ returns, our analysis is based on linearly interpolated five-minute returns; i.e., $R_{t,n}$, where the subscripts refer to the n'th five-minute time interval on day t.[7] This is done to mitigate the effects of non-synchronous quotations and spuriously induced autocorrelation. Andersen and Bollerslev (1997a, 1998a) argue in a similar context that, due to the scarcity of quotes and the practice of spread positioning (shading) by dealers in the FX market, the use of standard discrete-time statistical techniques with returns calculated over intervals covering less than five minutes is problematic.

Turning to the results, the average absolute five-minute returns over the Japanese trading segment, defined as 9:00-3:30 Tokyo time, for the twenty and sixty days before and after the deregulation are displayed in Panels A and B of Figure 1. The large degree of noise is evident from the jagged nature of the graphs. Nonetheless, consistent with the qualitative features emphasized by ILM, the overall return volatility does appear to increase following the regulatory change. Furthermore, the discrepancy between the before and after graphs is particularly striking for the lunch period, where the volatility roughly doubles. Meanwhile, it is clear that any other conclusions concerning the overall change in the shape based on visual inspection of the average patterns in Figures 1.A and 1.B will be subject to a considerable amount of uncertainty. In addition to erratic peaks induced by outliers, potential time-variation in volatility at the daily frequency will render the level of the pattern hard to identify from samples spanning only twenty or sixty days.

The simplest possible remedy for such small sample estimation error uncertainty is to prolong the daily windows around the event. Panel C of Figure 1 displays the identical volatility patterns, based on absolute five-minute ¥-$ returns over the two-year samples immediately before and after the regulatory change. The reduced impact of outliers is evident. Moreover, visual inspection of

[7] For a more detailed description of the method of data capture and return calculations we refer to Andersen and Bollerslev (1997a, 1998a) and Dacorogna et al. (1993).

Figure 1.C indicates that, aside from a heightened volatility from 12:00-1:30, the pre- and post-event average intraday volatility patterns appear almost identical.[8] The following section addresses the apparent contradiction between the graphical displays in Figure 1 and the seemingly strong statistical evidence put forth by ILM.

II. On the Distribution of Intraday Variance Ratio Tests

Variance ratio statistics have routinely been employed in the empirical market microstructure literature as a way to assess the rate of information flow through time and across market structures.[9] This is also the methodology adopted by ILM.

To illustrate, consider the first hypothesis investigated by ILM, which also happens to produce some of their strongest statistical rejections. If private information plays no role in the price formation process and the ¥-$ returns are serially uncorrelated,[10] then the return volatility during the Tokyo lunch period from 12:00-1:30 should be identical in the before and after regimes. That is, the null hypothesis of no private information may be stated in terms of the variance ratio,

$$V_L^o / V_L^c \;=\; 1, \tag{1}$$

where V_L^c and V_L^o denote the lunch period variability for the closed and open regimes, respectively. If the returns are i.i.d. normally distributed and the null hypothesis is true, the corresponding sample variance ratio statistic, $\hat{V}_L^o / \hat{V}_L^c$, should be the realization of an $F_{no-1,nc-1}$ distribution, where *no* and *nc* refer to the number of returns used in the computation of each of the sample variances. For large *no* and *nc* this $F_{no-1,nc-1}$ distribution is well approximated by a normal distribution with a unit mean

---

[8] It is noteworthy, that the 2-year patterns in Figure 1.C are almost uniformly above the 20-day and 60-day before patterns in Figures 1.A and 1.B which form the basis for the inference in ILM. This suggests that the apparent discrepancy in the level of volatility over the short windows before and after the event may be due to an unusually low level of overall volatility just prior to the event, rather than an elevation in the level of volatility following the deregulation.

[9] For instance, Amihud and Mendelson (1987, 1991) and Stoll and Whaley (1990) examine the performance of different trading mechanisms on the basis of variance ratios of open-to-open versus close-to-close returns, while Oldfield and Rogalski (1980), French and Roll (1986), Harvey and Huang (1991) and Jones, Kaul and Lipson (1994) analyze the rate of information flow during trading and non-trading periods by comparing variance ratios for open-to-close versus close-to-open returns. Jones and Kaul (1994) survey the literature, and also argue that commonly reported variance ratio statistics for individual stock returns may be severely biased due to small sample sizes, measurement errors, and the reliance on cross-sectional means as a convenient summary measure.

[10] Variance ratio statistics have also been used to test for mean reversion in returns by comparing short- and long-horizon return sample variances spanning the identical time period; see e.g., Poterba and Summers (1988), Lo and MacKinlay (1988), Richardson and Smith (1991), and Smith (1994).

and a variance of $2\cdot nc^2\cdot(no+nc\text{-}2)\cdot[no\cdot(nc\text{-}2)^2\cdot(nc\text{-}4)]^{-1}$; see, e.g., Johnson and Kotz (1970).

Based on linearly interpolated one-minute ¥-$ returns, ILM report a 20-day before-versus-after sample variance ratio of 2.27. With a sixty day window the sample variance ratio equals 2.13. Comparing these statistics to the fractiles in the $F_{1799,1799}$ and $F_{5399,5399}$ distributions, both are overwhelmingly significant. ILM further note that "... the usual small sample bias that plagues equity market studies is irrelevant here: the number of one-minute observations in our sample period of 20 days is 1,799."[11] Meanwhile, for the five-minute returns the 20-day and 60-day variance ratios equal 2.43 and 3.49, both of which exceed their one-minute counterparts, and both are in the extreme right tails of the corresponding $F_{359,359}$ and $F_{1079,1079}$ distributions.[12] Consequently, our five-minute return variance ratio statistics appear to further solidify the strong statistical evidence in ILM against the null hypothesis of identical lunch period variances in the before and after regimes.

Unfortunately, the critical values underlying this inference are predicated on the returns being normally distributed with a constant variance, whereas high-frequency speculative returns exhibit temporally dependent volatility clustering. Several competing volatility models have been proposed for characterizing this phenomenon at the daily and lower frequencies; see, e.g., Bollerslev, Engle and Nelson (1994) and the references therein. To illustrate what profound effects such volatility clustering may have on the distribution of standard variance ratio statistics, suppose that the conditional variance for the five-minute returns is determined by a standard GARCH(1,1) model,

$$\sigma^2_{t,n} \;=\; \omega \;+\; \sigma^2_{t,n-1}\cdot(\alpha\cdot z^2_{t,n-1} + \beta) \tag{2}$$

where $R_{t,n} \equiv \sigma_{t,n}\cdot z_{t,n}$ and $z_{t,n}$ is i.i.d. N(0,1).[13] Figures 2.A and 2.B display the $F_{359,359}$ and $F_{1079,1079}$ distributions for the 20-day and 60-day $\hat{V}^o_L/\hat{V}^c_L$ variance ratio statistics under the ideal assumption of

---

[11] With only 15.3 quotes available on average for each ninety-minute lunch period interval in the twenty day pre-event sample, and 39.8 quotes in the post-event sample, the one-minute returns underlying the 20-days variance ratio statistics in ILM entail a considerable amount of interpolation. Consequently, the effective sample size for the variance computation is much less than 20·90 = 1,800. Of course, the observed discrepancy in quote intensity across the two regimes may itself be informative.

[12] This discrepancy may be explained by a downward bias in the one-minute statistics due to linear interpolation of the non-synchronous quotes.

[13] While the GARCH(1,1) model provides a good description of the volatility clustering at the interdaily level, the model clearly misses the complex component structure and systematic intraday volatility patters documented by Andersen and Bollerslev (1997a, 1998a) and Müller et al. (1997) among others. Nonetheless, as a first order approximation, it provides a useful illustration of the effects of time varying volatility on standard test statistics involving high-frequency returns.

i.i.d. normally distributed returns, i.e., $\alpha = \beta = 0$, along with the distribution that obtains when the five-minute ω, α, and β parameters are fixed at values implied by typical daily GARCH(1,1) estimates.[14,15] It is evident that even though the *unconditional* variance is constant, the F-distributions provide extremely poor approximations to the true sampling distributions for the variance ratio statistics in the presence of time-varying *conditional* volatility.[16] In the empirically relevant GARCH(1,1) distribution, the p-value for the 20-day before-versus-after variance ratio statistic of 2.43 equals 0.082. Similarly, the p-value for the 60-day variance ratio statistic of 3.49 equals 0.025. While the latter statistic would be significant at the usual five percent critical level, the true significance level is markedly lower than suggested by the $F_{1079,1079}$ distribution.

This lack of robustness is confirmed by a simple empirical bootstrap based on the rolling variance ratio statistics from December 2, 1986, up until the twenty or sixty days immediately preceding the December 22, 1994, deregulation. Excluding weekends, holidays, and days with gaps in the data transmission, a total of 1,918 and 1,878 such daily variance ratio statistics are available for the 20-day and 60-day horizons, respectively. The relatively close correspondence between these unconditional distributions, given by the solid lines in Figures 2.A and 2.B, and the GARCH(1,1) distributions for $\hat{V}_L^o/\hat{V}_L^c$ is striking. Evaluating the actual 20-day $\hat{V}_L^o/\hat{V}_L^c$ test statistic in the empirical bootstrap distribution implies a p-value of 0.205 for the null hypothesis of no structural change on December 22, 1994, vis-a-vis the behavior of the Tokyo lunch period volatility in the eight years preceding the lifting of the trading restrictions. The 60-day variance ratio statistic of 3.49 has a corresponding p-value of 0.043.[17] Thus, counter to the inference based on the traditional F-

---

[14] With a sample of daily ¥-$ returns from October 1, 1987 through September 30, 1992, Andersen and Bollerslev (1998b) reports $\hat{\omega} = 0.026$, $\hat{\alpha} = 0.104$, and $\hat{\beta} = 0.844$, corresponding to an unconditional annualized volatility of approximately eleven percent. From the temporal aggregation results in Drost and Nijman (1993), the implied parameters at the five-minute frequency are $\omega = 0.308 \cdot 10^{-6}$, $\alpha = 0.00924$, and $\beta = 0.991$.

[15] The GARCH(1,1) distributions are based on a total of 50,000 simulations. The normal random variables, $z_{t,n}$, were generated by the RNDNS subroutine in the GAUSS computer language.

[16] The current investigation of the impact of conditional heteroskedasticity on non-overlapping return variance ratio statistics is conceptually distinct from, and have no direct implications for, the short-horizon versus long-horizon return variance ratio statistics employed in testing for serially correlated returns, as investigated by Lo and MacKinlay (1989).

[17] It is noteworthy that the mean of the rolling 20-day and 60-day variance ratio statistics equals 1.71 and 1.21, respectively. Taken at face value, both of these statistics suggest a systematic increase in the Tokyo lunch period volatility from 1986 through 1994. This would be consistent with the migration of private order flow to the Singapore and Hong Kong markets which, as noted by ILM, may have been the primary motivation behind the lifting of the trading restrictions in an attempt by the Tokyo foreign exchange dealers to stem the tide; see also the survey of Far Eastern foreign exchange dealers in Cheung and Wong (1997). This

distributions, the true statistical evidence for a significant increase in the Tokyo lunch period volatility following the regulatory change, that can be gleaned from simple non-parametric variance ratio statistics based on event windows spanning only a few months, is at best tentative once we account for the noisy, non-normally distributed character of high-frequency returns.

The results in this section highlight the need for more powerful and flexible parametric techniques that could be employed more generally in empirical market microstructure studies. The next section develops such new robust statistical procedures.

III. Intraday Volatility Patterns: Estimation and Testing

The preceding sections establish that the main difficulties surrounding estimation of intraday volatility patterns are due to the relatively large outliers and the strong serial correlation of volatility in high-frequency return series. The former feature renders sample volatility measures for short intraday return intervals very noisy. The pronounced serial correlation in volatility, on the other hand, complicates the estimation of the overall volatility level. The flexible Fourier form regression proposed by Andersen and Bollerslev (1997a, 1998a) provides a simple procedure for estimation of the intraday volatility pattern that accommodates these inherent features of high-frequency return series. This section first summarizes the approach and next develops a general and robust inference procedure for hypothesis testing in this setting.

A. Flexible Fourier Form Regression

A few stylized empirical observations guide the flexible Fourier form (FFF) regression approach.[18] First, the systematic intraday volatility movements dwarf any predictable intraday return components associated with variation in the conditional mean. Second, the impact of outliers may be controlled via a robustifying log-transformation. Third, the strong serial correlation in volatility at the daily level does not appear to distort the overall intraday shape. Fourth, the shape of the intraday pattern is remarkably stable over time, enabling the use of relatively long event windows. Fifth, the intraday volatility pattern evolves quite smoothly over the trading day, and

reasoning is somewhat flawed, however. While the mean of $nc\cdot(nc-2)^{-1}$ for the variance ratio statistics that obtains under the ideal assumption of normally distributed returns is close to unity for large values of nc, the 60-day $\hat{V}_L^c/\hat{V}_L^o$ statistic for the more realistic five-minute GARCH(1,1) model equals 1.20 due to a Jensen's inequality effect.

[18] See Andersen and Bollerslev (1997a, 1998a, 1998c) for an in-depth discussion of these characteristics.

sharp discontinuities that may exist generally arise from predictable changes in the trading environment such as regulated market openings and closures, or the release of significant economic news during a specific time interval.

The first observation above suggests defining intraday return innovations by simply subtracting the associated sample mean. For convenience, we continue to denote this, now demeaned, return series as $R_{t,n}$, with the subscripts referring to the n'th intraday return interval on day t. The second through fourth observation suggest the following representation for this intraday return process,

$$R_{t,n} \;=\; \sigma_{t,n} \cdot s_{t,n} \cdot z_{t,n} \,, \tag{3}$$

where $z_{t,n}$ denotes an i.i.d. zero mean, unit variance error term, $\sigma_{t,n}$ signifies the influence of the overall level of volatility for day t, and $s_{t,n}$ represents the components associated with the intraday pattern as well as any other predictable factor affecting the expected return volatility over the n'th time interval. Note also that the third observation justifies treating all terms on the right hand side of equation (3) as statistically independent. Squaring the returns, invoking the robustifying log-transformation, and rearranging terms results in the following decomposition,

$$2\,log\,|R_{t,n}| \,-\, log\,\sigma^2_{t,n} \;=\; c \;+\; 2\,log\,s_{t,n} \;+\; u_{t,n} \,, \tag{4}$$

where $c \,=\, E[log\; z^2_{t,n}]$ and $u_{t,n} \;=\; log\; z^2_{t,n} \,-\, E[log\; z^2_{t,n}]$. Equation (4) may be viewed as a regression relating the deviation between the intraday squared returns and a daily volatility factor to the single explanatory variable, $s_{t,n}$. Hence, the role of $s_{t,n}$ is to capture the systematic intraday volatility movements that are unrelated to the daily ARCH effects. This is exactly the type of intraday volatility component that market microstructure theories tend to focus on.

Of course, the regression suggested by equation (4) can only be implemented if we obtain an observable proxy for the stochastic volatility component, $\sigma_{t,n}$, and a specific parameterization of $s_{t,n}$ in terms of measurable variables. The first issue is readily addressed. Given a sample of daily returns, it follows from Nelson (1990, 1992) that any reasonable volatility model may be used to extract a consistent daily volatility factor, say $\hat{\sigma}_t$. Since this daily volatility factor is largely independent of the intraday pattern, the normalization $\hat{\sigma}_{t,n} \equiv \hat{\sigma}_t / N^{1/2}$, where N denotes the number of intraday returns, readily converts this estimate into the required intraday volatility factor. The second issue requires more work. Market microstructure theory has little to say regarding any specific functional form for the intraday pattern, suggesting the use of a flexible nonparametric type representation for $s_{t,n}$. The flexible Fourier form introduced by Gallant (1981) is ideally suited for modeling periodic patterns and has the advantage of readily incorporating dummy variables to accommodate any discontinuities associated with, e.g., market openings. Defining the regressand in equation (4) as $Y_{t,n} \equiv 2\log|R_{t,n}| - \log\hat{\sigma}_t^2 + \log N$, this leads to the following practical FFF-regression,

$$Y_{t,n} = f(\theta;t,n) + \hat{u}_{t,n}, \tag{5}$$

where the error process $\{\hat{u}_{t,n}\}$ is stationary, and the explanatory variables on the right-hand-side of equation (5) is determined by the flexible Fourier form,

$$f(\theta;t,n) = \Sigma_{j=0}^{J} \mu_j \cdot n^j + \Sigma_{d=1}^{D} \lambda_d \cdot I_d(t,n) + \Sigma_{p=1}^{P} [\delta_{c,p} \cdot \cos(2\pi pn/N) + \delta_{s,p} \cdot \sin(2\pi pn/N)], \tag{6}$$

and $\theta \equiv (\mu_0, \mu_1, \ldots, \mu_J, \lambda_1, \ldots, \lambda_D, \delta_{c,1}, \ldots, \delta_{c,P}, \delta_{s,1}, \ldots, \delta_{s,P})$. The polynomial and the sinusoids in equation (6) capture the overall smooth variation in the intraday pattern, while the (zero-one) indicators, $I_d(t,n)$, allow for discontinuities in response to predetermined events that occur in the interval (t,n). The strength of the associated volatility impact is governed by the λ-coefficients, whereas the overall shape of the average intraday pattern is determined by the μ- and δ-coefficients.

The robustness properties of the FFF-regression approach are noteworthy. As long as the average intraday fluctuations are specified correctly and the associated errors, $\hat{u}_{t,n}$, are stationary, the OLS-regression estimates for θ, defined by equations (5) and (6), will be consistent under very general conditions.[19] Of course, the regression errors will likely be conditionally heteroskedastic and serially correlated. However, valid standard errors that accommodate heteroskedasticity and autocorrelation of unknown form are readily derived from standard procedures for robust asymptotic

[19] In particular, it follows that by increasing the orders J and P any continuous pattern can be approximated arbitrarily well asymptotically. Also, the estimate for the daily volatility factor in $Y_{t,n}$ is only included to help alleviate the conditional heteroskedasticity in the return series, and need not even be consistent for the true daily return volatility factor. The only caveat is a generated regressor problem induced by the estimation of $\hat{\sigma}_t$, but this effect is negligible for any reasonable volatility estimator. A detailed discussion of these issues is provided in the appendix to Andersen and Bollerslev (1998a).

covariance matrix estimation, say $AVar(\hat{\theta})$, for OLS-regressions.[20]

B. Testing for Changes in Intraday Patterns and Shapes

The FFF regression technique provides a relatively simple framework for characterizing the intraday return volatility patterns that permeate financial markets. It is thus natural to adopt the identical econometric methodology for the purpose of empirically discriminating among competing market microstructure theories in an event type setting.

In order to develop a formal inference procedure for general hypotheses regarding the shape of the intraday pattern, we require that each hypothesis be quantified in terms of a scalar ordinal measure, say $g(\theta) \equiv F[f(\theta;t,n)]$. The mapping F should be designed to operationalize the feature of interest in terms of the underlying FFF parameters that characterize the shape of the pattern. This is typically straightforward. For instance, in the context of the "Tokyo Experiment" and the variance ratio statistics discussed in section II, interest centers on the level of volatility during the Japanese lunch period. In this case, the mapping F may be defined as the area under the curve $f(\theta;t,n)$ between 12:00-1:30 Tokyo time, or more formally, the integral of the function $f(\theta;t,n)$ with respect to the corresponding continuous values of n ranging from time 12:00-1:30. Similarly, the steepness of the intraday pattern at a point-in-time, or the average steepness during a particular time-interval, is readily expressed in terms of $\partial f(\theta;t,n)/\partial n$ for a specific value of n, or the integral of $\partial f(\theta;t,n)/\partial n$ for continuous values of n over the relevant time interval(s).

Of course, in practice the parameters θ entering $g(\theta)$ will have to be estimated. However, it follows by the Mean Value Theorem that, when evaluated at the OLS FFF-regression estimate for θ, the function $g(\hat{\theta})$ is asymptotically normally distributed around the true value. Also, by the delta-method, the asymptotic variance of $g(\hat{\theta})$ is consistently estimated by,

$$AVar(g(\hat{\theta})) = G(\hat{\theta}) \cdot AVar(\hat{\theta}) \cdot G(\hat{\theta})', \tag{7}$$

where $G(\theta) \equiv \partial g(\theta)/\partial \theta'$. Furthermore, with the pre- and post-event FFF parameters, say θ^b and θ^a, estimated from non-overlapping samples, the estimation error is asymptotically independent across the two samples. A simple t-test for the null hypothesis of no structural change, or H_0: $g(\theta^b) = g(\theta^a)$,

[20] In particular, it is straightforward to implement either the Newey and West (1987) or the Andrews (1991) heteroskedasticity and autocorrelation consistent covariance matrix estimator in this context.

versus any of the alternative hypotheses, H_A: $g(\theta^b) \neq g(\theta^a)$, $g(\theta^b) > g(\theta^a)$, or $g(\theta^b) < g(\theta^a)$, is therefore readily constructed as,

$$\zeta = [g(\hat{\theta}^a) - g(\hat{\theta}^b)] / [AVar(g(\hat{\theta}^a)) + AVar(g(\hat{\theta}^b))]^{1/2}. \tag{8}$$

The standard normal approximation for the ζ-test statistic that obtains under the null hypothesis of identical pre- and post-event shapes is based solely on the asymptotic normality of the OLS estimates, $\hat{\theta}^a$ and $\hat{\theta}^b$. With the large intraday samples typically available for estimation purposes, and a relatively small number of parameters in θ, this approximation should prove very accurate in most cases. Our reevaluation of the ILM evidence in the following section further illustrates the ideas.

IV. The Tokyo Experiment and ILM Revisited

The statistical inference presented by ILM in support of the notion of private information in the FX market is based exclusively on variance ratio statistics. However, as the discussion in section II makes clear, the corresponding test statistics should be interpreted carefully. Meanwhile, the FFF regression technique developed in the previous section provides a simple robust procedure for testing for specific changes in the intraday return volatility pattern. This section implements these new robust procedures. Beyond reassessing the direct ILM evidence, we also complement their 20-day and 60-day before-versus-after comparisons with the results from our much longer, and less noisy, 2-year sample of high-frequency returns before and after the December 22, 1994, deregulation.

A. *Estimated Intraday Volatility Patterns*

This section presents formal estimates of the intraday ¥-$ return volatility pattern over the Japanese market segment for various sample windows before and after the regulatory change in the Tokyo FX market. Estimation is performed by means of the robust FFF-regression discussed in section III.A. Practical implementation of the approach requires a specific representation of the daily volatility factor, σ_t, and a particular parameterization of the flexible Fourier form in equation (5). First, the values for σ_t were obtained from an MA(1)-GARCH(1,1) model estimated from a daily sample of ¥-$ returns covering the weekdays over the period December 2, 1986, to November 29, 1996. This popular model provides a good first approximation to the daily return volatility process, and should help alleviate the conditional heteroskedasticity at the daily level. Second, the sharp discontinuities in the intraday pattern induced by the Tokyo lunch period led us to use separate specifications before, during and after lunch. Moreover, since the estimation only covers a fraction of the full trading day in the FX market, the sinusoids are less compelling as regressors, and simple third order polynomials were deemed sufficient to capture all significant variation within each of the segments. Finally, we allow for a Japanese market opening effect at 9:00 AM and 1:30 PM Tokyo time, as such terms, a priori, are expected to be important in the pre-event samples, and they were found to be highly significant for the corresponding DM-$ return series in Andersen and Bollerslev (1998a). Thus, the actual FFF-regression is based on the following simplification of equation (6),

$$\begin{aligned} f(\eta;t,n) \quad = \quad & I_M(n) \cdot \Sigma_{j=0}^{3}\, \mu_{M,j} \cdot n^{j} \quad + \quad I_L(n) \cdot \Sigma_{j=0}^{3}\, \mu_{L,j} \cdot (n\text{-}36)^{j} \\ & + \quad I_A(n) \cdot \Sigma_{j=0}^{3}\, \mu_{A,j} \cdot (n\text{-}54)^{j} \quad + \quad \Sigma_{d=1}^{2}\, \lambda_d \cdot I_d(t,n), \end{aligned} \tag{9}$$

where $\mu_{M,j}$ denotes the polynomial coefficients for the morning pattern, 9:00-12:00 ($n \in \{1, \ldots ,36\}$), $\mu_{L,j}$ are the lunch coefficients of relevance for 12:00-1:30 ($n \in \{37, \ldots ,54\}$), $\mu_{A,j}$ are afternoon coefficients relating to the 1:30-3:30 period ($n \in \{55, \ldots ,78\}$), and $I_M(n)$, $I_L(n)$ and $I_A(n)$ are indicator variables taking on the value of unity for the interval, n, belonging to the morning, lunch or afternoon periods respectively, and zero otherwise.[21] Finally, the $I_d(t,n)$ market opening indicators are zero everywhere except for the 9:00-9:05 AM and 1:30-1:35 PM intervals respectively, where they equal one.

The individual FFF-coefficients are void of direct economic interpretation,[22] so we present the estimated FFF-regression results by displaying the implied volatility pattern in the log-absolute return dimension captured by the $Y_{t,n}$ - regressor in section III.A. The robustifying impact of the log-transformation is evident from Figure 3. Although the 20-day pre- and post-event figures remain jagged, the overall shape is quite transparent in the 60-day displays, and the 2-year samples convey remarkable stability in the intraday volatility pattern outside of the Tokyo lunch period.

The ability of low-order polynomials to provide an adequate characterization of the smooth intraday patterns is obvious. It would be straightforward to convert this pattern into corresponding average absolute or squared intraday return patterns under reasonable auxiliary assumptions, as

---

[21] The subtraction of 36 and 54, respectively, for the intervals in the Tokyo lunch and afternoon periods in equation (9) reflects a normalization to prevent the regression coefficient from becoming numerically very small.

[22] The estimated FFF coefficients and their robust asymptotic standard errors are available from the authors upon request. The standard errors were calculated with a Newey and West (1987) covariance matrix estimator and a lag length of 78, or one Japanese trading day, for the 20-day and 60-day estimates, and 390, or five Japanese trading days, for the 2-year estimates.

illustrated in Andersen and Bollerslev (1998a). More importantly, the parsimonious representation of all the basic features of the volatility pattern in terms of a simple OLS-regression allows for powerful and robust inference procedures. The following subsection illustrates this point by providing formal tests for the various hypothesis explored by ILM in this setting.

B. Testing for Changes in Intraday Volatility Patterns

If volatility is caused solely by public information and the flow of public information did not change, the volatility over the lunch period should not be affected by the introduction of lunch period trading. On the other hand, the presence of either private information or mispricing would increase the lunch period volatility following the lifting of the trading restrictions. The first hypothesis investigated by ILM tests for equality of the lunch period volatilities before and after the change in regime.[23] However, instead of the variance ratio statistics employed by ILM, the open versus closed volatilities may be estimated more precisely by the area under the intraday volatility curves during the Tokyo lunch period, say $A_L^o = g(\theta^o)$ and $A_L^c = g(\theta^c)$, where $g(\theta)$ is defined by the integral of $f(\theta;t,n)$ over the relevant range of n.[24] The corresponding one-sided test for no change against the alternative of a heightened private information revelation during the Tokyo lunch period may then be stated as,

$$H_{10} : A_L^o = A_L^c \qquad \textit{versus} \qquad H_{1A} : A_L^o > A_L^c .$$

The outcome of this test, as detailed in Table I, is directly in line with the conclusion reached by ILM, and the estimates for $f(\theta;t,n)$ in Figure 3 above.[25] For all three horizons, the null hypothesis of no change is clearly rejected. Still, the p-value for the 20-days before-versus-after test of 0.005 is notably larger than all p-values reported by ILM. Meanwhile, recall from the discussion in section II above that, when properly evaluated, the standard 20-days variance ratio statistic for the lunch

---

[23] Our numbering of the different hypotheses follows that of ILM. We refer to their paper for a more elaborate discussion of the underlying market microstructure motivations.

[24] For the FFF regression in equation (9), this integral translates into $\mu_{L,0}\, n_L + (\mu_{L,1}/2)\, n_L^2 + (\mu_{L,2}/3)\, n_L^3 + (\mu_{L,3}/4)\, n_L^4$, where $n_L = 18$ denotes the number of five-minute intervals in the Tokyo lunch period. Similar expressions for the remaining hypothesis tests described below are available from the authors on request.

[25] Consistent with the notation in equation (9), all of the ζ-test statistics in Table I are based on the open minus close (or after minus before) values for the function $g(\theta)$; i.e., $g(\hat{\theta}^o) - g(\hat{\theta}^c)$. The p-values for the alternative hypotheses reflect this convention.

period is not even significant at the five percent level. As such, these results highlight the increased test power afforded by the semi-parametric FFF procedures vis-a-vis the standard non-parametric variance ratio methodology.

If the total amount of private information produced did not change as a result of the lifting of the trading restrictions, lunch period trading should redistribute some of this same information more evenly throughout the day. This suggests that a flattening of the typical U-shape in the intraday volatility should occur. In a U-shaped curve, the average slope of the second half of the curve, say S_A, is positive, whereas the average slope of the first half of the curve, say S_M, is negative. The difference in the average slopes, say $F \equiv S_A - S_M = g(\theta)$, thus provides a natural measure of the curvature. The average slopes S_M and S_A are easily quantified by the integral of $\partial f(\theta;\sigma_t,n)/\partial n$ over the morning and afternoon trading segments, respectively. The second hypothesis investigated by ILM may therefore be re-stated as,

$$H_{20}: F^o = F^c \qquad versus \qquad H_{2A}: F^o < F^c .$$

Whereas the variance ratio statistics reported by ILM seemingly provide overwhelming support for such a flattening of the U-shape, none of the robust tests in Table I are significant. In fact, for the 60-days and 2-years horizons, the ζ-test statistics are in the wrong tail of the distribution.

When trading is restricted during the lunch period and before 9:00, traditional market microstructure theories predict the existence of a separate U-shaped volatility pattern over the morning trading session. This morning U-shape should disappear when the trading restrictions are lifted. Let S_{EM} and S_{LM} denoted the average slope of the early-morning and late-morning curves,[26] respectively, while $F_M \equiv S_{LM} - S_{EM}$. Again, these statistics may be estimated directly from the integral of the first derivative of the FFF functional over the relevant time intervals. The third hypothesis stipulating the disappearance of the morning U-shape then translates into,

$$H_{30}: F_M^o = F_M^c \qquad versus \qquad H_{3A}: F_M^o < F_M^c .$$

The variance ratio tests reported in ILM strongly rejects this hypothesis. ILM also assert this finding as " ... perhaps the most compelling single fact of the four ..." in favor of the private information

[26] The early-morning and late-morning segments are defined as 9:00-10:30 and 10:30-12:00, respectively.

hypothesis. However, for none of the three horizons is the robust test for H_{30} in Table I anywhere close to significant. Visual inspection of the estimated volatility patterns in Figure 3 also suggests that the evidence in favor of the disappearance of a morning U-shape is rather dubious.

The final ILM test concerns the apparent increase in the volatility towards the end of the Japanese trading day. If information is short-lived, and trading is restricted over the lunch period, privately informed investors would have an incentive to trade early before their information might be revealed to others. Consequently, some volatility from the morning will "move" to the afternoon, causing an upward tilt in the full day volatility U-shape. This hypothesis is succinctly stated as,

$$H_{40}: S_A^o = A_A^c \qquad \text{versus} \qquad H_{4A}: S_A^o > S_A^c,$$

where, as above, S_A denotes the average slope of the FFF volatility curve during the afternoon trading segment. None of the three tests in Table I suggests that such an upward tilt occurred. Again, these results are in sharp contrast to the apparent strong statistical evidence provided by the variance ratio tests in ILM.

In addition to these four hypotheses investigated by ILM, we offer two additional tests for changes in the intraday volatility pattern based directly on our before and after FFF regression estimates. The market microstructure theories cited by ILM in support of the role of private information also predict a distinct market opening effect in the volatility. The estimated FFF curves include a dummy variable for the 9:00-9:05 time interval, λ_I, to take account of this effect. Thus, if private information is important, this market-opening effect should also diminish in response to the lifting of the trading restrictions. Our H_M hypothesis formalizes this idea,

$$H_{M0}: \lambda_I^o = \lambda_I^c \qquad \text{versus} \qquad H_{MA}: \lambda_I^o < \lambda_I^c.$$

The results in Table I indicate that there is no evidence for a diminished volatility at the opening. In fact, not only are all of the before and after estimates for λ_I individually significant, but for all three horizons, the estimate for λ_I^o actually exceed that of λ_I^c.[27]

Prior to December 22, 1994, the average intraday volatility pattern also exhibit a sharp peak

[27] The presence of a spike in volatility at 9:00 Tokyo time may be due to the simultaneous opening of a number of Japanese financial markets at this time. The price innovation at the opening in equity and bond markets may well impact FX rates.

immediately following the Tokyo lunch at 1:30 local time. The λ_2 dummy variable for the 1:30-1:35 interval captures this effect. As above, this post-lunch opening effect should diminish in response to the lifting of the lunch period trading restrictions. This is tested by the H_L hypothesis,

$$H_{LO}: \lambda_2^o = \lambda_2^c \qquad versus \qquad H_{LA}: \lambda_2^o < \lambda_2^c .$$

Consistent with the discussion of the patterns in section IV.A, the formal tests in Table I do provide marginal support for a diminished spike in post-lunch volatility. The significance of this hypothesis is clearly related to the evidence of heightened lunch-time volatility which is supported by our earlier results. As FX trading picks up over the Tokyo lunch period, the post-lunch opening should increasingly become a non-event.[28] At the same time, the magnitude of this effect is trivial relative to the overall daily volatility.

V. Concluding Remarks

The empirical market microstructure literature has documented pronounced intraday patterns in a number of key financial market variables. These striking regularities have inspired theorists to explain how such features may arise from the interaction of distinct customer groups, trading professionals, and specialists or market makers in a given institutional setting. As such, this literature may help shed light on intrinsic differences between existing financial markets, the optimal design of market mechanisms, our ability to disentangle short-lived effects from more permanent price shocks, as well as the implementation of better trading and hedging strategies by market participants.

Meanwhile, the empirical analysis of high-frequency data presents a set of unique statistical problems. This paper develops new tools for robust empirical analysis of intraday patterns that are capable of handling these complications. We illustrate the importance of these procedures by revisiting the "Tokyo Experiment" investigated by ILM. We confirm the finding of enhanced volatility during the Tokyo lunch period following the lifting of the trading restrictions. However, none of the other market microstructure hypotheses stipulated by ILM are supported by the robust tests when applied to our longer pre- and post-event windows.

[28] Although the diminished post-lunch spike is consistent with an important role for private information, there are other plausible explanations, see, e.g., Hsieh and Kleidon (1996) for a discussion.

Although most of our statistical tests provide evidence counter to the ILM conclusions, we do not interpret the findings as rejecting the relevance of private information for the spot foreign exchange interbank market per se. In fact, the heightened volatility during the lunch period may be interpreted as providing support for the notion of private information. At the same time, the regulatory change do not appear significant enough to induce a detectable shift in the pattern outside of the Tokyo lunch period. At a more general level, the empirical analysis in the paper clearly illustrate the importance of properly accounting for all of the forces at work in the high-frequency data context. As such, the framework developed here should be useful for the investigation of a wide range of qualitative market microstructure theories.

References

Admati, A.R. and P. Pfleiderer (1988), "A Theory of Intraday Patterns: Volume and Price Variability," Review of Financial Studies, 1, 3-40.

Amihud, Y. and H. Mendelson (1987), "Trading Mechanisms and Stock Returns: An Empirical Investigation," Journal of Finance, 42, 533-553.

Amihud, Y. and H. Mendelson (1991), "Volatility, Efficiency and Trading: Evidence from the Japanese Stock Market," Journal of Finance, 46, 1765-1789.

Andersen, T.G. and T. Bollerslev (1997a), "Intraday Periodicity and Volatility Persistence in Financial Markets," Journal of Empirical Finance, 4, 115-158.

Andersen, T.G. and T. Bollerslev (1997b), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," Journal of Finance, 52, 975-1005.

Andersen, T.G. and T. Bollerslev (1998a), "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies," Journal of Finance, 53, 219-265.

Andersen, T.G. and T. Bollerslev (1998b), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," International Economic Review, forthcoming.

Andersen, T.G. and T. Bollerslev (1998c), "Towards a Unified Framework for High- and Low-Frequency Return Volatility Modeling," Statistica Neerlandica, forthcoming.

Andrews, D.W.K. (1991), "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," Econometrica, 59, 817-858.

Ball, R. and P. Brown (1968), "An Empirical Evaluation of Accounting Income Numbers," Journal of Accounting Research, 6, 159-178.

Barclay, M., R. Litzenberger, and J. Warner (1990), "Private Information, Trading Volume, and Stock Return Variances," Review of Financial Studies, 3, 233-253.

Bollerslev, T., R.F. Engle and D.B. Nelson (1994), "ARCH Models," in Handbook of Econometrics, Vol.IV, (R.F. Engle and D. McFadden, eds.), Amsterdam: North Holland Press.

Cheung, Y.W. and C.Y.P. Wong (1997), "Foreign Exchange Markets in Hong Kong, Tokyo, and Singapore," manuscript, Department of Economics, University of California, Santa Cruz.

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

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

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

Evans, M.D.D. (1998), "The Microstructure of Foreign Exchange Dynamics," manuscript, Department of Economics, Georgetown University.

Fama, E., Fisher, L., Jensen, M. and R. Roll (1969), "The Adjustment of Stock Prices to New Information," International Economic Review, 10, 1-21.

Foster, F.D. and S. Viswanathan (1990), "A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Security Markets," Review of Financial Studies, 3, 593-624.

Foster, F.D. and S. Viswanathan (1993), "Variations in Trading Volume, Return Volatility, and Trading Costs: Evidence on Recent Price Formation Models," Journal of Finance, 48, 187-211.

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

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

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

Hsieh, D. and A. Kleidon (1996), "Bid-Ask Spreads in Foreign Exchange Markets: Implications for Models of Asymmetric Information," in The Microstructure of Foreign Exchange Markets (J.A. Frankel, G. Galli and A. Giovannini, eds.), Chicago: University of Chicago Press.

Ito, T., R.K. Lyons and M.T. Melvin (1998), "Is There Private Information in the FX Market? The Tokyo Experiment," Journal of Finance, 53, 1111-1130.

Johnson, N.L. and S. Kotz (1970), Continuous Univariate Distributions-2, New York: John Wiley and Sons.

Jones, C.M., G. Kaul, and M.L. Lipson (1994), "Information, Trading and Volatility," Journal of Financial Economics, 36, 127-153.

Jones, C.M. and G. Kaul (1994), "On the Use of Variance Ratios," manuscript, Department of Finance, School of Business Administration, University of Michigan.

Lo, A.W. and A.C. MacKinlay (1988), "Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test," Review of Financial Studies, 1, 41-66.

Lo, A.W. and A.C. MacKinlay (1989), "The Size and Power of the Variance Ratio Test in Finite Samples: A Monte Carlo Investigation," Journal of Econometrics, 40, 203-238.

Müller, U.A., M.M. Dacorogna, R.D. Davé, R.B. Olsen, O.V. Pictet, and J.E. von Weizsäcker (1997), "Volatilities at Different Time Resolutions - Analyzing the Dynamics of Market Components," Journal of Empirical Finance, 4, 213-239.

Nelson, D.B. (1990), "ARCH Models as Diffusion Approximations," Journal of Econometrics, 45, 7-38.

Nelson, D.B. (1992), "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," Journal of Econometrics, 52, 61-90.

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

O'Hara, M. (1995), Market Microstructure Theory, Cambridge, MA: Blackwell Publishers.

Oldfield, G.S. Jr., and R.J. Rogalski (1980), "A Theory of Common Stock Returns over Trading and Nontrading Periods," Journal of Finance, 35, 729-751.

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.

Poterba, J.M. and L.H. Summers (1988), "Mean Reversion in Stock Prices, Evidence and Implications," Journal of Financial Economics, 22, 27-58.

Richardson, M. and T. Smith (1991), "Tests of Financial Models in the Presence of Overlapping Observations," Review of Financial Studies, 4, 227-254.

Smith, T. (1994), "Econometrics of Financial Models and Market Microstructure Effects," Journal of Financial and Quantitative Analysis, 29, 519.

Stoll, H.R. and R.E. Whaley (1990), "Stock Market Structure and Volatility," Review of Financial Studies, 3, 37-71.

Wei, S.J. and J. Kim (1997), "The Big Players in the Foreign Exchange Market: Do They Trade on Information or Noise?" NBER Working Paper No.6256.

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

Table I

Tests for Changes in the Intraday ¥-$ Return Volatility Pattern
and the Relevance of Private Information in the Foreign Exchange Market

Hypothesis	Horizon	ζ-statistic	p-value
	2-years	8.00	0.000
H_1	60-days	4.65	0.000
	20-days	2.57	0.005
	2-years	0.16	0.564
H_2	60-days	0.34	0.633
	20-days	-1.28	0.100
	2-years	-0.36	0.359
H_3	60-days	0.16	0.564
	20-days	0.20	0.579
	2-years	0.12	0.452
H_4	60-days	-0.55	0.709
	20-days	1.20	0.115
	2-years	0.70	0.758
H_M	60-days	0.40	0.655
	20-days	0.48	0.684
	2-years	-1.62	0.053
H_L	60-days	-1.38	0.084
	20-days	-1.57	0.058

The table reports robust FFF based tests for the null hypothesis of no change in the intraday ¥-$ volatility pattern following the lifting of the trading restrictions during the Tokyo lunch period on December 22, 1994. The tests are based on the ζ-statistic defined in equation (8), and comparisons involving non-weekend five-minute ¥-$ returns for the 20 days, 60 days, and 2 years before December 22, 1994, versus after January 4, 1995. The H_1 hypothesis tests for an increase in the total volatility during the Tokyo lunch period; H_2 tests for a flattening of the U-shape over the full trading day; H_3 tests for a flattening of the U-shape during the morning trading segment; H_4 test for an upward tilt in the full-day volatility U-shape; H_M tests for an increase in the volatility peak at the morning opening; H_L tests for an increase in the volatility peak at the post-lunch opening.

Figure I

The figure graphs the non-weekend, average absolute five-minute ¥-$ returns for the 20 days, 60 days and 2 years before December 22, 1994, and after January 4, 1995.

Figure II

The figure displays the empirical distribution of the rolling variance ratio statistics for the eight years preceding the December 22, 1994, regulatory change along with the simulated distribution from a high-frequency GARCH(1,1) model and the standard F-distributions.

Figure III

The figure compares the non-weekend average robust absolute ¥-$ return regressand of equation (5) to the corresponding flexible Fourier form fit for the 20 days, 60 days and 2 years before December 22, 1994, and after January 4, 1995.

FIGURE 1.A

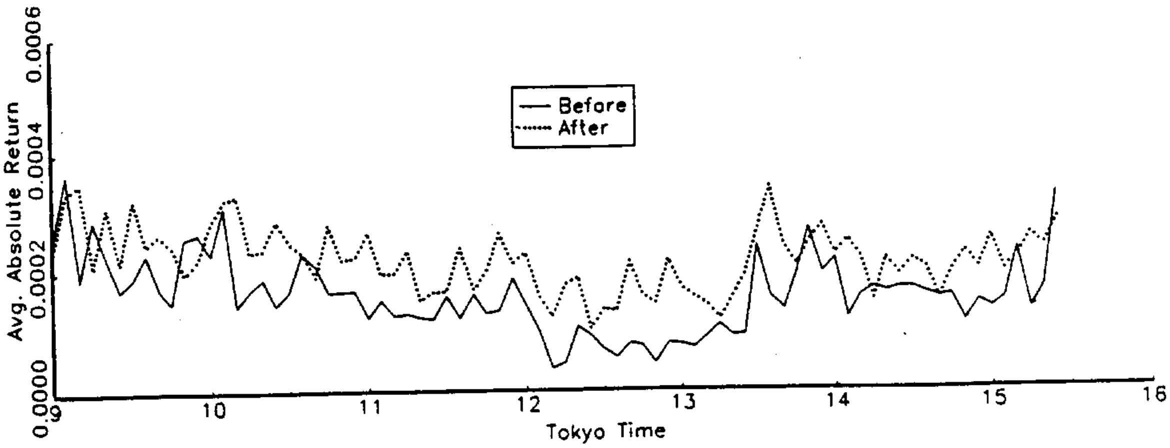

FIGURE 1.B

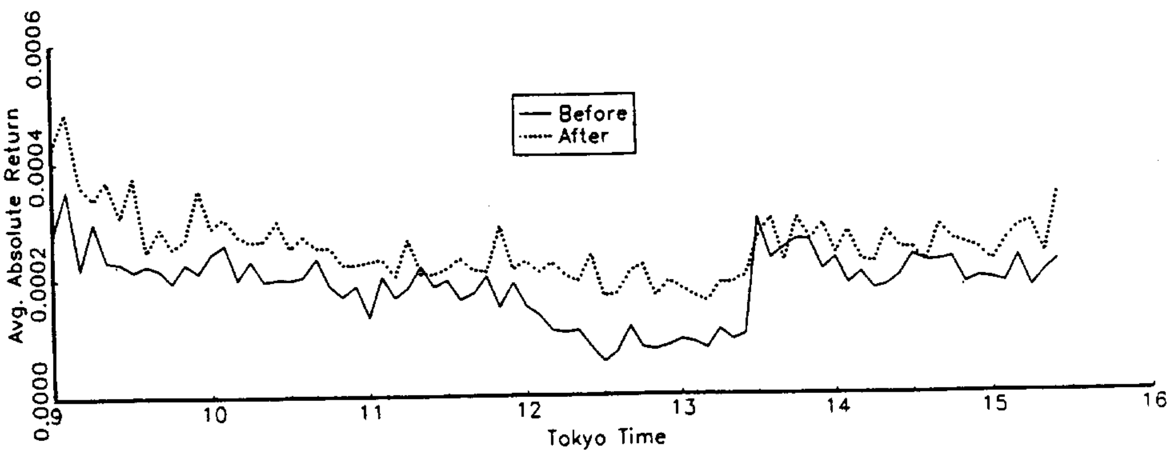

FIGURE 1.C

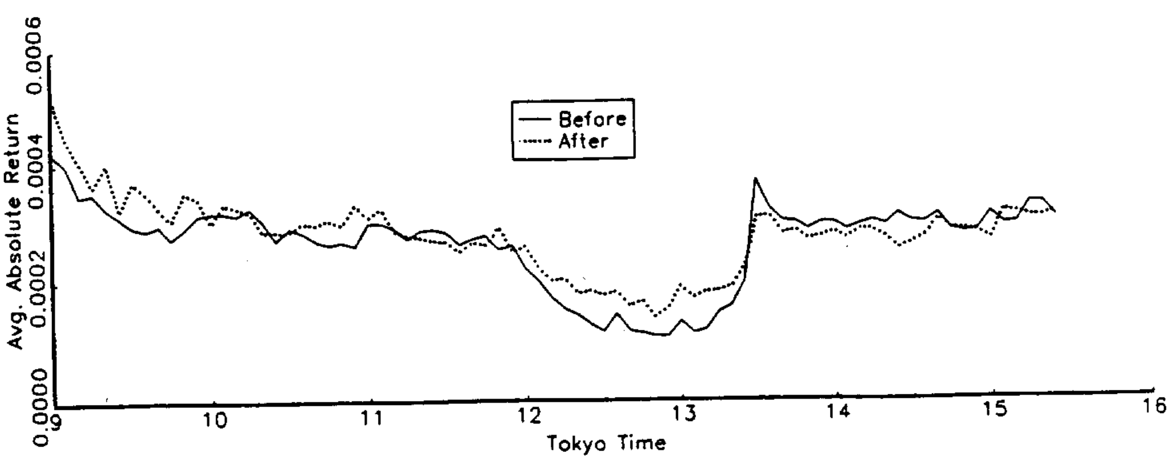

FIGURE 2.A

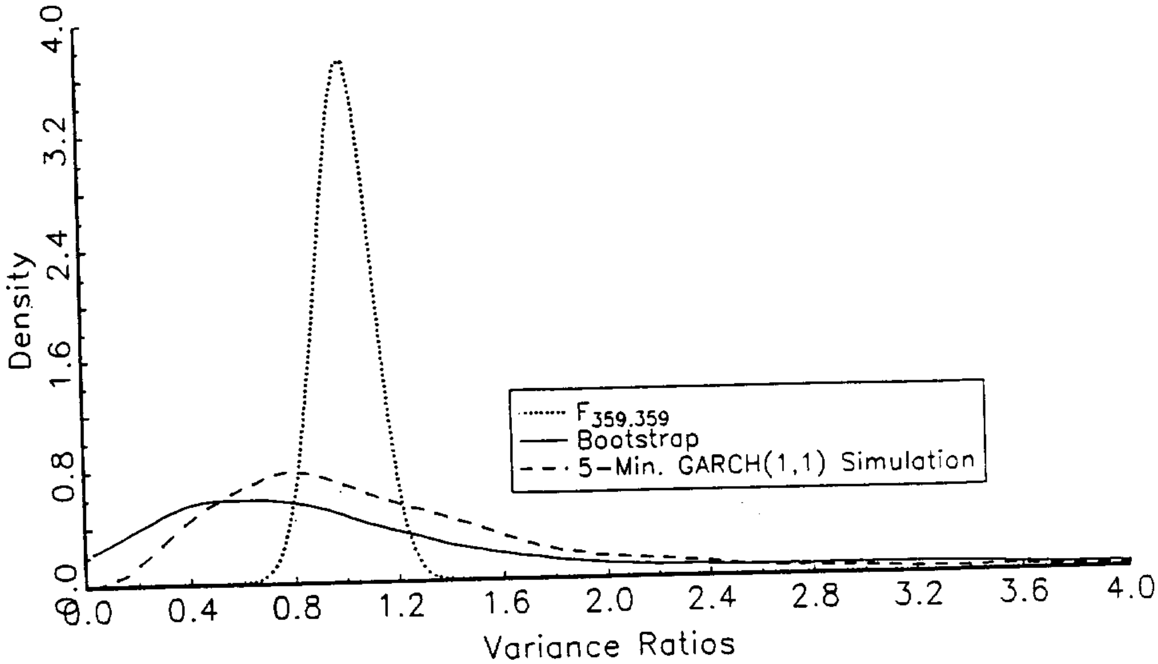

FIGURE 2.B

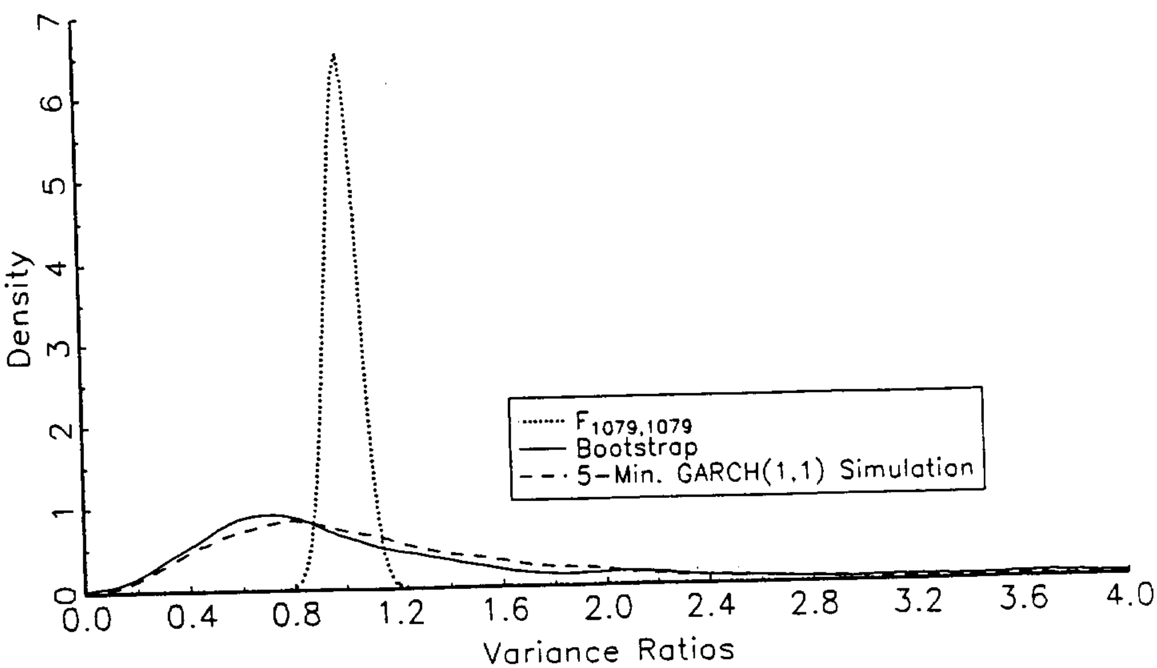

FIGURE 3.A

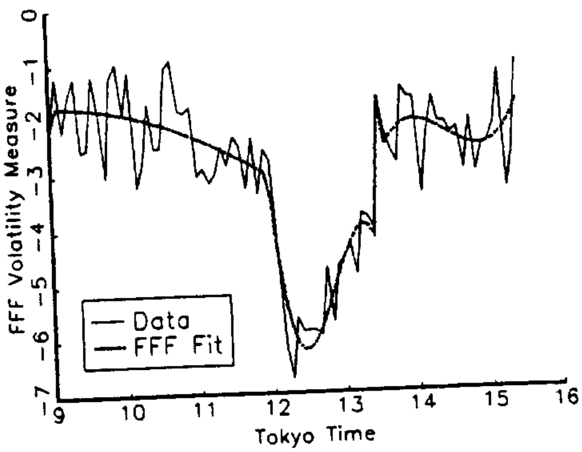

FIGURE 3.B

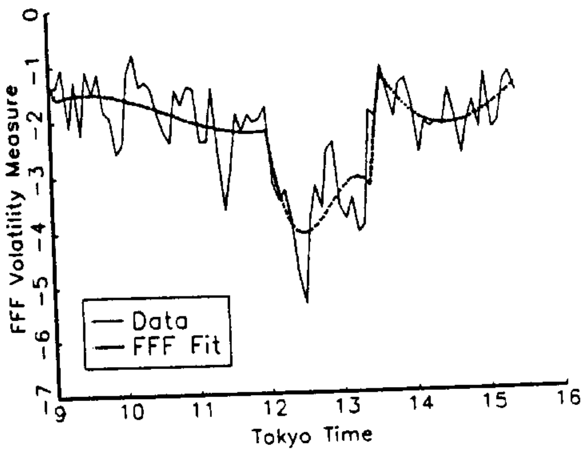

FIGURE 3.C

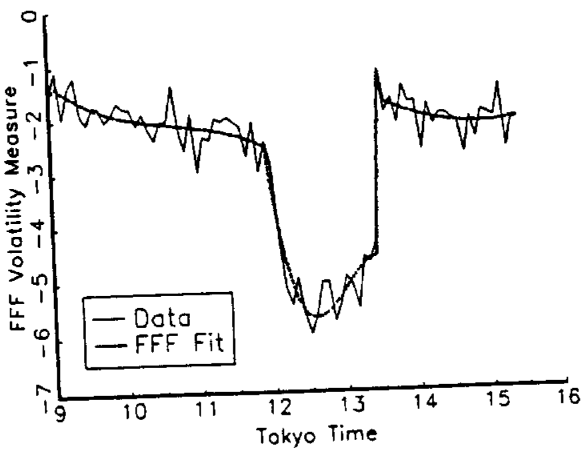

FIGURE 3.D

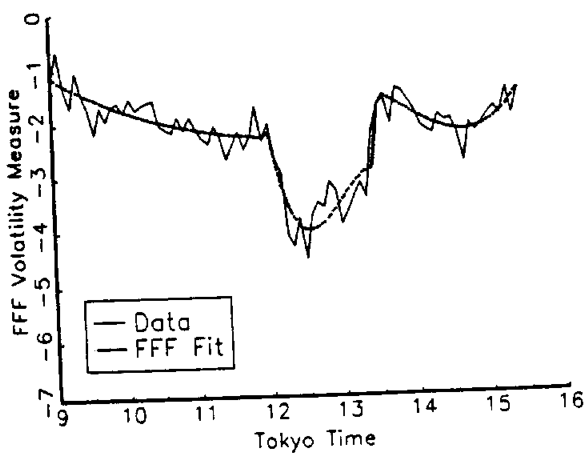

FIGURE 3.E

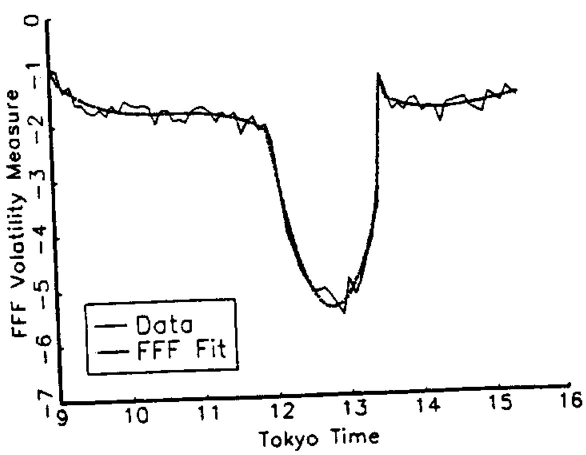

FIGURE 3.F