

Is There Private Information in the FX Market?

The Tokyo Experiment

Takatoshi Ito*

Richard K. Lyons

Michael T. Melvin

This draft: January 1996

First draft: October 1995

Abstract

It is a common view that private information in the foreign exchange market does not exist. This paper provides evidence against this view. The evidence comes from the opening of lunch-hour trading in the Tokyo market: in December 1994 the restriction on interdealer trading between 12:00 and 1:30 was abolished. We find that lunch-hour volatility is higher when the market is trading, implying the presence of private information (since the flow of public information did not change with the trading rules). In addition, the shift in regime is instructive because it induces interpretable changes in the U-shape of volatility over the trading day. We find that opening lunch-hour trade changes the volatility U-shape in three ways: the curve is (1) flatter, (2) tilted upward, and (3) higher on average. These effects on the U-shape have implications for the nature of the private information involved.

JEL: F31

Keywords: Volatility, Information, Microstructure.

Correspondence

Professor Richard K. Lyons

Haas School of Business, UC Berkeley

Berkeley, CA 94720-1900

Tel: 510-642-1059, Fax: 510-643-1420

E-mail: lyons@haas.berkeley.edu

* Respective affiliations are Hitotsubashi University, IMF, and NBER; UC Berkeley and NBER; and Arizona State University. We thank Olsen & Associates for generously providing data. Lyons thanks the National Science Foundation and the Berkeley Program in Finance for financial assistance.

Is There Private Information in the FX Market?

The Tokyo Experiment

1. Introduction

It is a common view that private information in the foreign exchange market does not exist. This paper evaluates this view by examining how trading affects volatility. In particular, we examine the Tokyo foreign exchange (FX) market, which until recently was restricted from trading over the lunch break from 12:00 to 1:30. For that ninety minutes, the Tokyo interbank market shut down. On December 22, 1994, this restriction was abolished.¹ The removal of the restriction thus provides a natural experiment regarding the effects of trading on volatility.

It is well known that the variance of equity returns is more than ten times larger when markets are open than when closed (see Fama (1965), Oldfield and Rogalski (1980), and French and Roll (1986), among others). French and Roll posit three possible sources: (1) public information that arrives during business hours, (2) private information that enters price through trading, and (3) pricing errors that occur during trading. They examine stock market closures during which there is no change in the flow of public information to determine the importance of the public information source. They find that return volatility decreases significantly during these closures, which is inconsistent with public information being the sole source of volatility. Rather, private information appears to be the main source of high trading-time volatility (they find only a small role for pricing errors).

The analysis of this paper has two stages. The first stage is similar in spirit to the equity analysis of French and Roll (1986) in that we compare volatility within a given interval across regimes. Our analysis differs, however, in an important way: we are addressing a market whose information structure is more skewed toward public information. In fact, it is a common view that all participants in the FX market are projecting on the same public information set (except, perhaps, at times of central bank intervention). Nevertheless, we find results similar to French

¹ This trading hour restriction was imposed by The Committee of Tokyo Foreign Exchange Market Customs, which is composed of representatives from commercial banks, foreign banks, and FX brokers. It was thus a voluntary regulation by market participants rather than Ministry of Finance guidance. The restriction was introduced in 1972 after the yen exited the Bretton Woods system. The reason the restriction was abolished—according to news reports—was to regain the substantial volume that had migrated to other, unrestricted locations (e.g., Singapore and Hong Kong). The decision to lift the restriction was made on December 21, and implemented on the following day. Restrictions on trading before 9 a.m. and after 3:30 p.m. were also abolished on December 22.

and Roll's: lunch-hour volatility is twice as high with trading, without any change in the release of public information. Given the perceived predominance of public information, this result is striking.²

The second stage of our analysis moves beyond comparing lunch-hour volatility across regimes. In the second stage, we examine the regime shift's effect on the pattern of volatility over the trading day. This allows us to link our analysis to the literature on the U-shape of intraday volatility. That work has established that volatility over the day is U-shaped in many markets.³ And recent theory in the area can rationalize the U-shape, most of it relying on some type of private information.⁴ Empirical work, however, has only recently moved from documenting the phenomenon to determining if private information is indeed involved. Here we address the link between the U-shape and information by comparing the U-shape before and after the introduction of lunch-hour trading. For example, most models relying on private information predict a flattening of the U-shape: allowing trading induces greater revelation during the lunch hour, with a smaller share remaining for the morning and afternoon.

For concreteness, we offer a working definition of the term private information. By private information we mean any non-public information that is price relevant. Non-public information is any information that is not common knowledge in the usual sense. As for price relevance, we stress that price relevant information is not necessarily fundamental. To understand this point, it is helpful to establish a taxonomy of information types that qualify as private by our definition. For the taxonomy, consider a standard two-period, three date trading model in which trading occurs at $t=0$ and $t=1$, and at $t=2$ the fundamental value F is realized (e.g., final payoffs occur). Now, by non-public fundamental information we have in mind signals of the final payoff F . We offer two possibilities in the context of FX. The first involves (unsterilized) intervention: the dealer who receives a central bank order has information that other market participants do not (see Peiers (1995)). The second involves aggregation of information in orders tied to trade in goods: if a country's trade balance is a component of fundamental value, then dealers receive signals of this component long before published statistics are available (see Lyons (1995b)).⁵ In contrast, by price relevant private information that is not fundamental we have in

² The fact that many markets exhibit relatively low volatility over the lunch period does not weaken this result since we are measuring the *change* in volatility from the change in regime; we are not making a statement about the absolute level, so controlling for a "natural" decline over the lunch period is not necessary.

³ Evidence of the U-shape pattern of intraday volatility appears in Wood, McInish, and Ord (1985), Harris (1986), and Andersen and Bollerslev (1994), among many others.

⁴ See Admati and Pfleiderer (1988), Foster and Viswanathan (1990), and Hong and Wang (1995), among others.

⁵ A refinement that is relevant to both of these types of order-flow information lies in its distribution, for example its distribution across dealers. Since the transparency of FX trades is low, order flow remains private information even after being conveyed to counterparties through trading (see, e.g., Lyons (1995b) and Romer (1993)).

mind information that is unrelated to F but is relevant to the interim price at $t=1$. We offer two possibilities in the context of FX. The first involves digestion effects that arise when risk sharing is incomplete: for example, dealer positions uncorrelated with F will be price relevant if it is known that positions are closed at day's end and that these (net) trades affect price (see Lyons (1995a)). The second involves the presence of differential prior beliefs in FX: Kraus and Smith (1994) show that beliefs about beliefs affect $t=1$ prices even when it is common knowledge that no participant has private information about final payoffs at $t=2$.

The two bodies of work most closely related to this paper are those concerning (i) volatility dynamics in the FX market and (ii) volatility effects of opening/closing in other markets. The first emphasizes volatility dynamics across both trading centers and time.⁶ A central result is that volatility generally propagates from one trading center to the next (e.g., a volatile Tokyo market is typically followed by a volatile New York market — the so-called “meteor shower” effect). The seminal work in the second area on opening/closing effects is the French and Roll (1986) paper on the closing of the NYSE. One later paper in this second area that is particularly relevant is Ito and Lin (1992). That paper examines the effects on stock price volatility of the lunch closing of the Tokyo Stock Exchange. (Note that there is no change in the lunch break in that case; the authors address the effects of its presence.) Surprisingly, they find no evidence of an intraday U-shape in volatility. This is surprising because a U-shape is evident on the NYSE, but the NYSE does not break for lunch.

The rest of the paper is divided into four sections. Section 2 summarizes the relevant theory and information specifications. Section 3 describes the data. Section 4 presents our results. And Section 5 concludes.

⁶ See, e.g., Engle, Ito, and Lin (1990), Baillie and Bollerslev (1991), Ito, Engle, and Lin (1992), Dacorogna et al. (1993), and Hogan and Melvin (1994).

2. Experiment Design

2.a. Stage 1: Testing for the presence of private information

The first stage of our analysis is in the spirit of French and Roll (1986). In this first stage we test whether FX volatility is caused solely by public information. To do this, we test whether lunch-hour volatility is affected by the introduction of trading. If there is no change in the flow of public information at the time the restriction is removed, then volatility should not be affected as long as volatility is caused solely by public information. The alternative is that this market is not driven purely by public information, in which case volatility should increase with the introduction of trading.

Clearly, an important concern in interpreting the results is whether the flow of public information changed across the two regimes. If it does, and if the change over the lunch-hour is an increase, then an increase in volatility might be spuriously attributed to private information. We offer three arguments that the flow of public information did not change with the regime shift. First, the decision to abolish the restrictions on trading was not part of a broader policy-reform package; rather, it was an isolated change in regime. Second, the decision was not the work of the Ministry of Finance (MOF). It came from The Committee of Tokyo Foreign Exchange Market Customs (see footnote 1). It is thus difficult to argue that it might have correlated with MOF initiatives more broadly. Finally, we have reviewed the schedule for public release of relevant macro data and find no changes within our sample either toward or away from the midday period.⁷

To summarize our first hypothesis, its test, and the alternative, we have:

H.1: Volatility in the FX market is caused solely by public information.

Test: $V_t^Y(\text{closed}) = V_t^Y(\text{open})$

AH.1: $V_t^Y(\text{closed}) < V_t^Y(\text{open})$

where V_t^Y is the return variance of the yen/dollar rate over the lunch period from 12:00-1:30, Tokyo time. The closed sample corresponds to the period before removal of the trading

⁷ For example, we checked the number of news reports over the 12-1:30 period on the Reuters Money Market Headline News screen. In the thirty trading days before the lunch-hour opening, there were 18.5 reports on average. In the thirty trading days after the lunch-hour opening, there were 17.5 reports on average (precise dates correspond to our sampling convention, described below).

restriction (i.e., before December 22, 1994). The open sample corresponds to the period after removal of the restriction.

Our second stage-one hypothesis is directed at the fact that the Tokyo experiment is not a pure regime shift (closed to open) since there were alternative markets where Tokyo participants could trade during the restricted ninety minutes. Singapore, in particular, is open during the Tokyo trading day, and can execute FX transactions originating in Tokyo over the lunch break.⁸ Nevertheless, under the restriction overall depth in yen/dollar was sharply lower over the break. If depth went to zero over the break, our test would be more powerful. That depth does not go to zero does not, however, bias our test.

There is a way to measure this power loss, and at the same time check the robustness of our test of hypothesis 1. Note that in Tokyo yen/dollar trading is much higher relative to mark/dollar trading than in Singapore or Hong Kong (data provided below). This relative importance provides a way to measure the importance of Singapore and Hong Kong as alternative trading venues. Per above, the power of our test depends on the degree to which depth drops from the Tokyo break. Taking Tokyo off line causes a much sharper drop in the depth of yen/dollar than in the depth of mark/dollar (since Tokyo trading is far more yen/dollar intensive). Though we expect the effect on volatility in these two markets to have the same sign, the magnitude should be greater in the case of yen/dollar.

To summarize our second hypothesis and test, we have:

H.2: If volatility is induced from trading, then mark/dollar volatility over the lunch-hour should increase when the Tokyo market opens, but by less than the increase in yen/dollar volatility.

Test:

$$V_i^M(\text{closed}) < V_i^M(\text{open})$$

and

$$V_i^M(\text{closed})/V_i^M(\text{open}) > V_i^Y(\text{closed})/V_i^Y(\text{open})$$

where V_i^M is the return variance of the mark/dollar rate over the lunch period from 12:00-1:30, Tokyo time.

⁸ Japanese newspapers report that there was indeed migration of lunch-hour trading to Singapore and Hong Kong before the restriction was abolished.

2.b. Stage 2: Identifying types of private information from changes in the U-shape

The second stage of our analysis moves beyond simply documenting higher lunch-hour volatility. The objective of this stage is to discriminate between the types of private information that might cause the volatility effects we find in testing hypotheses 1 and 2. We do this by examining how the U-shape of intraday volatility is affected by the introduction of lunch-hour trading.

There are three key dimensions to shifts in the U-shaped curve: (1) does the curve flatten or curl? (2) does the curve rise or fall? and (3) does the curve tilt up or down? The answers to these questions provide valuable evidence regarding the nature of the private information involved.

Our third hypothesis and test summarizes the implications of recent theory with respect to flattening the U-shape:

H.3: If volatility is caused by the revelation of information in the trading process, then opening lunch-hour trading should flatten the U-shape of intraday volatility.

Test:
$$V_l^Y(closed)/V_m^Y(closed) < V_l^Y(open)/V_m^Y(open)$$

and
$$V_l^Y(closed)/V_a^Y(closed) < V_l^Y(open)/V_a^Y(open)$$

where V_m^Y and V_a^Y are the yen/dollar return variances over the morning and afternoon periods, respectively (10:30-12:00 and 1:30-3:00, Tokyo time).

The second way the U-shaped curve might shift is in its overall level. A recent paper that introduces a helpful taxonomy here is Romer (1993). Romer distinguishes between “internal news”, which must be aggregated by the trading process, and “external news”, which is reflected in price without trading. We add to this the distinction between short-lived and long-lived information. Private information can be short-lived if a future public signal is expected to reveal it. This induces a front loading of informed trading within a given trading period. In contrast, long-lived private information produces a more even trading pattern, as in Kyle (1985). Our fourth hypothesis employs these information categories to generate implications for the average level of volatility over the day.

H.4: If volatility is caused by the revelation of information in the trading process, and that information is wholly short-lived internal news, then opening lunch-hour trading should not effect average volatility over the day.

Test:
$$V^Y(closed) = V^Y(open)$$

AH.4: If the information revealed by the trading process is partly long-lived internal news, then opening lunch-hour trading should increase volatility over the day.

where V^Y is the yen/dollar return variance over all three subperiods (10:30-3:00, Tokyo time).

The fifth and last of our hypotheses addresses the third way in which the U-shaped curve might shift: tilting either up or down. If volatility is redirected toward the afternoon, we should find a tilting up; if it is redirected toward the morning, a tilting down.

H.5: If the information revealed by the trading process is partly short-lived, then opening lunch-hour trading should raise afternoon volatility relative to morning volatility, tilting the U-shape upward.

Test:
$$V_m^Y(closed)/V_a^Y(closed) > V_m^Y(open)/V_a^Y(open)$$

2.c. An additional consideration

There is a complication that derives from the date of the regulatory change (December 22). Traditionally, the Christmas week is one of the least representative since most businesses and currency dealers take vacation (from just before Christmas through New Year's day). Hence, a standard before-versus-after comparison is not a clean test because there is a strong seasonal in the holiday period following the change in regime in our data.

We use two methods to correct for this seasonal. With the first, we omit from the after sample the business days between December 22 and January 2. In the end, we find no substantive effect on our results from this correction. (This is the correction used for our reported results.) The second method we examined is comparing before and after results over intervals exactly one year apart. Though this clearly eliminates the seasonal, the resulting larger gap between the sub-samples introduces other factors that are difficult to control. In any case, lunch-hour volatility is significantly higher in the after sample using this method as well.

3. Data

3.a. Exchange rate data

The yen/dollar and mark/dollar rates we use are the indicative spot quotes posted on the Reuters FFX screen between 29 September 1994 and 28 March 1995 (the first and last dates of our longest sampling interval). Our source for these series is Olsen & Associates (Zurich). Each bid-offer quote is time stamped to the second. (For more details regarding the capture and cleaning of the data, see Dacorogna et al. (1993). For more on their indicative nature, see Goodhart (1989) and Goodhart, Ito, and Payne (1996)).

Though the indicative quotes have their shortcomings, our method of use minimizes our exposure. Specifically, the variance ratios require only the exchange rate series themselves (bid-offer midpoint). This is the dimension of the data that Goodhart, Ito, and Payne (1996) find provides the most accurate measure of the underlying market analogue (cf the series for the indicative spread and the frequency of quote entry).

3.b. The Tokyo and Singapore markets

According to the tri-annual survey of the Bank for International Settlements (BIS), the daily turnover of the Tokyo FX market in the month of April 1995 was \$161 billion. Of this, roughly 80 percent was interbank (i.e., interdealer), either done directly or through brokers (the remaining 20% being customer-dealer trades).⁹ By currency, yen/dollar accounts for 76% and the mark/dollar accounts for 12%. In Singapore, the daily turnover was \$105 billion over the same period. By currency, yen/dollar accounts for 30% and mark/dollar accounts for 32%.¹⁰

⁹ We write “roughly” 80% because the BIS report includes a third category of participant called *other financial institutions*. This category includes nonreporting banks, which in many countries includes investment banks, some of which are important in dealing. That dealers are included in this category is evidenced by the significant brokered trading of this category: only dealers have access to brokers in this market.

¹⁰ Singapore’s currency breakdown is from the 1992 BIS Survey (the 1995 results for the breakdown in Singapore are not yet available to us).

4. Results and Discussion

4.a. Stage 1 results: *The volatility effects from the regime shift*

Our stage one analysis includes the testing of hypotheses 1 and 2. The test of hypothesis 1 is a test of whether volatility rises over the lunch hour from 12 to 1:30 after the opening of trade (i.e., after the removal of the restriction on 22 December 1994). We test this by calculating the ratio of the lunch-hour variance when open to the lunch-hour variance when closed, and testing whether this ratio is greater than one. Table 1 presents the results.

The first column presents the number of trading days in both the closed and open samples (i.e., the before and after samples). For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted—see section 2.c.). Note that 20 trading days corresponds roughly to one trading month. The second column presents the variance ratios, open to closed. The variance ratios are calculated from returns measured as the change in the log of the bid-offer midpoint. P-values for the null that the variance ratios equal one are not included because in all three cases they are zero to two decimal places (i.e., less than 0.5%).

These results are strong evidence that lunch-hour volatility does indeed rise with the opening of trade. In fact, return variance roughly doubles. This is a clear rejection of the null that volatility in FX is caused solely by public information.

Our test of hypothesis 2 is a test of whether mark/dollar volatility also rises over the lunch hour, but by less than the rise of yen/dollar volatility. Recall that this implication is motivated by the presence of alternative trading venues available during the Tokyo market closure. We expect a smaller rise in mark/dollar volatility because mark/dollar is relatively unimportant in Tokyo, implying a smaller effect from bringing Tokyo back on line. Table 2 presents the results.

As in Table 1, the first column presents the number of trading days in both the closed and open samples. The second column presents the variance ratios, open to closed, in the case of mark/dollar over the lunch hour. P-values for the null that these ratios equal one are in parentheses. The last column presents the difference between the corresponding ratio from Table 1 for the yen/dollar rate and the ratio for the mark/dollar rate in column two. (P-values are not presented for this column because the resulting statistic is the difference between two F-statistics, requiring Monte Carlo methods, which we have not yet performed).

Note that volatility also rises significantly in the case of mark/dollar, but by less than yen/dollar. Thus, we cannot reject hypothesis 2, providing further evidence that volatility is induced from trading, and that the source is not public information.

4.b. Stage 2 results: The pattern of volatility over the day

Stage 2 of our analysis moves beyond showing that lunch-hour volatility rises with the opening of trade. Here, we examine how the U-shape of volatility over the day is affected by the lunch-hour opening. There are three questions we address: (1) does the curve flatten or curl? (2) does the curve rise or fall? and (3) does the curve tilt up or down? The answers to these questions provide evidence regarding the type of private information involved.

Our third hypothesis concerns the first of these questions. If volatility is caused by the revelation of information in the trading process, then opening lunch-hour trading should flatten the U-shape of intraday volatility. Table 3 presents our results. This table presents relative variance ratios for the yen/dollar rate over three intraday periods: morning, lunch, and afternoon (10:30-12:00, 12:00-1:30, and 1:30-3:00 Tokyo time, respectively). The second column presents the lunch-to-morning variance ratio after the opening relative to the lunch-to-morning variance ratio before the opening. The third column presents the lunch-to-afternoon variance ratio after the opening relative to the lunch-to-afternoon variance ratio before the opening. P-values for the ratio=1 null are zero to two decimal places for the component variance ratios (i.e., less than 0.5%).

The results clearly indicate a flattening of the U-shape. To see this, note that the denominators imply that lunch return variance is substantially less than both morning and afternoon variance in the closed sample (i.e., the before sample). After trade is opened, however, lunch variance rises relative to both morning and afternoon variance, as indicated by the numerators being higher than the denominators in all cases. Because the numerators remain less than one, the U-shaped curve is not flattened completely.

Our fourth hypothesis concerns whether the intraday volatility curve rises or falls on average. In formulating it, we referred to information types that have different implications for the average level of volatility over the day. Specifically, if the information driving trading-time volatility is wholly short-lived internal news, then opening lunch-hour trading should not effect average volatility over the day. In contrast, if the information revealed by the trading process is partly long-lived internal news, then opening lunch-hour trading should increase volatility over the day. Table 4 presents our results.

The second column implies that the U-shaped volatility curve does rise on average. This column presents the return variance ratio for the yen/dollar rate over all three periods (10:30-3:00 Tokyo time). In all cases, the P-value for the null that the ratio equals one is zero to two decimal places. (Note, however, that these relative increases in variance are much smaller than those that occur over the lunch hour, per Table 1.) Thus, we have evidence that short-lived information is not the whole story, at least to the extent that the opportunity to profit from such information is dissipated by the end of the trading day.

Our fifth hypothesis concerns whether the intraday volatility curve tilts up or down with the opening of lunch-hour trade. If volatility is redirected toward the afternoon, we should find a tilting up; if it is redirected toward the morning, a tilting down. For interpretation, if the information revealed by the trading process is partly short-lived, then opening lunch-hour trading should raise afternoon volatility relative to morning volatility, tilting the U-shape upward. Table 5 presents our results.

Columns two and three indicate that the U-shaped volatility curve does indeed tilt upward. Column two presents the morning variance ratio. Column three presents the afternoon variance ratio. Note from column two that there is evidence that the morning variance actually falls slightly after opening lunch-hour trade. In contrast, the afternoon variance clearly rises after opening lunch-hour trade. This result, in conjunction with the previous result, suggests that short-lived information is indeed being revealed in the trading process, though it is not the only type of information being revealed as internal news (in the sense of Romer (1993)).

4.c. Discussion

So how do these results inform us about trading in FX? Consider three ways of thinking about the stage one results. One view regards FX as perfectly competitive, driven by public information, and nearly infinitely deep (order flow has virtually no price impact). Order flow is liquidity-motivated under this view, coming from customers and from dealers trading for inventory control purposes. In this case, opening midday trade in Tokyo does not affect volatility. A second view admits a role for private information in FX, but considers the drop in depth as resulting primarily from less liquidity trading in Tokyo since the informed would simply choose to route their trades to Singapore over the break. With little decline in informed trading midday, the opening of midday trade in Tokyo should have little impact on volatility. A third view also admits a role for private information, but regards the drop in depth as largely from reduced trading by the informed. (An example is information that is communicated in bilateral

transaction patterns among familiar players. In this case, transactions in Tokyo are inherently different from transactions in Singapore, and the informed may not trade in Singapore because of unfamiliarity with the counterparties.) This reduction in the trading of the informed over the lunch break reduces volatility. Our results support this third view.

5. Conclusions

We reject the view that private information in FX does not exist. Our strong rejection comes from the fact that lunch-hour return variance roughly doubles after the opening of trade. Thus, the view that traders of FX project exclusively on the same public information needs to be relaxed. Relaxing this view is made easier by recognizing that private, price-relevant information need not be private fundamental information (per our introduction).

Examining the regime shift's effect on the U-shape of volatility over the day yields further results. In particular, we find that opening lunch-hour trade alters the U-shape in three ways: the curve is (1) flatter, (2) tilted upward, and (3) higher on average. The flattening of the curve is supportive of the presence of private information that must be aggregated in the trading process (or internal news, in the sense of Romer (1993)). The flattening occurs because allowing trading induces greater revelation during the lunch hour, with a smaller share remaining for the morning and afternoon. The fact that the curve tilts sharply upward is supportive of a predominant role for short-lived private information. The tilt upward occurs because opening trade over lunch reduces the incentive to front-load informed trading in order to avoid a lunch-hour signal that dissipates its value. Finally, the raising of volatility on average suggests that at least some of the private information being aggregated through trading is long-lived. Here, by long-lived we mean information whose value is not fully dissipated at the end of the trading day by the arrival of public information (see, e.g., Foster and Viswanathan (1990)).

Table 1

Lunch-Hour Yen/Dollar Variance Ratio

This table presents the return variance ratio for the yen/dollar rate over the lunch-hour (12-1:30 Tokyo time). The numerator of the ratio is the lunch-hour return variance after the lunch market opened (i.e., after December 22, 1994); the denominator is the lunch-hour return variance before the opening. The returns are calculated as the change in the log of the bid-offer midpoint. The left-hand column describes the length of the before and after samples. For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted). P-values for the ratio=1 null are zero to two decimal places in all cases (i.e., less than 0.5%).

<u>Sample</u>	<u>Yen Variance Ratio: Open/Closed</u>
20 days	2.27
40 days	2.04
60 days	2.13

Table 2

Lunch-Hour Mark/Dollar Variance Ratio

This table presents the return variance ratio for the mark/dollar rate over the lunch-hour (12-1:30 Tokyo time). The numerator of the ratio is the lunch-hour return variance after the lunch market opened (i.e., after December 22, 1994); the denominator is the lunch-hour return variance before the opening. The returns are calculated as the change in the log of the bid-offer midpoint. The left-hand column describes the before and after samples. For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted). The second column presents the variance ratio for the mark/dollar rate. The third column presents the difference between the corresponding ratio from Table 1 for the yen/dollar rate and the ratio for the mark/dollar rate in column two. P-values for the ratio=1 null are in parentheses.

Mark Variance Ratio:

<u>Sample</u>	<u>Open/Closed</u>	<u>Yen Ratio - Mark Ratio</u>
20 days	1.11 (0.01)	2.27 - 1.11 = 1.16
40 days	1.23 (0.00)	2.04 - 1.23 = 0.81
60 days	1.96 (0.00)	2.13 - 1.96 = 0.17

Table 3

Does Lunch-Hour Trading Flatten the U-Shape?

This table presents relative variance ratios for the yen/dollar rate over three intraday periods: morning, lunch, and afternoon (10:30-12:00, 12:00-1:30, and 1:30-3:00 Tokyo time, respectively). The returns are calculated as the change in the log of the bid-offer midpoint. The left-hand column describes the before and after samples. For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted). The second column presents the lunch-to-morning variance ratio after the lunch-hour market opened relative to the lunch-to-morning variance ratio before the opening. The third column presents the lunch-to-afternoon variance ratio after the lunch-hour market opened relative to the lunch-to-afternoon variance ratio before the opening. P-values for the ratio=1 null are zero to two decimal places for the component variance ratios (i.e., less than 0.5%).

Sample	Lunch/Morning Open	Lunch/Afternoon Open
	----- Lunch/Morning Closed	----- Lunch/Afternoon Closed
20 days	0.64 / 0.22 = 2.91	0.56 / 0.38 = 1.47
40 days	0.66 / 0.26 = 2.54	0.62 / 0.37 = 1.68
60 days	0.67 / 0.32 = 2.09	0.51 / 0.32 = 1.59

Table 4

Does Lunch-Hour Trading Raise the U-Shape?

This table presents the return variance ratio for the yen/dollar rate over all three periods (10:30-3:00 Tokyo time). The numerator of the ratio is the return variance after the lunch-hour market opened (i.e., after December 22, 1994); the denominator is the return variance before the opening. The returns are calculated as the change in the log of the bid-offer midpoint. The left-hand column describes the length of the before and after samples. For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted). P-values for the ratio=1 null are zero to two decimal places in all cases (i.e., less than 0.5%).

All-Period Variance Ratio:	
<u>Sample</u>	<u>Open/Closed</u>
20 days	1.20
40 days	1.11
60 days	1.30

Table 5

Does Lunch-Hour Trading Tilt the U-Shape?

This table presents the return variance ratio for the yen/dollar rate over the morning and afternoon periods (10:30-12:00 and 1:30-3:00 Tokyo time, respectively). The numerator of the ratio is the return variance after the lunch-hour market opened (i.e., after December 22, 1994); the denominator is the return variance before the opening. The returns are calculated as the change in the log of the bid-offer midpoint. The left-hand column describes the before and after samples. For example, the dates of the 20 day before sample are the weekdays from November 24 to December 21, 1994 and the dates of the 20 day after sample are the weekdays from January 4, 1995 to January 31, 1995 (holiday week omitted). The second column presents the variance ratio for the morning period. The third column presents the variance ratio for the afternoon period. P-values for the ratio=1 null are in parentheses.

	Morning Variance Ratio:	Afternoon Variance Ratio:
<u>Sample</u>	<u>Open/Closed</u>	<u>Open/Closed</u>
20 days	0.77 (0.00)	1.59 (0.00)
40 days	0.81 (0.00)	1.22 (0.00)
60 days	1.01 (0.41)	1.35 (0.00)

References

- Admati, A., and P. Pfleiderer (1988): "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3-40.
- Andersen, T., and T. Bollerslev (1994): "Intraday Seasonality and Volatility Persistence in Foreign Exchange and Equity Markets," Northwestern University Department of Finance Working Paper No. 186.
- Baillie, R., and T. Bollerslev (1991): "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates," *Review of Economic Studies*, 58, 565-585.
- Berry, T., and K. Howe (1994): "Public Information Arrival," *Journal of Finance*, 49, 1331-1346.
- Bollerslev, T., and M. Melvin (1994): "Bid-Ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis," *Journal of International Economics* 36, 355-372.
- Dacorogna, M., U. Muller, R. Nagler, R. Olsen, and O. 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.
- Ederington, L., and J. Lee (1993): "How Markets Process Information: News Releases and Volatility," *Journal of Finance*, 48, 1161-1191.
- Engle, R., T. Ito, and W. Lin (1990): "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market," *Econometrica*, 58, 525-542.
- Fama, E. (1965): "The Behavior of Stock Market Prices," *Journal of Business*, 38, 34-105.
- Foster, F., and S. Viswanathan (1990): "A Theory of Interday Variations in Volumes, Variances, and Trading Costs in Securities Markets," *Review of Financial Studies*, 3, 593-624.
- Foster, F., 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., and R. Roll (1986): "Stock Return Variance: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics* 17, 99-117.
- Goodhart, C. (1989): "'News' and the Foreign Exchange Market," Manchester Statistical Society Working Paper, October.
- Goodhart, C., T. Ito, and R. Payne (1996): "One Day in June 1993: A Study of the Working of Reuters' Dealing 2000-2 Electronic Foreign Exchange Trading System," in J. Frankel et al. (eds.), *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago.

- Harris, L. (1986): "A Transaction Data Survey of Weekly and Intraday Patterns in Stock Returns," *Journal of Financial Economics*, 16, 99-117.
- Hogan, K., and Melvin, M. (1994): "Sources of Meteor Showers and Heat Waves in the Foreign Exchange Market," *Journal of International Economics*, 37, 239-247.
- Hong, H., and J. Wang (1995): "Trading and Returns Under Periodic Market Closures," MIT working paper, July.
- Ito, T., and W. Lin (1992): "Lunch Break and Intraday Volatility of Stock Returns: An Hourly Data Analysis of Tokyo and New York Stock Markets," *Economics Letters* 39, 85-90.
- Kraus, A., and M. Smith (1994): "Beliefs about Beliefs," University of British Columbia working paper, September.
- Kyle, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315-1335.
- Lyons, R. (1995a): "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics* 39, 321-351.
- Lyons, R. (1995b): "A Simultaneous Trade Model of the Foreign Exchange Hot Potato," *Journal of International Economics*, forthcoming.
- Oldfield, G., and R. Rogalski (1980): "A Theory of Common Stock Returns Over Trading and Non-trading Periods," *Journal of Finance*, 35, 729-751.
- Peiers, B. (1995): "Informed Traders, Intervention, and Price Leadership: A Deeper View of the Microstructure of the Foreign Exchange Market," UCLA mimeo, October.
- Romer, D. (1993): "Rational Asset-Price Movements Without News," *American Economic Review*, 83, 1112-1130.
- Slezak, S. (1994): "A Theory of the Dynamics of Security Returns Around Market Closures," *Journal of Finance*, 49, 1163-1211.
- Wood, R., T. McNish, and J. Ord (1985): "An Investigation of Transaction Data on NYSE Stocks," *Journal of Finance*, 40, 723-741.