

Variance-ratio Statistics and High-frequency Data: Testing for Changes in Intraday Volatility Patterns

TORBEN G. ANDERSEN, TIM BOLLERSLEV, and ASHISH DAS*

ABSTRACT

Variance-ratio tests are routinely employed to assess the variation in return volatility over time and across markets. However, such tests are not statistically robust and can be seriously misleading within a high-frequency context. We develop improved inference procedures using a Fourier Flexible Form regression framework. The practical significance is illustrated through tests for changes in the FX intraday volatility pattern following the removal of trading restrictions in Tokyo. Contrary to earlier evidence, we find no discernible changes outside of the Tokyo lunch period. We ascribe the difference to the fragile finite-sample inference of conventional variance-ratio procedures and a single outlier.

MOST FINANCIAL ACTIVITY VARIABLES, including return volatility and bid-ask spreads, display pronounced intraday patterns. For markets with well-defined daily openings and closures, including equity and bond markets, these activity measures typically conform to a distorted U-shape over the trading day (see, e.g., Wood, McInish, and Ord (1985) and Harris (1986)). In contrast, markets with round-the-clock trading, such as the foreign exchange (FX) interbank market, produce more complex patterns, related to the ebb and flow of activity across the global financial centers (see, e.g., Dacorogna et al. (1993) and Andersen and Bollerslev (1998a)). These striking regularities have inspired a large literature, seeking to explain how such patterns may arise from the interaction of distinct customer groups and market makers. The rationalizations center on the prevalence of private (asymmetric) versus public information, the risk and inventory cost borne by market makers, and the institutional setting. Early contributions include Admati and Pfleiderer (1988) and Foster and Viswanathan (1990).

* Department of Finance, J. L. Kellogg Graduate School of Management, Northwestern University (Andersen and Das) and Department of Economics, Duke University and NBER (Bollerslev). An earlier version of this paper circulated under the title "Testing for Microstructure Effects in Volatility: Revisiting the Tokyo Experiment." This work is supported by a grant from the National Science Foundation to the NBER. We are also grateful to Olsen & Associates for making the intraday exchange rate quotations available, and to Mike Melvin for forwarding his version of the data. Finally, we acknowledge comments from brown bag seminar participants at Northwestern University and London School of Economics, Richard Lyons, a referee, and the editor René Stulz.

Empirical work in the area confronts several obstacles. First, key theoretical concepts are often inherently unobservable, such as the importance of asymmetric information. Second, theories usually do not provide quantitative predictions, but merely qualitative characterizations of the intraday pattern. Third, comparisons across market structures are complicated by the fact that they differ along many dimensions simultaneously. Consequently, much empirical work has resorted to regime-shift and event studies, where the impact of exogenous changes in the environment may be gauged within a specific setting, thus providing an opportunity to discriminate among theories that offer contrasting predictions in the wake of the change. Examples include modifications in trading hours, the release of firm-specific information, or regularly scheduled macroeconomic announcements.¹

Many empirical market microstructure studies rely on standard variance-ratio procedures. For instance, Amihud and Mendelson (1987, 1991) and Stoll and Whaley (1990) study the performance of different trading mechanisms, whereas Forster and George (1996) examine the effect of cross-listings of U.S. securities on the basis of variance ratios of open-to-open versus close-to-close returns. Similarly, Oldfield and Rogalski (1980), French and Roll (1986), Harvey and Huang (1991), and Jones, Kaul, and Lipson (1994) among others, analyze the rate of information flow during trading and nontrading periods by comparing variance ratios for open-to-close versus close-to-open returns. For surveys of this literature we refer to Jones and Kaul (1994) and Ronen (1997). Whereas these earlier studies are based on interdaily returns, the recent contribution of Ito, Lyons, and Melvin (1998) provides an intriguing extension to the high-frequency data setting by specifying a string of variance-ratio statistics designed to test for specific types of changes in the *intraday* volatility pattern. Specifically, they study the impact of the removal of intraday FX trading restrictions on Japanese based banks over the Tokyo lunch period, and their approach has already motivated several other studies, including the tests for market closure effects in the Hong Kong equity index by Ho and Lee (1998).

Because the number of intraday observations is typically very large, it may appear that standard variance-ratio statistics or comparable tests derived from asymptotic distribution theory should provide good approximations in the high-frequency setting and that test power should be correspondingly high. Such intuition is deceptive. Besides the pronounced intraday volatility pattern, which motivates much of the market microstructure literature, high-frequency returns embody highly persistent conditionally heteroskedastic components along with discrete information arrival effects.

¹ The event study methodology was introduced by Ball and Brown (1968) and Fama et al. (1969), but applications to market microstructure issues are more recent. For example, Barclay et al. (1990) follow French and Roll (1986) in exploring the discrepancy in volatility over closed-versus-open trading periods. Likewise, announcement effects have long been explored by event study techniques, but the use of high-frequency data is fairly new (see, e.g., Ederington and Lee (1993)).

Combined, these factors generate serial correlation and extreme outliers in the intraday squared returns that render standard variance-ratio procedures unreliable. Recognizing these complications, this paper develops new robust tools for inference in the high-frequency data setting. Our approach is based on an extension of the Fourier Flexible Form (FFF) regression framework proposed by Andersen and Bollerslev (1997a, 1998a). This methodology is effective in modeling the intraday volatility patterns, and we demonstrate that the identical framework may be exploited more generally in testing for changes in the intraday patterns, expressed in terms of functionals of the underlying FFF parameters.

Upon applying these new robust procedures to test for a change in the intraday volatility pattern of the U.S. dollar–Japanese yen FX market following the removal of trading restrictions for Japanese-based banks on December 22, 1994, we confirm that volatility increases over the Tokyo lunch period but find no discernible change in the pattern outside of lunch. Meanwhile, we also document that the highly significant variance-ratio statistics in Ito et al. (1998) for shifts in the volatility pattern across the entire Japanese trading day are driven by a single questionable indicative quote. Removing this one observation reverses the inference.² Of course, this does not dispel the relevance of private information per se—it may simply be that the lifting of the restrictions has an insignificant impact on the optimal trading strategies of informed agents—but it is clear that our results weaken the Ito et al. (1998) case for private information, based on the “Tokyo Experiment,” substantially.³

The rest of the paper is structured as follows. Using a simple simulation design, Section I illustrates the nonrobustness of the standard variance-ratio methodology in the high-frequency data setting. Section II first reviews the FFF approach for estimating intraday volatility patterns, and then sets out the new procedures to test for qualitative changes in such patterns. Section III applies these tools to the Ito et al. (1998) hypotheses and also discusses the economic implications of our findings, and Section IV concludes. The appendix provides supplementary material further documenting the fragility of the standard variance-ratio methodology in the high-frequency setting.

² The observation highlights the difficulty in constructing automatic outlier detection filters and the dangers of relying exclusively on such to guard against the impact of outliers, as has been common in the literature. We emphasize the need to carefully inspect influential outliers and the development of robust procedures that will deliver reliable inference in the face of a few erroneous observations.

³ The importance of correct inference concerning this question is underscored by the fact that the Ito et al. (1998) findings have motivated much recent academic work seeking to explain the apparent change in the volatility patterns throughout the trading day along with explicit private information sources (see, e.g., Peiers (1997), Wei and Kim (1997), Covrig and Melvin (1998), Hau (1998), Sapp (1998), and Payne (1999) for examples of this style of work. The Ito et al. (1998) study also prompted a write-up in the May 31, 1997, *Economist*, entitled “Forex with Rice.”

I. On the Distribution of Intraday Variance-ratio Tests

Variance-ratio statistics are routinely employed in empirical market microstructure studies to assess return volatility and the rate of information flow over time and across market structures. Because the number of intraday observations is typically very large, it may seem plausible that standard variance-ratio statistics and comparable tests derived from asymptotic distribution theory should work very well, for example, “. . . 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 . . .” (Ito et al. (1998), p. 1122). Unfortunately, this intuition is misleading.

For concreteness, we first explore the inference provided by traditional variance-ratio statistics and standard asymptotic theory in testing the generic null hypothesis of no change in return volatility over a specific intraday interval following some exogenous event. This null hypothesis may be stated formally in terms of the variance-ratio,

$$V^a/V^b = 1, \quad (1)$$

where V^b and V^a denote the interval specific return variability for the before-the-event and after-the-event regimes, respectively. If the high-frequency returns are i.i.d. normally distributed and the null hypothesis is valid, the corresponding sample statistic, \hat{V}^a/\hat{V}^b , represents a realization of an $F_{na-1, nb-1}$ distributed random variable, where na and nb are the number of returns included in the computation of the respective sample variances. When na and nb are large, this F -distribution is approximately normal (see, e.g., Johnson and Kotz (1970)).

It is well known, however, that high-frequency returns are characterized by pronounced volatility clustering effects. To convey the distortion this may induce in the distribution of the variance-ratio statistic,⁴ suppose that the conditional variance for the intraday return, $R_{t,n}$, over the n th intraday time interval on day t , is determined by the GARCH(1,1) model,

$$\sigma_{t,n}^2 = \omega + \sigma_{t,n-1}^2 \cdot (\alpha \cdot Z_{t,n-1}^2 + \beta), \quad (2)$$

where $R_{t,n} \equiv \sigma_{t,n} \cdot Z_{t,n}$, and $Z_{t,n}$ is i.i.d. $N(0,1)$.⁵

⁴ The current investigation of the impact of conditional heteroskedasticity on nonoverlapping return variance ratio statistics is conceptually distinct from, and has no direct implications for, the short-horizon versus long-horizon variance-ratio statistics employed in testing for serially correlated returns, as investigated by Lo and MacKinlay (1988, 1989), Richardson and Smith (1991), Smith (1994), and Poterba and Summers (1988), among others.

⁵ Numerous models have been proposed for characterizing the volatility clustering at daily and lower frequencies (see, e.g., Bollerslev, Engle, and Nelson (1994)). Although the GARCH(1,1) model provides a good description at the interdaily level, it misses some long-memory features, the complex component structure, and systematic intraday volatility patterns documented by, for example, Andersen and Bollerslev (1997a, 1997b, 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.

We calibrate the high-frequency return process to estimates based on daily returns from markets with minimal microstructure frictions. Specifically, we simulate five-minute returns that are consistent with daily GARCH parameters of $\hat{\omega} = 0.026$, $\hat{\alpha} = 0.104$, and $\hat{\beta} = 0.844$.⁶ The trading day comprises a full 24 hours, or 288 5-minute intervals, and we test for equality of the return volatility over a specific $1\frac{1}{2}$ hour (18 5-minute intervals) period. Inspired by the event study methodology, we compare the return variances over fairly short samples around the event. Our leading illustration uses a 20-day “before” and “after” window. Using 18 5-minute returns (out of the daily total of 288) per day, the sample variances, \hat{V}^b and \hat{V}^a , are each based on $20 \cdot 18 = 360$ nonoverlapping observations. Finally, we exclude observations from days just around the event to avoid the impact of any particular turmoil associated with the regime switch, or an inherent uncertainty about the exact timing of the event. We separate the before and after event windows by 10 days, or 2,880 5-minute returns, corresponding to a full 2 weeks of (weekday) trading.⁷

The solid line in Figure 1 displays the $F_{359,359}$ distribution for the \hat{V}^a/\hat{V}^b statistic under the ideal assumption of i.i.d. normally distributed returns, or $\alpha = \beta = 0$. The dashed line conveys the simulated distribution of the same variance-ratio statistic for the empirically more realistic GARCH(1,1) model described above. Although the *unconditional* variance is constant, so that the null hypothesis is true, it is evident that the F -distribution provides an extremely poor approximation to the true sampling distributions for the variance-ratio statistic in the presence of persistent time-varying *conditional* volatility, or *nonnormally* distributed returns. Clearly, if the GARCH process provides a reasonable approximation to the properties of the intraday returns, then standard inference based on variance-ratio statistics will reject the null hypothesis in the majority of cases, even when it is valid, that is, test size is badly distorted. For example, a realization of the variance-ratio statistic of 2.43 is in the extreme upper tail of the $F_{359,359}$ distribution, but is at best mildly significant in the simulated GARCH(1,1) distribution with an associated p -value of 0.082.⁸

⁶ These specific parameter estimates were obtained from daily yen–dollar returns for October 1, 1987 to September 30, 1992, in Andersen and Bollerslev (1998b). From Drost and Nijman (1993), the corresponding implied parameters at the five-minute frequency are then $\omega = 0.308 \cdot 10^{-6}$, $\alpha = 0.00924$, and $\beta = 0.991$. If anything, the implied volatility persistence at the daily level is below typical estimates in the literature, so the associated inference problems likely constitute a conservative estimate of the distortions encountered in actual applied work with high-frequency data.

⁷ This also alleviates a potential statistical problem. Given the volatility dependence, the sample variances for the intraday interval obtained during the before and after regimes are not strictly independent. Leaving a gap between the two windows helps to reduce this problem. Our particular choices of daily trading cycle, window length, and separation of before and after event samples are guided by the approach in Ito et al. (1998). We study their findings in more detail in Section III. However, none of these specific choices has a material impact on the qualitative conclusions of our simulation study.

⁸ This particular value corresponds to the actual variance-ratio statistic for a significant increase in Tokyo lunch period volatility in the FX market following the lifting of trading restrictions on Japanese banks over lunch.

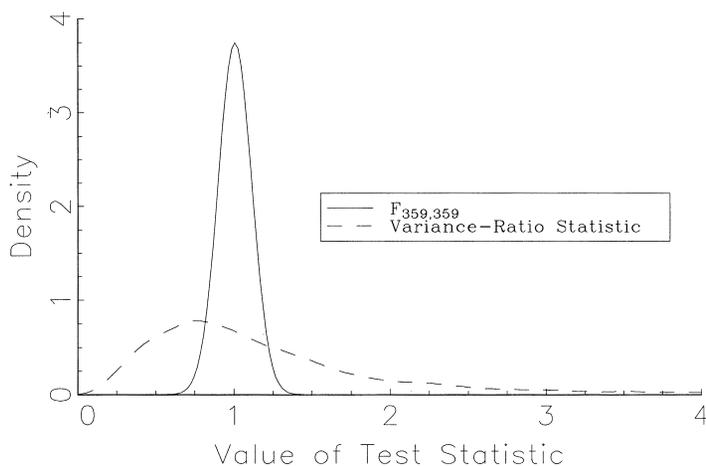


Figure 1. Variance-ratio statistics. The figure displays the distribution of the variance-ratio statistic in equation (1) testing the null hypotheses of constant unconditional variance based on simulated five-minute returns. The statistic is calculated from 18 5-minute return observations per day over 20 days before and after. The returns are generated by the high-frequency GARCH(1,1) model described in the text. The solid line labeled $F_{359,359}$ gives the F -distribution that obtains with i.i.d. normally distributed returns.

The sample variance is based on squared returns and may be heavily influenced by a few extreme observations. Volatility measures constructed from absolute returns are less sensitive to outliers. Hence, an equally valid, but presumably more robust, test for a change in volatility may be obtained from the absolute returns. To illustrate, let A^b and A^a denote the average absolute returns in the before and after regimes. The null hypothesis of constant volatility may then be stated as

$$A^a - A^b = 0. \quad (3)$$

Valid asymptotic standard errors for the sample mean absolute returns, say $A \text{Var}(\hat{A}^a)$ and $A \text{Var}(\hat{A}^b)$, are readily obtained from the standard procedures for robust covariance matrix estimation in the context of GMM (see, e.g., Newey and West (1987) and Andrews (1991)). Thus, assuming that the estimation errors for the nonoverlapping pre- and postevent samples are asymptotically independent, the hypothesis in equation (3) may be tested by means of a standard t -test,

$$Z = [\hat{A}^a - \hat{A}^b] / [A \text{Var}(\hat{A}^a) + A \text{Var}(\hat{A}^b)]^{1/2}. \quad (4)$$

If the use of squared deviations and the presence of influential outliers are the two main reasons behind the poor performance of the variance-ratio statistic in Figure 1, this test should be better behaved.

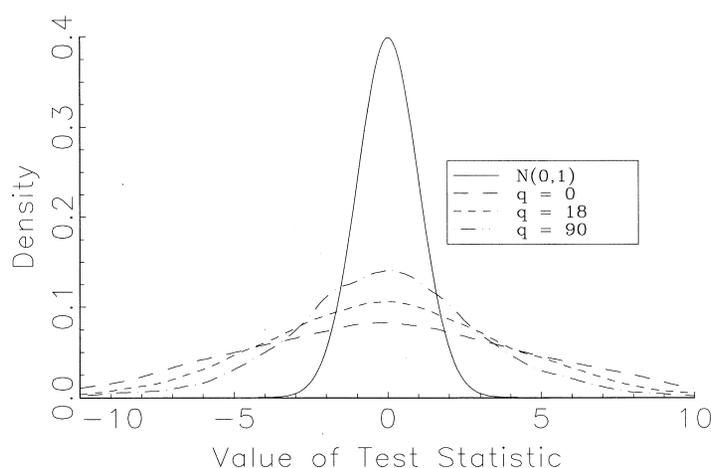


Figure 2. Absolute return volatility tests. The figure displays the distribution of the absolute return volatility test in equation (4). The statistic is calculated from 18 5-minute return observations per day over 20 days before and after. The values of q indicate the order of the augmentation lag in the Newey–West correction. The returns are generated by the high-frequency GARCH(1,1) model described in the text. The solid line labeled $N(0,1)$ gives the standard normal distribution that obtains with i.i.d. normally distributed returns.

Figure 2 shows the simulated distribution of the t -statistic in equation (4) testing the constancy of the 20-day $1\frac{1}{2}$ -hour interval absolute returns generated by the same 5-minute GARCH(1,1) process underlying Figure 1. The distribution denoted by $q = 0$ refers to the test statistic using the sample variance of the absolute returns in place of $A \text{Var}(A^a)$ and $A \text{Var}(A^b)$, while $q = 18$ and $q = 90$ give the distributions using a Newey–West covariance matrix estimator with a lag-length extending for 1-day, or 18 5-minute intervals, and 5-day, or 90 5-minute intervals, respectively. The results are ubiquitous. The strong conditional heteroskedasticity in the high-frequency returns renders standard statistical inference procedures problematic. Even though the test statistics are calculated from 360 observations, the effective sample size is very small, and the statistic with 90 augmentation lags is still far from being standard normally distributed.

These results motivate our development of new robust semiparametric techniques that are applicable in the high-frequency data setting more generally. We now turn to this task.

II. Intraday Volatility Patterns: Estimation and Testing

Statistical inference concerning intraday volatility patterns is severely complicated by nonnormality of the high-frequency returns and strong serial correlation in volatility. The former feature renders sample volatility measures over intraday intervals very noisy, whereas the latter renders the

estimation of the overall volatility level difficult. The FFF regression described in Andersen and Bollerslev (1997a, 1998a, 1998c) provides a simple estimation strategy that readily accommodates these features within a coherent framework, thus allowing for robust inference and reliable hypothesis testing procedures in empirically realistic settings.

A. Fourier Flexible Form Regression

The FFF regression approach is guided by the observations that the systematic intraday volatility patterns generally appear remarkably stable through time and are largely independent of the strong serial correlation in volatility across days. This suggests the following representation of the demeaned high-frequency return process,⁹

$$R_{t,n} = \sigma_{t,n} \cdot s_{t,n} \cdot z_{t,n}, \quad (5)$$

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. Squaring the returns, invoking a robustifying log transformation, and rearranging terms results in the following decomposition:

$$2 \log |R_{t,n}| - \log \sigma_{t,n}^2 = c + 2 \log s_{t,n} + u_{t,n}, \quad (6)$$

where $c = E[\log z_{t,n}^2]$ and $u_{t,n} = \log z_{t,n}^2 - E[\log z_{t,n}^2]$. Equation (6) 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, $s_{t,n}$ captures 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, implementation of the regression suggested by equation (6) requires a specific proxy for the stochastic volatility component, $\sigma_{t,n}$, and a parameterization of $s_{t,n}$ in terms of observables. The first issue is readily addressed. The results of Nelson (1990, 1992) suggest that any reasonable high-frequency GARCH-like model will produce a (continuous-record asymptotics) consistent estimate of the daily volatility factor, say $\hat{\sigma}_t$. Because the 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 a corresponding proportionate intraday measure.

⁹ Aside from negative serial correlation induced by bid-ask bounce effects, most high-frequency return series are effectively mean zero.

The second issue requires more work. Theory provides few guidelines regarding the functional form for the intraday pattern, suggesting the use of a flexible representation for $s_{t,n}$. The FFF introduced by Gallant (1981) is ideally suited for characterizing periodic patterns, and has the advantage of readily incorporating dummy variables to accommodate discontinuities associated with, for example, market openings. Defining the regressand in equation (6) as $Y_{t,n} \equiv 2 \log|R_{t,n}| - \log \hat{\sigma}_t^2 + \log N$, we have the FFF regression

$$Y_{t,n} = f(\boldsymbol{\theta}; t, n) + \hat{u}_{t,n}, \tag{7}$$

where the error process $\{\hat{u}_{t,n}\}$ is stationary, and the explanatory variables on the right-hand-side of equation (7) are determined by the FFF

$$f(\boldsymbol{\theta}; t, n) = \sum_{j=0}^J \mu_j \cdot n^j + \sum_{d=1}^D \lambda_d \cdot I_d(t, n) + \sum_{p=1}^P [\delta_{c,p} \cdot \cos(2\pi pn/N) + \delta_{s,p} \cdot \sin(2\pi pn/N)], \tag{8}$$

and $\boldsymbol{\theta} \equiv (\mu_0, \mu_1, \dots, \mu_J, \lambda_1, \dots, \lambda_D, \delta_{c,1}, \dots, \delta_{c,P}, \delta_{s,1}, \dots, \delta_{s,P})$. The polynomial and the sinusoids in equation (8) capture the overall smooth variation in the intraday pattern, whereas the (zero-one) indicators, $I_d(t, n)$, allow for discontinuities associated with predetermined events in the interval (t, n) .

The robustness properties of the FFF-regression approach are noteworthy. As long as the average intraday fluctuations are correctly specified, which is guaranteed asymptotically by expanding the orders of J and P , and the associated errors, $\hat{u}_{t,n}$, are stationary, the OLS estimates for $\boldsymbol{\theta}$, defined by equations (7) and (8), remain consistent under quite general conditions; a more detailed discussion is provided in the appendix to Andersen and Bollerslev (1998a). Of course, the regression errors are generally conditionally heteroskedastic and serially correlated. However, valid standard errors for the OLS estimates are readily obtained from robust procedures for estimation of the parameter covariance matrix, denoted $\mathbf{A} \mathbf{Var}(\hat{\boldsymbol{\theta}})$, referenced in Section I.

B. Testing for Changes in Intraday Patterns and Shapes

This section develops our formal procedure for general hypothesis testing regarding the shape of the intraday pattern. The approach requires that the relevant hypothesis is quantified in terms of a scalar ordinal measure, say $g(\theta) \equiv F[f(\boldsymbol{\theta}; t, n)]$, where the mapping F operationalizes the feature of interest in terms of the underlying FFF parameters. This is typically straightforward. For instance, if interest centers on the level of volatility during a specific intraday interval, the mapping F may be defined as the area under the curve $f(\boldsymbol{\theta}; t, n)$ over this interval, or more formally, the integral of $f(\boldsymbol{\theta}; t, n)$

with respect to the continuous values of n ranging over the interval. Similarly, the steepness of the pattern at a point in time, or the average steepness during a particular interval, is readily expressed in terms of $\partial f(\boldsymbol{\theta}; t, n)/\partial n$ for a specific value of n , or the integral of $\partial f(\boldsymbol{\theta}; t, n)/\partial n$ for continuous values of n over the relevant interval(s). The flexibility afforded by arbitrary specifications of the mapping F sets it apart from the standard variance-ratio statistics which only measures average volatility over an interval (“box-shaped” patterns). The FFF test may be refined to specific features of the pattern.

Of course, to implement the test, the parameters $\boldsymbol{\theta}$ entering $g(\boldsymbol{\theta})$ must be estimated. However, standard asymptotic theory implies that, when evaluated at the FFF-regression estimate for $\boldsymbol{\theta}$, the function $g(\hat{\boldsymbol{\theta}})$ is asymptotically normal around the true value with an asymptotic variance of

$$A \text{Var}(g(\hat{\boldsymbol{\theta}})) = \mathbf{G}(\hat{\boldsymbol{\theta}}) \cdot \mathbf{A} \text{Var}(\hat{\boldsymbol{\theta}}) \cdot \mathbf{G}(\hat{\boldsymbol{\theta}})', \quad (9)$$

where $G(\boldsymbol{\theta}) \equiv \partial g(\boldsymbol{\theta})/\partial \boldsymbol{\theta}'$. Moreover, because the pre- and post-event parameters, say $\boldsymbol{\theta}^b$ and $\boldsymbol{\theta}^a$, are estimated from nonoverlapping samples, the estimation errors will be asymptotically independent. Thus, an asymptotically normal t -test for the null hypothesis of no structural change, or $H_0: g(\boldsymbol{\theta}^b) = g(\boldsymbol{\theta}^a)$, versus any of the alternatives, $H_A: g(\boldsymbol{\theta}^b) \neq g(\boldsymbol{\theta}^a)$, $g(\boldsymbol{\theta}^b) > g(\boldsymbol{\theta}^a)$, or $g(\boldsymbol{\theta}^b) < g(\boldsymbol{\theta}^a)$, is readily constructed as

$$Z = [g(\hat{\boldsymbol{\theta}}^a) - g(\hat{\boldsymbol{\theta}}^b)]/[A \text{Var}(g(\hat{\boldsymbol{\theta}}^a)) + A \text{Var}(g(\hat{\boldsymbol{\theta}}^b))]^{1/2}. \quad (10)$$

Given the large intraday samples typically available for OLS estimation of $\hat{\boldsymbol{\theta}}^a$ and $\hat{\boldsymbol{\theta}}^b$, and the small number of parameters required in $\boldsymbol{\theta}$, the normal approximation should be fully adequate.

To illustrate, consider the hypothesis of constant unconditional volatility explored by the simulations in Section I. Within the FFF approach, this hypothesis implies equality of the integrals of $f(\hat{\boldsymbol{\theta}}^a; t, n)$ and $f(\hat{\boldsymbol{\theta}}^b; t, n)$ with respect to n , for n ranging over the relevant $1\frac{1}{2}$ hour interval. In keeping with the empirical analysis in Section III, we take $f(\boldsymbol{\theta}; t, n)$ to be a third-order polynomial, with the order of the augmentation lag in the Newey–West correction fixed at one day, or $q = 18$.¹⁰ The long dashed line in Figure 3, labeled “Homoskedastic FFF”, gives the simulated distribution under the five-minute GARCH(1,1) data-generating process, with the volatility factor, $\hat{\sigma}_t^2$, fixed at the constant unconditional daily variance. As for the standard variance-ratio statistic and the test based on high-frequency absolute returns depicted in Figures 1 and 2, respectively, the normal distribution provides a poor approximation to the true sampling distribution of the FFF test

¹⁰ Because we did not include any intraday components in the simulated high-frequency GARCH(1,1) model, all of the estimated higher-order terms in the FFF expansion should be zero asymptotically. However, very similar simulation results were obtained when we included a nontrivial periodic pattern to be estimated.

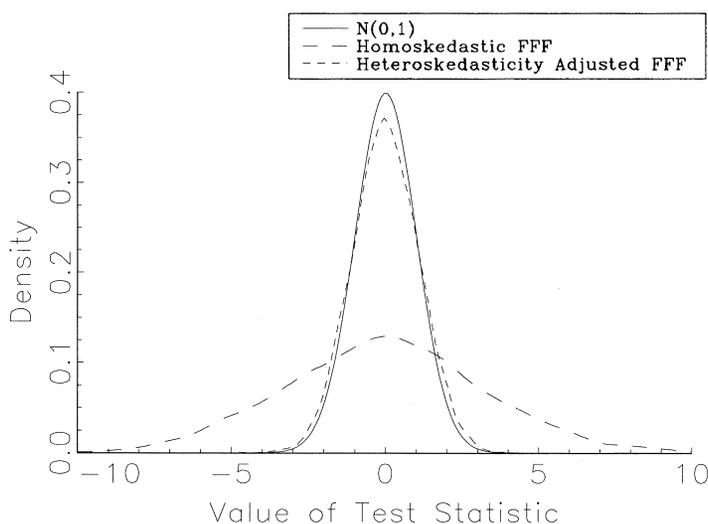


Figure 3. Absolute return volatility tests. The figure displays the distribution of the FFF volatility test in equation (10). The statistic is calculated from 18 5-minute return observations per day over 20 days before and after. The FFF regression is based on a third-order polynomial, along with a lag length of 18, or one day, in the Newey–West correction. The returns are generated by the high-frequency GARCH(1,1) model described in the text. The long dashed line labeled “Homoskedastic FFF” gives the distribution of the test statistic with the daily volatility factor assumed to be constant, and the short dashed line labeled “Heteroskedasticity Adjusted FFF” estimates the daily volatility from the sum of the previous days five-minute return volatilities. The solid line refers to the asymptotic standard normal distribution.

based on a constant σ_t^2 . Meanwhile, on estimating the *daily* volatility factor entering $Y_{t,n}$ by the sum of the one-day lagged five-minute volatilities, that is, $\hat{\sigma}_t^2 = \sum_{n=1, \dots, N} \sigma_{t-1,n}^2$, the close correspondence between the distribution of the FFF test given by the short dashed line labeled “Heteroskedasticity Adjusted FFF” and the normal distribution is striking. It is evident that the robust FFF procedure provides a much more reliable indication of the statistical significance than standard variance-ratio type procedures. The empirical illustration below reinforces this conclusion.

III. The Tokyo Experiment Revisited

The Ito et al. (1998) inference regarding a shift in the intraday yen–dollar volatility pattern around the lifting of Japanese lunch-period trading restrictions is based predominantly on variance ratios. However, Section I shows that such test statistics must be interpreted very carefully. Meanwhile, the FFF approach provides a simple robust procedure for testing for specific changes in the intraday volatility pattern. This section implements the new approach. Beyond revisiting the Ito et al. (1998) evidence, we also complement

their 20-day and 60-day before-versus-after analysis with results based on a much longer two-year sample of high-frequency returns before and after the December 22, 1994, deregulation.

A. Data Description

The FX interbank market is an over-the-counter market with no public reporting of trades or transaction prices, so information about high-frequency features must be obtained by other means. A common source is the indicative quotes supplied by commercial information providers. In particular, many recent studies, including Ito et al. (1998) and the present paper, utilize data from Reuters' multicontributor FAFX screen, captured and distributed by Olsen & Associates. However, whereas Ito et al. (1998) rely on one-minute returns, our analysis is based on five-minute returns. The choice of a longer interval reflects our desire to mitigate the effects of nonsynchronous quotations and spuriously induced autocorrelation; a more detailed discussion of this issue and the general properties of indicative FX quotes is given by Andersen, Bollerslev, and Das (1998), and the exact procedure for construction of the return series is provided by Müller et al. (1990) and Dacorogna et al. (1993).

B. The Estimates of the Intraday Yen–Dollar Volatility Patterns

This section presents FFF-regression estimates of the intraday yen–dollar volatility pattern over the Japanese segment of the FX interbank market for various sample windows before and after the regulatory change. The daily volatility factor, $\hat{\sigma}_t$, is obtained by an MA(1)-GARCH(1,1) model estimated from daily yen–dollar returns covering December 2, 1986 to November 29, 1996. The model provides a good first approximation to the volatility process, and thus should alleviate the influence of the interdaily conditional heteroskedasticity. The sharp discontinuities in the intraday pattern induced by the Tokyo lunch period led us to use a separate FFF specification before, during, and after lunch. Moreover, because the estimation covers only a fraction of the full trading day in the FX market, the sinusoids are less compelling as regressors, and simple third-order polynomials, augmented with Japanese market opening dummies at 9:00 and 1:30 Tokyo time, were deemed sufficient to capture all significant variation within each segment. Thus, the FFF regression takes the form

$$\begin{aligned}
 f(\boldsymbol{\theta}; t, n) = & I_M(n) \cdot \sum_{j=0}^3 \mu_{M,j} \cdot n^j + I_L(n) \cdot \sum_{j=0}^3 \mu_{L,j} \cdot (n - 36)^j \\
 & + I_A(n) \cdot \sum_{j=0}^3 \mu_{A,j} \cdot (n - 54)^j + \sum_{d=1}^2 \lambda_d \cdot I_d(t, n), \quad (11)
 \end{aligned}$$

where $\mu_{M,j}$ denotes the polynomial coefficients for the morning pattern, 9:00–12:00 ($n \in \{1, \dots, 36\}$), $\mu_{L,j}$ are the lunch coefficients of relevance for 12:00–1:30 ($n \in \{37, \dots, 54\}$), $\mu_{A,j}$ are afternoon coefficients relating to the 1:30–3:30 period ($n \in \{55, \dots, 78\}$), and $I_M(n)$, $I_L(n)$, and $I_A(n)$ are indicator variables equaling unity if the interval, n , belongs to the morning, lunch, or afternoon, respectively, and zero otherwise. Finally, the $I_d(t, n)$ ($d \in \{1, 2\}$), market opening indicators are zero everywhere except for the 9:00–9:05 and 1:30–1:35 intervals, respectively, where they equal unity.

Because the individual FFF-coefficients are void of economic interpretation,¹¹ we convey the results through the estimated patterns in the log-absolute return dimension, as defined by the $Y_{L,n}$ regressor in equation (7). Although the two panels for the 20-day windows in Figure 4 appear somewhat jagged, the overall shape is quite transparent in the 60-day graphs, and the longer two-year samples display remarkable stability in the volatility pattern outside of the lunch period. The ability of third-order polynomials to adequately characterize the smooth intraday patterns is transparent.

C. Testing for Changes in the Yen–Dollar Intraday Volatility Patterns

If volatility is caused solely by public information and the public information flow is unaltered, volatility should not be affected by the introduction of lunch-period trading. On the other hand, the presence of private information or mispricing would increase volatility over lunch following the deregulation. The first variance-ratio test reported by Ito et al. (1998) soundly rejects equality of the lunch period volatilities before and after the change in regime. However, as discussed above, the conventional critical values from the F -distribution can be very misleading in this context. Instead the open-versus-closed volatilities may be assessed more reliably by the area under the estimated intraday volatility curves during the Tokyo lunch period, as defined by the integral of $f(\theta; t, n)$ over the relevant range of n .¹² Consistent with the visual impression from Figure 4, the results reported in Table I imply strong rejections of the null hypothesis of no change versus the alternative of higher volatility following the relaxation of the Tokyo trading restrictions.¹³ Thus, the results from the FFF regression are actually much stronger than indicated by the p -value of the 20-day variance-ratio statistic

¹¹ 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–West covariance matrix estimator and a lag length of 78, or 1 Japanese trading day, for the 20-day and 60-day estimates, and 390, or 5 Japanese trading days, for the two-year estimates.

¹² For the FFF regression in equation (11), 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. This is the same integral that underlies the simulations in Section II.B. Similar expressions for the remaining hypothesis tests described below are available from the authors on request.

¹³ Consistent with the notation in equation (10), all of the Z -tests in Table I are based on the “open” minus “close” (or “after” minus “before”) values for the function $g(\theta)$; that is, $g(\theta^a) - g(\theta^b)$. The reported p -values also reflect this convention.

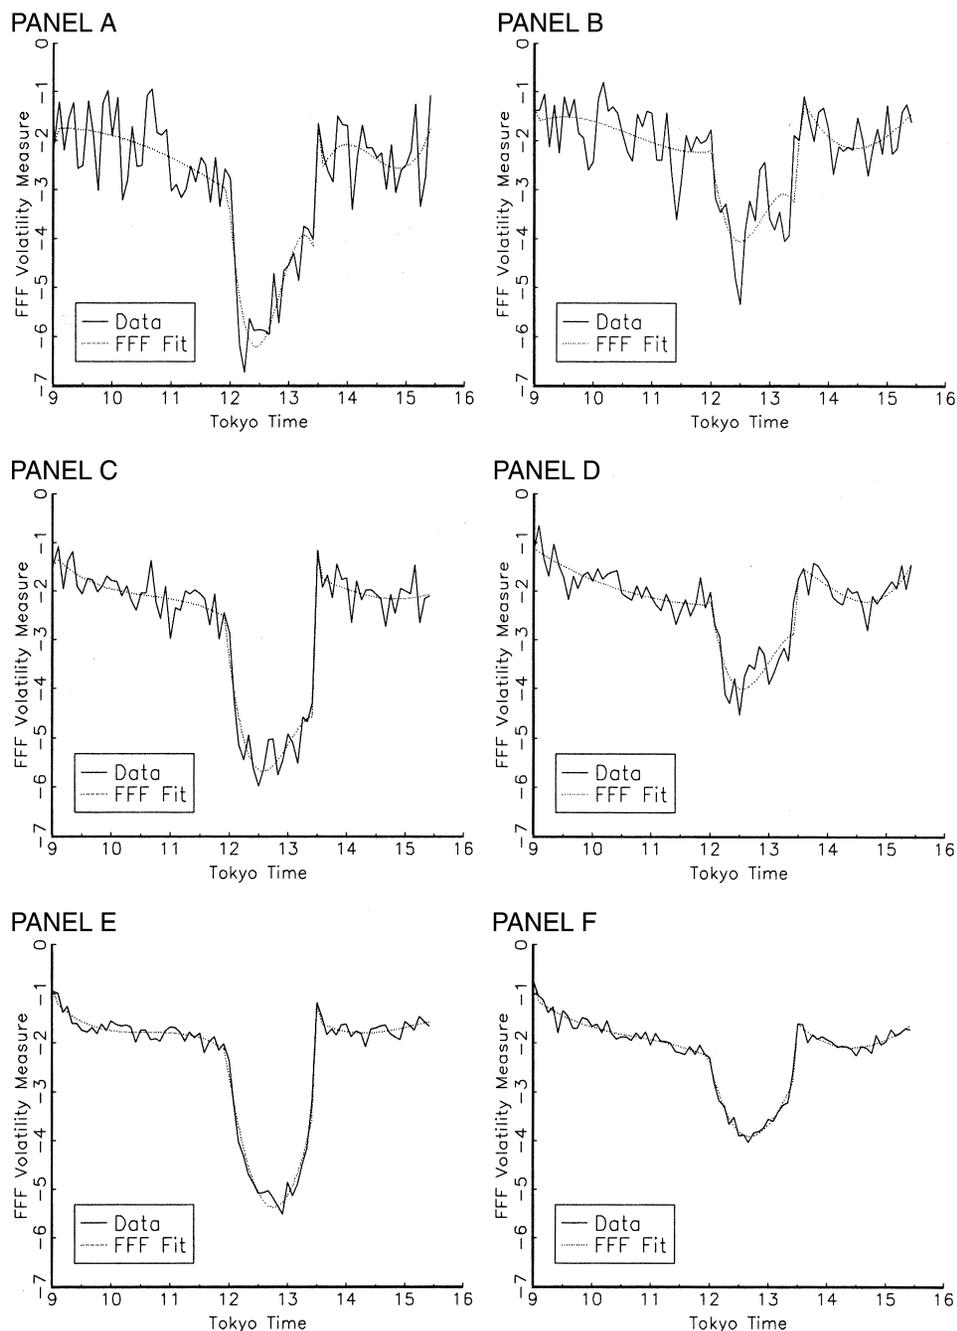


Figure 4. FFF intraday volatility estimates. The figure compares the average log-absolute yen-dollar return as defined by the left-hand side of equation (7) to the corresponding FFF fit from equation (11) for samples covering 20 days, 60 days and 2 years before December 22, 1994, and after January 4, 1995.

Table I

Tests for Changes in the Intraday Yen–Dollar Volatility Pattern

The table reports the FFF-based tests for the null hypothesis of no change in the intraday yen–dollar volatility pattern following the lifting of the trading restrictions on December 22, 1994, as defined by the Z -statistic in equation (10) and the FFF regression in equation (11). The tests are calculated from the nonweekend five-minute yen–dollar 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 tests for an upward tilt in the afternoon volatility U shape.

Hypothesis	Horizon	Z -statistic	p -value
H_1	2 years	8.00	0.000
	60 days	4.65	0.000
	20 days	2.57	0.005
H_2	2 years	0.16	0.564
	60 days	0.34	0.633
	20 days	-1.28	0.100
H_3	2 years	-0.36	0.359
	60 days	0.16	0.564
	20 days	0.20	0.579
H_4	2 years	0.12	0.452
	60 days	-0.55	0.709
	20 days	1.20	0.115

when evaluated in the simulated GARCH(1,1) distribution in Section I (cf., footnote 7). This reflects the much improved test power obtained by the FFF approach vis-à-vis the properly size-corrected variance-ratio statistics explored in Section I. As such, this result represents a statistically robust confirmation of the prior evidence on relative lunch period volatility in Ito et al. (1998).

Unfortunately, the evidence in favor of the remaining Ito et al. (1998) hypotheses does not survive our more robust testing methodology. Whereas the variance-ratio statistics in Ito et al. (1998) all rely on “box-like” comparisons, the FFF approach allows us to target specific characteristics of the volatility pattern that are more in line with the underlying theoretical considerations. Hence, our remaining tests are not exactly equivalent to those of Ito et al. (1998), but should be superior in terms of power, as illustrated by the lunch volatility hypothesis H_1 above.

If the amount of private information is unaffected by the deregulation, Ito et al. (1998) argue that trading over lunch should result in a more smooth incorporation of information throughout the day. This suggests a flattening of the typical U-shape in the intraday volatility. We capture the U-shape by the average slope of the second part of the curve, say S_A , versus the average slope of the first part, S_M . The difference, defined as $F \equiv S_A - S_M = g(\theta)$,

provides a natural measure of curvature.¹⁴ 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. Whereas the variance-ratio statistics reported by Ito et al. (1998) seemingly provide overwhelming support for a flattening of the U-shape, none of the robust H_2 tests for the difference between the before and after slopes are significant. In fact, for the 60-day and 2-year horizons, the Z -statistics are in the wrong tail of the distribution.¹⁵

When trading is restricted during lunch 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. Letting S_{EM} and S_{LM} denote the average slopes of the early morning (9:00–10:30) and late morning (10:30–12:00) curves, respectively, and defining $F_M \equiv S_{LM} - S_{EM}$, these statistics may again 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 a test for equality of F_M in the before and after samples. Ito et al. (1998) strongly reject this hypothesis, finding it “. . . perhaps the most compelling single fact of the four . . .” in favor of the private information hypothesis. However, for none of the three horizons are the robust tests for H_3 in Table I close to significant. Visual inspection of Figure 4 does not indicate any structural break in the average volatility pattern over the Tokyo morning period either.¹⁶

The final Ito et al. (1998) test concerns the apparent increase in volatility towards the end of the Japanese trading day. If information is short-lived, and trading is restricted over lunch, 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. Because we already rejected the notion that the volatility pattern changed during the

¹⁴ To be precise, this test is constructed differently from the corresponding test of the H_2 hypothesis in Ito et al. (1998). They operationalize the stipulated change in the U shape as a test for an increase in the volatility over lunch relative to that over either the morning or the afternoon. It is evident from Figure 4 that this test is endowed with strong power to reject the null hypothesis of no change because of the increase over lunch, even if the volatility pattern outside of lunch is unchanged. Consequently, it will, by construction, almost surely reject the null hypothesis. Because the Ito et al. (1998) theory emphasizes the shift in the pattern over the entire trading day, our test for H_2 instead focuses on the asserted shift in the U-shape pattern outside of the lunch period.

¹⁵ These findings are consistent with the graphical displays in Figure 4. Ito et al. (1998) present a graph that conveys an entirely different impression. However, Figure A1.A discussed in the appendix shows that their nonrobust variance-ratio estimators, and corresponding figures, are distorted by a *single* questionable indicative quote.

¹⁶ Ito et al. (1998) present somewhat different specifications of this hypothesis, but convey the evidence solely in terms of a figure. However, Figure A1.B in the appendix again shows that the Ito et al. (1998) volatility pattern based on one-minute returns is completely transformed if one eliminates a *single* questionable quote from their sample.

morning period, we focus on a change in the afternoon pattern, as measured by the average slope of the FFF volatility curve after lunch, or S_A in the notation above. None of the H_4 tests in Table I suggest that such an upward tilt occurred. Again, these results are squarely at odds with the apparent strong statistical evidence provided by the variance-ratio tests in Ito et al. (1998).

D. Further Discussion of the Ito et al. (1998) Findings

The main objective of Ito et al. (1998) is to make a case for the presence of private information in the FX interbank market. How do we then interpret our findings of an increase in volatility over lunch, but no important changes outside of the lunch period? Obviously, we cannot exclude the possibility that the impact of the deregulation is insufficient to alter the trading strategies of informed agents outside of the lunch period, even in the presence of asymmetric information so, by construction, our findings cannot refute the existence of private information per se.

Motivated by the analysis in French and Roll (1986), Ito et al. (1998) offer an additional test for the presence of private information by considering a bound on the fraction of the return variance during the lunch period that may arise from temporary pricing errors. The bound is given by the ratio of the lunch period variance implied by the 90-minute price change over the full lunch period relative to that implied by the corresponding one-minute returns. Ito et al. (1998) report that for the 60-day windows, this bound drops following the deregulation, suggesting that private information contributes to the postperiod increase in volatility. Unfortunately, this result is not robust either. As detailed in Table AII in the appendix, the ordering of the bounds is reversed if one considers the 20-day or 40-day windows. The decrease in the bound for the 60-day window is thus strictly due to observations that are between 40 to 60 days removed from the time of the deregulation. The fact that the indicative quote intensity over lunch rises dramatically following the regulatory change further complicates these comparisons. Given the importance of spurious negative autocorrelation in the one-minute returns induced by artificial bouncing in the indicative quotes, an increase in quote frequency and an associated rise in the fraction of nonzero returns will artificially inflate measured volatility; for additional discussion along these lines we refer to Andersen, Bollerslev, and Das (1998).

In sum, given the lack of direct evidence in favor of any change in the volatility pattern outside of the Tokyo lunch period, we conclude that the Tokyo Experiment provides little useful evidence on this general debate.

IV. Concluding Remarks

The empirical market microstructure literature has documented pronounced intraday patterns in a number of key financial variables. These striking regularities have inspired alternative theoretical developments that

seek to shed light on a variety of issues, including the intrinsic differences between equity, bond and foreign exchange markets, the optimal design of market mechanisms, the possibility of disentangling short-lived price effects from more permanent shocks, as well as the actual trading strategies of market participants.

Meanwhile, much of the pertinent empirical evidence and statistical tests rely on standard variance-ratio procedures. Such tests can be seriously misleading in the high-frequency data context. The present paper develops new tools for robust empirical analysis of intraday volatility patterns that are capable of handling the unique complications that arise in this situation. We illustrate the importance of the issue by revisiting the Tokyo Experiment. We confirm that volatility increases over lunch following the deregulation, but find no support for qualitative shifts in the volatility pattern outside of the lunch period. Although this evidence runs counter to the majority of the findings in Ito et al. (1998), we reiterate that our results cannot reject the relevance of private information for the spot FX interbank market per se. The regulatory change may simply not be significant enough to induce a detectable shift in the pattern outside of lunch. At a more general level, our analysis highlights the importance of robust inference procedures in the high-frequency data setting. As such, the methods developed here should be useful for the investigation of a range of market microstructure issues with direct and testable implications for the shape of intraday volatility patterns.

Appendix

This appendix highlights some of the problems with the one-minute return series and the standard variance-ratio methodology employed by Ito et al. (1998) that are relevant for judging the new FFF methodology and the empirical results presented in Section III.

A. On the Influence of Indicative Quote Errors

The Ito et al. (1998) indicative quote sample contains an influential outlier on December 1, 1994. Table A1 provides an excerpt of the data file around this observation, recorded at 02:26:56 GMT, or 11:26:56 Tokyo time. The bid-ask spread is one full yen, or 10-fold the typical quoted spread. Moreover, it is evident that the ask accounts for this phenomenon, as the bid is in line with the surrounding quotes, and the following ask quote drops back in line with the general level of the yen-dollar rate at the time. In fact, the identical one-yen spread is also recorded at 2:31:10 GMT, but at that time the error code indicates a potential data problem. Still, even if the 02:26:26 quote is accurate, it is unlikely to reflect underlying transaction prices, and it is thus economically irrelevant.

Unfortunately, the Ito et al. (1998) inference is highly sensitive to such outliers. Figure A1 reproduces Figures 1 and 2 in Ito et al. (1998), using their exact approach, including an additional “before” graph obtained from

Table AI

An Excerpt from Reuters FXXF Quote Series on December 1, 1994

The table provides an excerpt from the FXXF series of indicative yen-dollar quotes on Thursday, December 1, 1994. The Greenwich Mean Time (GMT) period covered translates into 11:16:48–11:43:16 Tokyo time. The bid and ask quotes are given in number of yen per U.S. dollar. The data have been run through an automated filter designed by Olsen & Associates to assist users in the identification of questionable quotes. It generates the code 0 if a quote does not pass certain criteria, indicating a potential problem, whereas the code is 1 otherwise. The “code 1” quote that we identify as a questionable outlier is given in bold type.

Date	Time (GMT)	Bid	Ask	Code
1994-12-01	02:16:48	98.90	98.95	1
1994-12-01	02:19:56	98.88	98.98	1
1994-12-01	02:26:56	98.93	99.93	1
1994-12-01	02:27:06	98.93	98.94	0
1994-12-01	02:30:00	98.87	98.97	1
1994-12-01	02:30:10	98.92	99.93	0
1994-12-01	02:30:54	98.86	98.96	1
1994-12-01	02:31:10	98.93	99.93	0
1994-12-01	02:31:36	98.87	98.97	1
1994-12-01	02:41:34	98.92	98.97	1
1994-12-01	02:43:16	98.88	98.98	1

their one-minute series modified only by excluding the *one* questionable quote. The impact is dramatic. Panel A now simply suggests a parallel shift in the volatility level between the before and after samples. There is no sign of a qualitative shift in the pattern. Panel B conveys a similar message. The morning pattern shifts upward following the deregulation, rather than changing shape. The overall daily volatility level is simply higher in the ex post sample. This rather dramatic influence of a *single* moderate outlier underscores the lack of robustness of standard variance-ratio procedures in the high-frequency context.

B. On the Robustness of the French-Roll Variance Bounds

Ito et al. (1998) also provide so-called French-Roll bounds for the fraction of variance induced by pricing errors. Specifically, they argue that a sharp drop in the bound between the 60-day before and after sample signals the presence of private information. Our previous findings suggest that such variance ratios are estimated very imprecisely from a short 60-day sample. Table AII confirms this conjecture. The corresponding bounds for the 20- and 40-day samples are squarely at odds with the 60-day result. For both these samples, the bound increases following the deregulation. Our five-minute series provide similar findings. The drop in the bound across the 60-day samples is thus strictly due to observations that are 40 to 60 days removed from the event.

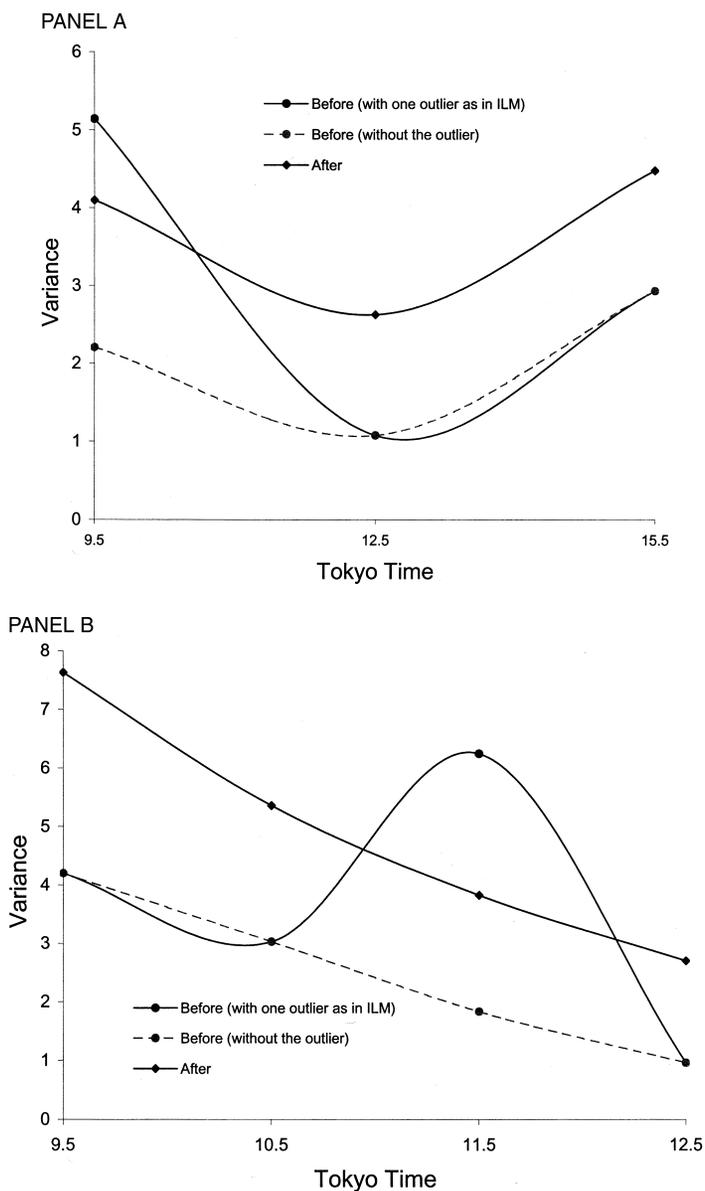


Figure A1. Ito et al. (1998) volatility patterns. The figure replicates the intraday volatility patterns of Ito et al. (1998) and compares them to the corresponding graphs obtained by deleting a *single* questionable indicative quote observed during late morning trading in Tokyo on Thursday, December 1, 1994. The computation relies on one-minute yen-dollar weekday returns over 20 days before December 22, 1994, and after January 4, 1995. The morning, lunch, and afternoon points in Panel A reflect sample variances over 10:30–12:00, 12:00–1:30, and 1:30–3:00 Tokyo time, respectively. In both panels, the lines are smoothed interpolations between the variance estimates. By construction, all but one of the point estimates for the variance on the before graphs coincide. The units of the variance are log changes of the bid-offer midpoint per minute (multiplied by 10^8).

Table AII
French-Roll Variance Bounds

The table reports the value of $1 - V_l/V_s$ estimated using different samples of one-minute and five-minute yen-dollar returns. V_l denotes the sample return variance over the Tokyo lunch period, 12:00–1:30 local time, computed from the log-price change over the entire lunch period, and V_s refers to the lunch period sample variance calculated from the high-frequency returns. The resulting statistic may be interpreted as an upper bound on the fraction of variance arising from mispricing. The samples consist of continuously compounded returns for the 20, 40, or 60 weekdays before December 22, 1994, and after January 4, 1995.

Sample	One-minute Returns	Five-minute Returns
20 Days		
Before	0.62	0.44
After	0.68	0.43
40 Days		
Before	0.72	0.49
After	0.78	0.64
60 Days		
Before	0.80	0.61
After	0.38	0.14

REFERENCES

- Admati, Anat R., and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.
- Amihud, Yakov, and Haim Mendelson, 1987, Trading mechanisms and stock returns: An empirical investigation, *Journal of Finance* 42, 533–553.
- Amihud, Yakov, and Haim Mendelson, 1991, Volatility, efficiency and trading: Evidence from the Japanese stock market, *Journal of Finance* 46, 1765–1789.
- Andersen, Torben G., and Tim Bollerslev, 1997a, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.
- Andersen, Torben G., and Tim Bollerslev, 1997b, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns, *Journal of Finance* 52, 975–1005.
- Andersen, Torben G., and Tim Bollerslev, 1998a, Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer-run dependencies, *Journal of Finance* 53, 219–265.
- Andersen, Torben G., and Tim Bollerslev, 1998b, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- Andersen, Torben G., and Tim Bollerslev, 1998c, Towards a unified framework for high- and low-frequency return volatility modeling, *Statistica Neerlandica* 52, 273–302.
- Andersen, Torben G., Tim Bollerslev, and Ashish Das, 1998, Testing for microstructure effects in volatility: Revisiting the Tokyo experiment, Working paper, Northwestern University.
- Andrews, Donald W. K., 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.
- Ball, Ray, and Philip Brown, 1968, An empirical evaluation of accounting income numbers, *Journal of Accounting Research* 6, 159–178.
- Barclay, Michael J., Robert H. Litzenberger, and Jerold B. Warner, 1990, Private information, trading volume, and stock return variances, *Review of Financial Studies* 3, 233–253.

- Bollerslev, Tim, Robert F. Engle, and Daniel B. Nelson, 1994, ARCH models, in Robert F. Engle and Daniel McFadden, eds.: *Handbook of Econometrics*, Vol. IV, (North Holland Press, Amsterdam).
- Covrig, Vincentiu, and Michael T. Melvin, 1998, Asymmetric information and price discovery in the fx market: Does Tokyo know more about the yen? Working paper, Arizona State University.
- Dacorogna, Michel M., Ulrich A. Müller, Robert J. Nagler, Richard B. Olsen, and Olivier V. Pictet, 1993, A geographical model for the daily and weekly seasonal volatility in the foreign exchange market, *Journal of International Money and Finance* 12, 413–438.
- Drost, Feike C., and Theo E. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 909–927.
- Ederington, Louis H., and Jae H. Lee, 1993, How markets process information: News releases and volatility, *Journal of Finance* 48, 1161–1191.
- Fama, Eugene, Lawrence Fisher, Michael C. Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information, *International Economic Review* 10, 1–21.
- Forster, Margaret M., and Thomas J. George, 1996, Pricing errors at the NYSE open and close: Evidence from international cross-listed stocks, *Journal of Financial Intermediation* 5, 95–126.
- Foster, F. Douglas, 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.
- French, Kenneth R., and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.
- Gallant, A. Ronald, 1981, On the bias in flexible functional forms and an essentially unbiased form: The Fourier flexible form, *Journal of Econometrics* 15, 211–245.
- Harris, Lawrence, 1986, A transaction data study of weekly and intraday patterns in stock returns, *Journal of Financial Economics* 16, 99–117.
- Harvey, Campbell, and Roger D. Huang, 1991, Volatility in foreign currency futures markets, *Review of Financial Studies* 4, 543–569.
- Hau, Harald, 1998, Competitive entry and endogenous risk in the foreign exchange market, *Review of Financial Studies* 11, 757–787.
- Ho, Richard Y., and Raymond S. Lee, 1998, Market closure effects on return, volatility and turnover patterns in the Hong Kong index futures market, *Journal of International Financial Markets, Institutions and Money* 8, 433–451.
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1998, Is there private information in the fx market? The Tokyo experiment, *Journal of Finance* 53, 1111–1130.
- Johnson, Normal L., and Samuel Kotz, 1970, *Continuous Univariate Distributions-2* (John Wiley and Sons, New York).
- Jones, Charles M., and Gautam Kaul, 1994, On the use of variance ratios, Working paper, University of Michigan.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson, 1994, Information, trading and volatility, *Journal of Financial Economics* 36, 127–153.
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41–66.
- Lo, Andrew W., and A. Craig 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, Ulrich A., Michel M. Dacorogna, Rakhil D. Davé, Richard B. Olsen, Olivier V. Pictet, and Jakob E. von Weizsäcker, 1997, Volatilities at different time resolutions—Analyzing the dynamics of market components, *Journal of Empirical Finance* 4, 213–239.
- Müller, Ulrich A., Michel M. Dacorogna, Richard B. Olsen, Olivier V. Pictet, Matthias Schwarz, and Claude Morgenegg, 1990, Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis, *Journal of Banking and Finance* 14, 1189–1208.
- Nelson, Daniel B., 1990, ARCH models as diffusion approximations, *Journal of Econometrics* 45, 7–38.
- Nelson, Daniel 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, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity consistent covariance matrix, *Econometrica* 55, 703–708.
- Oldfield, George S. Jr., and Richard J. Rogalski, 1980, A theory of common stock returns over trading and nontrading periods, *Journal of Finance* 35, 729–751.
- Payne, Richard, 1999, Informed trade in the spot foreign exchange markets: An empirical investigation, Working paper, London School of Economics.
- 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, James M., and Lawrence H. Summers, 1988, Mean reversion in stock prices, evidence and implications, *Journal of Financial Economics* 22, 27–58.
- Richardson, Matthew, and Tom Smith, 1991, Tests of financial models in the presence of overlapping observations, *Review of Financial Studies* 4, 227–254.
- Ronen, Tavy, 1997, Tests and properties of variance ratios in microstructure studies, *Journal of Financial and Quantitative Analysis* 32, 183–204.
- Sapp, Stephen G., 1998, Price leadership in the spot foreign exchange market, Working paper, Northwestern University.
- Smith, Tom, 1994, Econometrics of financial models and market microstructure effects, *Journal of Financial and Quantitative Analysis* 29, 519.
- Stoll, Hans R., and Robert E. Whaley, 1990, Stock market structure and volatility, *Review of Financial Studies* 3, 37–71.
- Wei, Shang J., and Jungshik Kim, 1997, The big players in the foreign exchange market: Do they trade on information or noise? Working paper, National Bureau of Economic Research.
- Wood, Robert A., Thomas H. McInish, and J. Keith Ord, 1985, An investigation of transaction data for NYSE stocks, *Journal of Finance* 25, 723–739.

