

Is There Private Information in the FX Market? The Tokyo Experiment

TAKATOSHI ITO, RICHARD K. LYONS, and MICHAEL T. MELVIN*

ABSTRACT

We provide evidence of private information in the foreign exchange market. The evidence comes from the introduction of trading in Tokyo over the lunch hour. Lunch-return variance doubles with the introduction of trading, which cannot be due to public information since the flow of public information did not change with the trading rules. We then exploit microstructure theory to discriminate between the two alternatives: private information and mispricing. Four key results support the predictions of private-information models. Three of these involve changes in the intraday volatility U-shape. The fourth is that opening trade causes mispricing's share in variance to fall.

IT IS A COMMON VIEW that all participants in the foreign exchange market are projecting on the same public information set. A corollary is that private information is irrelevant. This paper provides evidence against this view. The evidence comes from the Tokyo foreign exchange (FX) market, which until recently was restricted from trading over the lunch break (12:00 to 1:30).¹ For that ninety minutes the Tokyo interbank market shut down. In 1994, however, the restriction was abolished. This market opening provides new insights into why return volatility is so much higher during trading hours.²

*Respective affiliations are Hitotsubashi University and NBER; UC-Berkeley and NBER; and Arizona State University. We thank the following for valuable comments: an anonymous referee, Kathryn Dominguez, Martin Evans, Silverio Foresi, Jeff Frankel, Paul Pfleiderer, Andy Rose, Matt Spiegel, René Stulz, Avanidhar Subrahmanyam, Ingrid Werner, and seminar participants at Santa Cruz, LSE, Berkeley, FRBNY, Rutgers, Stanford, the IMF, Arizona, UCSD, Stockholm, ECARE, Geneva, Toulouse, Limburg, the NBER, and the 1997 AEA Meetings. Lyons thanks the National Science Foundation and the Berkeley Program in Finance for financial assistance, and the San Francisco Federal Reserve Bank for hospitality as a visiting scholar.

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. The reason the restriction was abolished—according to news reports—was to regain volume that had migrated to other, unrestricted locations (e.g., Hong Kong and Singapore). The decision to lift the restriction was made on December 21, 1994, 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.

² For evidence on equity returns see Fama (1965), Oldfield and Rogalski (1980), and French and Roll (1986), among others. For FX see Meese (1986) and Hertz, Kendall, and Kretzmer (1990).

Why trading increases volatility is central to the theory of price formation. Three candidate explanations have been proposed: (1) public information arrives primarily during trading hours, (2) private information induces trades that affect price during trading hours, and (3) errors in pricing are more likely to occur during trading hours. To discriminate among them, French and Roll (1986) examine stock market closures for which the flow of public information does not change. They find that return volatility decreases during these closures. Since public information cannot be the cause, they examine the two alternatives. Finding only a small role for pricing errors, they conclude that private information is the main source of high trading-time volatility on the NYSE.

The analysis of this paper has two stages. The first stage is similar in spirit to the analysis of French and Roll (1986) in that we compare volatility across regimes with an unchanged flow of public information. Our first-stage analysis differs in an important way, however: We are addressing a market whose information structure is skewed toward public information. In FX, for example, there is no plausible analog to the inside information so common to equity markets. Nevertheless, the result is similar to French and Roll's: *lunch return variance doubles when trading opens*. Given public information's role in this market, that this information is far from the whole story is all the more striking.

An unchanged flow of public information across lunch regimes is clearly important to our results. We offer four key supporting facts. First, abolishing the trading restriction was not part of a broader policy reform, nor was it the work of the Ministry of Finance; rather, it was an isolated change in regime, unlikely to have correlated with Ministry of Finance policy more generally. Second, we have reviewed the schedule of relevant macroannouncements and find no change over the lunch hour within our sample. Third, the timing of some public information is endogeneous and may be affected by trading rules, an issue that arises in the French–Roll paper as well since their closures were known in advance. Accordingly, we examine the flow of news reports on the Money Market Headline News screen from Reuters, a measure of public information flow commonly used in the literature (see Goodhart (1989)) and Peiers (1997), among many others). Specifically, we test whether lunch-hour news reports in the after sample increased. Before the opening there were on average 18.5 reports, versus 17.5 reports in the after sample (thirty trading days, with precise dates corresponding to our sampling convention described below); thus, there is no evidence of an increase. Finally, because news reports can have differential impact, we estimate a GARCH model, using dummies in the conditional variance equation for minutes with news. We find no difference in the value of the dummies in the before and after samples.³

³ There remains a refuge for the skeptic: suppose the after sample contains an ARCH-type increase in variance due to public information but our news reports miss it due to measurement error. We show below, however, that lunch variance nearly doubles *even after controlling for the increase in whole-day variance in the after sample*. If a lower-frequency volatility innovation is contaminating the after sample then the increase should disappear under this control.

Having eliminated public information as the cause of higher volatility, we discriminate between the alternatives—private information and pricing errors—by exploiting recent results in microstructure (a literature that scarcely existed at the time of the French–Roll paper). This second stage of our analysis is the main contribution of the paper. We draw specifically from work on intraday volatility patterns. It is a stylized fact that an intraday U-shape is present in many markets.⁴ Recent theory rationalizes the U-shape using private information of various types. Empirical work only recently, however, has progressed from documenting the U-shape to determining if private information is indeed involved (see, e.g., Foster and Viswanathan (1993)). Here we link the U-shape to information by comparing the U-shape before and after the introduction of lunch-hour trading. Private-information models make a number of predictions we are able to test. For example, in general they predict a flattening of the U-shape: lunch-hour trading induces greater revelation during that period, leaving a smaller share for the morning and afternoon.

The second-stage analysis produces four sharp results that support the private-information hypothesis. First, the volatility U-shape over the full day does flatten, as predicted by the private-information models. Second, the U-shape over the full day also tilts upward, which suggests information whose private value is temporary: an open lunch hour reduces the incentive to trade early because it reduces the likelihood that price will reflect that information before a position can be opened. Third, there is a clear U-shape in the morning when Tokyo closes over lunch. This morning U-shape disappears after the lunch opening, as the private-information models predict. Finally, we estimate mispricing's contribution to lunch-hour variance from the mean reversion in price (*à la* French and Roll). Its share in variance falls after the opening, which is inconsistent with the increase in lunch variance coming wholly from mispricing.

Of course, consistency with the private-information alternative does not rule out mispricing. Indeed, the question addressed here is whether private information exists in this market, not whether mispricing does not exist. Taken together, our results make a strong case for private information's existence: we know of no symmetric-information models that can explain all our results. To do so the model must not only predict how the U-shape changes upon opening, but also must account for the fall in mispricing's share in variance. At present, trading models that produce these phenomena rely on private information. The most commonly offered counterexamples do not work. For example, those who offer inventory models fail to recognize that the FX market is a multiple-dealer market, with individual dealer inventories that are not common knowledge; if inventories are price relevant then they are private information. Another model that is promising at first blush is that of Brock and Kleidon (1992). However, that model is based on a monopolist marketmaker who has the power to price discriminate; the FX market, in

⁴ Evidence of the U-shape pattern of intraday volatility appears in Wood, McNish, and Ord (1985), Harris (1986), and Andersen and Bollerslev (1997), among many others. For theory on U-shapes see, e.g., Admati and Pfleiderer (1988), Foster and Viswanathan (1990), and Slezak (1994).

contrast, is a competitive multiple-dealer market. Moreover, their model addresses U-shapes in spreads and volume, not volatility.

Before continuing, let us be explicit about our use of the term private information. *Private information is information that satisfies two criteria: (1) it is not common knowledge and (2) it is price relevant.* We consider this a natural definition for thinking about information-motivated trade. Note, however, that the price-relevance criterion is less stringent than most definitions in the literature. For example, it does not require price effects to be permanent. This contrasts with the definition used by French and Roll. Under their definition private information is identifiable precisely because its price effects are permanent, as opposed to mispricing, whose price effects are temporary. (This is implicit in their comparison of long-horizon variance with cumulative short-horizon variance: mean reversion from mispricing lowers long-horizon variance relative to short-horizon.) Under our definition, however, superior information about temporary price effects also qualifies as private information.

To make this distinction between temporary and permanent price effects more concrete, we provide a taxonomy of private information with some examples from FX. We root the taxonomy in theory by considering a canonical two-period trading model in which trading occurs initially at price P_0 , then again at P_1 , and then a terminal payoff F is realized at $t = 2$. In this framework, we shall refer to information on the terminal payoff F as *fundamental private information*.⁵ We offer two examples that are arguably fundamental by this definition. The first is aggregation of information in FX orders that derive from real trade: A country's trade balance is a component of exchange rate fundamentals, and dealers receive private signals of this component long before published statistics are available (Lyons (1997)). The second is central bank intervention: A dealer who receives a central bank's order has also received private information (Peiers (1997)).

In contrast to fundamental private information, we shall refer to information unrelated to the payoff F but relevant to interim prices P_0 and P_1 as *semifundamental private information*. The function determining P_0 and P_1 includes many arguments beyond expectations of the payoff F . One such argument is traders' risk aversion. Others include traders' trading constraints, the supply/distribution of the risky asset, and other features of the trading environment. Insofar as these affect P_0 or P_1 without altering expectations of F , superior knowledge of them qualifies as semifundamental private information.⁶ Consider two examples of semifundamental private

⁵ This information on F can take the direct form of signals of F , or it can take the indirect form of information about the trading environment that allows one to forecast F more accurately than the market at large (see Madrigal (1996) and note that his use of the term "fundamental" is even more restrictive than ours).

⁶ An earlier version of this paper referred to this semifundamental category as "nonfundamental" private information. An early reader, however, expressed consternation: Why, for example, should we consider the effect of risk aversion on price—even if temporary—as nonfundamental? The term "semifundamental" strikes a balance (we recognize that no pair of labels is beyond reproach).

information in the FX market. The first is superior knowledge of the distribution of dealer inventories (i.e., the distribution of the risky asset). Because the transparency of order flow in spot FX is low, a dealer often has superior knowledge of her own and others' inventories. If inventory risk earns a risk premium, as inventory-theoretic models predict, then superior knowledge of this kind allows a dealer to forecast interim price more accurately than the market at large (i.e., as average inventory across dealers becomes known, this induces a change in the risk premium, which requires a change in price even though terminal-payoff expectations remain unchanged).⁷ Our discussions with FX dealers certainly indicate that this first example is operative. Risk aversion, of course, is not the only source of semifundamental private information. Our second example recognizes that even when traders share common information on F they may still disagree on the meaning of this information, thereby affecting P_0 or P_1 . Consider, for example, the disagreement among financial analysts about the direction of the stock market despite access to the same economic data. Theory identifies the source of this disagreement as different prior beliefs, different models, or both (e.g., Harrison and Kreps (1978), Varian (1985), Harris and Raviv (1993), Kandel and Pearson (1995), and Kraus and Smith (1994)). Whatever the source, if one adds superior knowledge about others' beliefs, then one can forecast interim price more accurately than the market at large. In the FX market, it is a fact that dealers share their beliefs with certain other dealers, but do so rather selectively. Any price implications are likely to be temporary.

Our taxonomy highlights the semifundamental category for two reasons. First, we view it as especially relevant in FX, per our examples. Second, previous literature has neglected this class of private information. Information-theoretic models of trading are specified with private terminal-payoff information, our fundamental category. Empirical models follow suit. It is not clear, however, how much of the information effect found in the data comes from the fundamental category (see for example Hasbrouck (1988) and Madhavan and Smidt (1991)).

There are two branches of the empirical literature that closely relate to this paper. The first branch addresses the volatility effects of closure, though in markets other than FX. French and Roll (1986) on the closure of the NYSE is the seminal work in this branch. Other relevant papers include Barclay, Litzenberger, and Warner (1990), Amihud and Mendelson (1991),

⁷ Though this model is inventory-theoretic, we have added private information since individual dealer inventories are both price relevant and not common knowledge. For more on inventory effects on price see Ho and Stoll (1983) and O'Hara and Oldfield (1986) among many others. For a model of intraday patterns in inventory effects see Spiegel and Subrahmanyam (1995). Empirical evidence of inventory effects on price is provided in Lyons (1995).

To clarify further, consider the other traditional source of price effects from so-called liquidity trades (i.e., trades that are not motivated by information on F). In risk-neutral models with asymmetric information, liquidity trades affect the signal extracted from order flow, thereby affecting expectations and price. Superior knowledge of the liquidity trade in this risk-neutral case, however, allows one to estimate F more accurately than the market (as in Madrigal (1996)). It therefore qualifies as fundamental private information by our definition.

and Ito and Lin (1992). Barclay et al. examine weekend return volatility on the Tokyo Stock Exchange (TSE), exploiting a phase-out of half-day trading on Saturday. They show that weekend volatility fell after the phase-out. Regrettably, one cannot rule out public information as the cause; during the same period, Saturday announcements and other market activities were also phased out. Amihud and Mendelson and Ito and Lin also examine the TSE, but their focus is the lunch closure. Here again, one cannot control for the flow of public information: there was no change in lunch regime on the TSE so inferences are necessarily limited. Amihud and Mendelson show, among other things, that lunch volatility is significantly lower than in the morning and afternoon. They attribute this to "trading and the associated process of information dissemination," without being able to discriminate public from private information. Ito and Lin (1992) compare lunch volatility on the TSE with that on the NYSE, which does not break for lunch. They find the lunch dip in volatility is much deeper in Tokyo than in New York and attribute this to the suspension of trading. Turning now to the second branch of related empirical work, this branch addresses volatility patterns in the FX market and emphasizes the transmission of volatility across trading centers and time (e.g., Engle, Ito, and Lin (1990), Baillie and Bollerslev (1991), Dacorogna et al. (1993), and Hogan and Melvin (1994)). A central result is that volatility propagates from one trading center to the next—the so-called "meteor shower" effect; for example, a volatile Tokyo market is followed by a volatile New York market. Another especially relevant paper within this branch is that by Hsieh and Kleidon (1996), who use spread and volatility patterns measured simultaneously across different trading centers to evaluate models of asymmetric information. They conclude that standard models cannot fully account for the overlapping patterns they find.

The rest of this paper is divided into four sections. Section I summarizes the hypotheses we test, Section II describes the data, Section III presents our results, and Section IV concludes.

I. Testable Hypotheses

A. Stage 1: French–Roll Analysis of Lunch Volatility

The first stage of our analysis follows French and Roll (1986). First, we test whether FX volatility is caused solely by public information. To do this, we test whether the introduction of trading affects lunch volatility. If volatility is caused solely by public information then volatility will not be affected, for the flow of public information did not change. Summarizing our first hypothesis:

H1: Volatility in the FX market is caused solely by public information.

Test: $V_L^c = V_L^o$,

where V_L^c and V_L^o are the yen/dollar return variances over the lunch period (12:00–1:30 Tokyo time) for the closed and open regimes, respectively. The

closed regime is the period before removal of the trading restriction (i.e., before December 22, 1994). The open regime is the period after removal of the restriction.

Since rejecting H1 is consistent with both private information and mispricing, discriminating between them requires additional evidence. Stage two of our analysis generates the additional evidence from the changing pattern of volatility over the day. French and Roll, in contrast, use a different sort of evidence. Theirs derives from an identifying assumption: Mispricing has temporary effects on price but private information has permanent effects. To operationalize this, they decompose the return in period t , R_t , into two components, an information component I_t and an error component E_t :

$$R_t = I_t + E_t, \quad (1)$$

where the error component includes both pricing error and measurement error (they explicitly link the latter to bid-ask bounce). Since the error component's effect on price is temporary, the variance of returns over long holding periods reflects only the information component, whereas the cumulated variance over short subintervals includes both components. Letting V_l denote the return variance over the long holding period and V_s denote the cumulated variance over short subintervals, $1 - V_l/V_s$ provides an upper bound on the fraction of variance from mispricing (it is an upper bound because $V(E_t)$ includes the effect of bid-ask bounce)⁸:

$$1 - \frac{V_l}{V_s} = \frac{V(E_t)}{V(R_t)}. \quad (2)$$

French and Roll find that in the stock market, the upper bound on the fraction of variance from mispricing is about 12 percent. (This is based on a long holding period of six months and a short subinterval of one day.)

Though insightful, the French–Roll method for distinguishing mispricing from private information has its shortcomings. One part of their identifying assumption requires that private information's effects on price are permanent; per our introduction, however, there are classes of private information whose price effects are only temporary. In this instance their method overestimates the role of mispricing (at least, conditional on our definition of private information). The other part of their identifying assumption requires that mispricing's effects are temporary (dissipate within six months). It is possible, however, that mispricing is so persistent that mean reversion is not complete within six months (see Summers (1986) and Fama and French

⁸ This measure requires, of course, that the components are uncorrelated. Note also that, although our FX data are not plagued by the bid-ask bounce generated from transactions data, our price series may still include shading due to inventory control. Because the prices are input by different banks, this might induce a similar temporary bounce in the return series.

(1988)). In this instance their method underestimates mispricing. The sign of any net bias is unclear.

In light of the difficulty of interpreting the bound's level, here we address instead how opening trade changes the bound. This allows us to exploit information in the regime shift more fully. (Note that the way French and Roll measure the bound does not utilize their closure experiment in any way.) A rise in the bound is easily rationalized by greater mispricing: opening trade raises both the lunch variance and the fraction due to mispricing. A fall, on the other hand, is inconsistent with the no-private-information null. Under this null the variance of the information component, $V(I_t)$, is not changing since private information does not exist and the flow of public information is unchanged across regimes; it is thus impossible for total lunch variance to increase with a falling fraction due to the error component E_t because the change in total variance can only come from E_t .

B. Stage 2: Identifying Private Information from Changes in the Intraday U-Shape

The second stage of our analysis discriminates between private information and mispricing by examining volatility changes in the morning and afternoon. Specifically, from private-information-based theory we generate predictions for how intraday volatility responds to market opening. Since models of mispricing do not make the same predictions, if the predictions are borne out then this indicates private information is present.

The first hypothesis we test in stage 2 draws on the way private information is commonly modeled (see Admati and Pfleiderer (1988), among many others). In particular, if the amount of private information does not change when trading rules change, the result is that price impact is merely redistributed over time:

H2: If private information is revealed by trades then opening over lunch will flatten the volatility U-shape.

$$\text{Test:} \quad \frac{V_L^c}{V_M^c} < \frac{V_L^o}{V_M^o} \quad \text{and} \quad \frac{V_L^c}{V_A^c} < \frac{V_L^o}{V_A^o},$$

where V_M and V_A are the yen/dollar return variances over the morning and afternoon trading periods, respectively (as above, superscripts denote closed versus open regimes and the subscript L denotes lunch). Note that this is a joint test; a rejection may indicate changes in the amount of private information rather than nonexistent private information. It is therefore conservative as a test of whether private information exists: any alternative under which the amount of private information changes is an alternative under which private information exists.

Another prediction of private-information models relates to the morning. In particular, if trading is restricted over lunch then a morning-period U-shape should be present. Moreover, this morning U-shape should disappear when

the restriction is lifted (see the closure models of Slezak (1994) and Hong and Wang (1997)).

H3: If private information is revealed by trades then morning volatility will be U-shaped when trading is restricted over lunch and this U-shape will disappear when the restriction is lifted.

$$\begin{aligned} \text{Test:} \quad & V_{EM}^c > V_{MM}^c \quad \text{and} \quad V_{MM}^c < V_{LM}^c \\ & \text{and} \\ & V_{MM}^o \geq V_{LM}^o, \end{aligned}$$

where V_{EM} , V_{MM} , and V_{LM} are the yen/dollar return variances over the early-morning, mid-morning, and late-morning trading periods, respectively. The last relation comes from the full-day U-shape that emerges after the lunch opening.

Our fourth hypothesis addresses the horizon over which private information is privately valuable. Note that the duration of private value is distinct from whether information is fundamental or semifundamental: even fundamental private information, with its permanent price effects, can have short-lived private value if others share the same information (since it can be rapidly reflected in price due to competition between the informed; see Holden and Subrahmanyam (1992)). Thus, there is no clean mapping between type of private information and duration of private value (the latter being determined by the timing of price effects). And since here we only examine price effects, we are unable to test for specific types of private information directly.

The test we perform determines whether private information is of long-lived private value. The well-known model of Kyle (1985) provides a nice example of long-lived private value. In that model the informed trader chooses to trade over time such that his information is reflected in price evenly. Accordingly, here we test whether the opening of trade tilts the volatility U-shape. Under the long-lived null, opening trade should not affect morning volatility relative to afternoon.

H4: If the private value of information revealed by trades is long-lived, then opening over lunch will not tilt the volatility U-shape.

$$\text{Test:} \quad \frac{V_M^c}{V_A^c} = \frac{V_M^o}{V_A^o},$$

where, again, V_M and V_A are the return variances over the morning and afternoon trading periods, respectively. The alternative of short-lived private value predicts the U-shape will tilt upward. To see this, consider private information that may become common knowledge via public news. Then an open lunch tilts the U-shape upward because it reduces the incentive to trade early; it reduces the likelihood that price will reflect the information before a position can be opened. (It is precisely this feature that drives the private-information model of Foster and Viswanathan (1990).)

C. A Test of Robustness

The Tokyo experiment is not a pure regime shift (closed to open) because there were alternative markets where Tokyo dealers could route trades during the ninety-minute break. In fact, Japanese newspapers report that there was indeed migration of lunch-hour trading to Hong Kong and Singapore before the restriction was abolished. Thus, depth in yen/dollar trading fell sharply over the break, but did not go to zero. If depth had gone to zero our tests would be more powerful. They are not biased, however, by nonzero depth.

There is a way to measure this power loss, and at the same time check the robustness of our test of lunch volatility. Note that the ratio of yen/dollar trading to mark/dollar trading in Tokyo is much higher than in Singapore or Hong Kong (data are in the next section). 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). Thus, although we expect the effect of lunch opening on volatility in these two currencies to have the same sign, the effect should be larger in the case of yen/dollar. Summarizing our fifth hypothesis and test, we have:

H5: If private information is revealed by trades then opening over lunch will increase mark/dollar volatility, but by less than yen/dollar volatility.

$$\text{Test:} \quad V_L^c(DM) < V_L^o(DM) \quad \text{and} \quad \frac{V_L^o(DM)}{V_L^c(DM)} < \frac{V_L^o}{V_L^c},$$

where $V_L(DM)$ is the return variance of the mark/dollar rate over the lunch period.

II. Data

A. Exchange Rate Data

The yen/dollar and mark/dollar rates we use are the indicative spot quotes posted on the Reuters screen called "FXFX" between September 29, 1994, and March 28, 1995, the first and last dates of our longest sampling interval. (One of our robustness tests, described in the next subsection, makes use of an additional subsample of FXFX quotes from January 1994 to March 1994.) Our source is Olsen & Associates Research Institute for Applied Economics, in Zurich. Each bid-offer quote is time stamped to the second. The results we report are based on data with a periodicity of one minute, constructed by taking the prior quote closest to each minute.⁹

⁹ It is well known that the Reuters FXFX data contain significant noise that can be reduced by sampling at fixed time intervals. See Zhou (1996) for an analysis of this issue. Note that, for example in our 20-day sample, there are an average of 15.3 raw quotes per daily 12:00–1:30 interval in the before sample and 39.8 quotes in the after sample. For more detail on the capture and cleaning of the raw 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 et al. (1996) find provides the most accurate measure of the real-market analog (in contrast to the spread series and the volume series, where the latter is proxied by the frequency of quote entry).

B. Accounting for the Seasonal

There is a complication that derives from the date of the regulatory change (December 22). Traditionally, Christmas week is one of the least representative because 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 volatility in the holiday period following the change in regime.

We use two methods to correct for this seasonal volatility. With the first, we omit from the after sample the business days from December 22 to January 2. (This is the correction used for our reported results.) The second method compares before and after results over intervals exactly one year apart. Though this clearly eliminates the seasonal, the resulting larger gap between the subsamples introduces other factors that are difficult to control. In any case, lunch volatility is significantly higher in the after sample using this method as well.

We offer an additional fact that is inconsistent with the doubling of lunch variance being due to a seasonal or other spurious factor (e.g., an ARCH innovation in the after sample).¹⁰ In general, these factors would cause volatility to rise equally over the whole trading day. But we find sharp changes in the intraday pattern, with the variance in some subintervals significantly lower in the after sample. (Further, this falling-variance fact works against mispricing as the cause of higher lunch variance: it is unclear why larger errors over lunch should induce smaller errors in the morning or afternoon.)

C. Trading Volume in Tokyo and Other Asian Markets

In Asia, three major trading centers compete for business in an almost identical time zone: Tokyo, Singapore, and Hong Kong. According to the tri-annual survey of the Bank for International Settlements (1996) (BIS), the daily turnover of these three markets in the month of April 1995 was \$161 billion, \$105 billion, and \$90 billion, respectively, after adjusting for double counting. Yen/dollar trading accounts for 76 percent of turnover in Tokyo (\$122 billion), but accounts for only 29 percent in Hong Kong (\$26 billion). Though Singapore does not report a breakdown to the BIS, if one assumes

¹⁰ One possible concern is the Hanshin earthquake, which occurred at 5:46 a.m. on January 17, 1995. For the exchange rate, however, the earthquake was a nonevent; there is no evidence that daily or weekly volatility of the yen/dollar rate increased on the day or in the wake of the earthquake.

that Singapore's breakdown is similar to Hong Kong's then Tokyo clearly dominates yen/dollar trading with a share of more than two-thirds of the combined market. Consequently, the absence of Tokyo-based orders has a substantive effect on the depth of the yen/dollar market. In contrast, the mark/dollar market in Asia is more equally divided among the three trading centers. The share of mark/dollar transactions is only 12 percent in Tokyo (\$19 billion) versus 25 percent in Hong Kong (\$23 billion). Assuming, again, that Singapore and Hong Kong are similar in currency composition, the Tokyo's share of the combined mark/dollar market is less than one-third.

Data on transactions volume provide evidence that the lunch-hour restriction in Tokyo was effective. (The Bank of Japan collects these data daily; unfortunately, they do not collect the data on an intraday basis.) Daily interbank volume in Tokyo increased by roughly 30 percent after the opening of lunch-hour trade, an increase that is significant at the 1 percent level for all three of our sampling intervals. To clarify why it was effective, note that by law all FX orders in Japan must be intermediated by authorized currency-dealing banks (a requirement to this day under the so-called Foreign Exchange Law). Further, all of these authorized currency-dealing banks were subject to the lunch-hour trading restriction. Though customer orders received by authorized dealing banks over lunch could be transmitted overseas and executed, customers could not place orders with overseas dealers directly (here, by customer we mean any participant without currency-dealing authorization).

III. Results and Discussion

A. Stage 1 Results: Lunch Volatility Effects

Hypothesis 1 addresses whether lunch volatility changes after the opening of trade (i.e., after the removal of the restriction on December 22, 1994). To test this we calculate lunch variance when the market is open relative to that when it is closed, and test whether this ratio exceeds one.

Table I presents the results. The first column is 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 II). Note that 20 trading days correspond 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 over all days in the sample. (Note that the usual small sample bias that plagues equity market studies is irrelevant here: the number of one-minute observations in our shortest sample period of 20 days is 1,799.) The *p*-values for the null that the variance ratio equals one are not included because in all three cases they are zero to two decimal places (i.e., less than 0.5 percent).

Table I
Lunch-Hour Yen/Dollar Variance Ratio

This table presents the return variance ratio for the yen/dollar rate over the lunch hour (12:00–1:30 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). There are 1,799 observations in each 20-day lunch-hour sample. The p -values for the ratio = 1 null are zero to two decimal places in all cases (i.e., less than 0.5 percent).

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

These results are strong evidence that lunch 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. As a robustness check, we note that whole-day variance increases far less: the variance over all three trading periods increases by only about 20 percent. The variance ratios in Table I are also significantly higher than 1.2 at the 1 percent level. Controlling for whole-day variance in this way provides evidence against the view that the increase in lunch variance is spuriously due to, for example, a low-frequency ARCH innovation in the after sample.

We turn now to changes in the French–Roll bound on variance from pricing errors. When calculated from our before sample (i.e., closed), the resulting bound over the lunch hour is 80 percent. In the after sample (i.e., open) the resulting lunch-hour bound is lower, only 37 percent.¹¹ Given that total lunch variance goes up with the regime shift, the fall in this bound is inconsistent with the no-private-information null, under which the variance of the information component, $V(I_t)$, is not changing (private information does not exist and public information flow is unchanged across regimes). Because

¹¹ The numerator in our calculation of V_t/V_s is the actual variance of the lunch-period return (12:00–1:30) over 60 trading days (about three months) for both the closed and open samples. The denominator is the variance of the lunch-period return implied by the squared changes in one-minute returns. Whether this measure constitutes an upper bound is questionable because we use a shorter long holding period than do French and Roll. Accordingly, we refrain from use of the modifier upper. Note, however, that our inference here is based on *changes* in the bound, and is therefore not dependent on this being an upper bound.




Figure 1. The Intraday Volatility U-Shape: Before and After. This figure presents the return variance of the yen/dollar rate for 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 before and after samples each include 20 weekdays (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). The two lines are smoothed interpolations of the three variance estimates. The units of the variance are log changes of the bid-offer midpoint per minute (multiplied by 10^8).

the increase in total variance must then come from the error component E_t , it is not possible that the fraction of variance from the error component E_t simultaneously falls.¹²

B. Stage 2 Results: The Pattern of Volatility over the Day

Stage 2 of our analysis examines whether private information is present in the FX market by testing hypotheses H2–H5, all of which involve the U-shape of intraday volatility and how it is affected by the opening. To begin, Figure 1 provides a graphical overview of the effect on volatility over the day. It illustrates clearly the doubling of variance over lunch. The figure also foreshadows our results on the flattening and tilting of the U-shape (H2 and H4, respectively).

¹² More precisely, we are rejecting a three-part null here: (i) no private information, (ii) no change in the flow of public information, and (iii) no change in the persistence of pricing errors. We have already made the case for no change in the flow of public information. For (iii) to be the cause of the rejection would require a substantial increase in pricing-error persistence. There is no basis for believing this to be the case. For comparison, the morning bounds are statistically indistinguishable across the closed and open samples (roughly 60 percent); the same is true of the afternoon bounds (roughly 70 percent).

Table II
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. The *p*-values for the ratio = 1 null are zero to two decimal places for the component variable ratios (i.e., less than 0.5 percent).

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

Our second hypothesis, H2, posits that opening lunch-hour trade will flatten the volatility U-shape if private information is revealed by trades. Table II 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 opening relative to the lunch-to-morning variance ratio before opening. The third column presents the lunch-to-afternoon variance ratio after opening relative to the lunch-to-afternoon variance ratio before opening. The *p*-values for the ratio = 1 null are zero to two decimal places for the component variance ratios (i.e., less than 0.5 percent).

The results demonstrate a flattening of the U-shape. To see this, note that the denominators imply that lunch 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 (numerators > denominators in all cases). Because the numerators remain less than one, the U-shaped curve is not flattened completely.

Our third hypothesis addresses what is perhaps the strongest prediction of the private-information models. It posits that if private information is revealed by trades then morning volatility will be U-shaped when trading is restricted over lunch and this U-shape will disappear when the restriction is lifted. This prediction too is borne out in the data. Though statistically none of the relations expressed in the test of H3 can be rejected, in lieu of reporting them we present this result with Figure 2. The figure presents the estimated variance for four one-hour morning periods, both before and after the opening of lunch-hour trade. Note the pronounced variance increase in




Figure 2. Morning Volatility: Before and After. This figure presents the return variance of the yen/dollar rate for four one-hour intervals, from 9 a.m. to 1 p.m. Tokyo time. The returns are calculated as the change in the log of the bid-offer midpoint. The before and after samples each include 20 weekdays (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). The two lines are smoothed interpolations of the four variance estimates. The units of the variance are log changes of the bid-offer midpoint per minute (multiplied by 10^8).

the hour preceding the lunch break (11:00–12:00 Tokyo time). Once lunch-hour trade opens, however, this variance peak vanishes, and the familiar full-day U-shape appears.

Our fourth hypothesis addresses whether the intraday volatility curve tilts with the opening of lunch-hour trade. Under the null that the private value of information revealed by trades is long-lived, opening over lunch will not cause a tilt (i.e., it will have no effect on incentives to trade early versus late).

Table III presents the result: the curve does indeed exhibit a significant upward tilt. Column two presents the morning variance ratio and column three presents the afternoon variance ratio. Note from column two the evidence that morning variance actually falls slightly after opening lunch-hour trade. The afternoon variance, in contrast, clearly rises after opening lunch-hour trade. This upward tilt of the volatility curve is consistent with information whose private value is short-lived (whether this reflects private information that is semifundamental cannot be determined from these results).

Our test of Hypothesis 5 is a test whether mark/dollar volatility also rises over lunch, but by less than the rise of yen/dollar volatility. Recall that this implication comes from alternative trading venues being available during

Table III
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. The p -values for the ratio = 1 null are in parentheses.

Sample	Morning Variance Ratio:	Afternoon Variance Ratio:
	Open/Closed	Open/Closed
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)

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.

The result (not reported) is the following: dollar/mark volatility over lunch also rises significantly, but by less than yen/dollar. For example, in our 20-day sample the yen/dollar open/closed variance ratio is more than double the mark/dollar ratio, even though the mark/dollar ratio shows an increase that is significant at the 1 percent level. Thus, the prediction in Hypothesis 5 cannot be rejected, providing further evidence that our experiment is robust.

C. Some Perspective on the Results

Though not exhaustive, we offer three characterizations of the spot yen-dollar market. One polar view regards it as wholly driven by public information, with a large number of participants (large in a convergence sense). Under these conditions risk sharing is complete and order flow has no price impact because it is purely allocational (inventory effects on price are thus also ruled out). Opening trade in Tokyo would not affect volatility in this case. Our results do not support this view (though it is arguably predominant among open-economy macroeconomists).

A second view admits a role for private information—either permanent or temporary—but attributes the drop in trading over the closed lunch to fewer allocational trades in Tokyo, with informational trades being routed to Hong Kong or Singapore. Under this view the opening of lunch-hour trade would have little impact on volatility since the lunch closing does not affect infor-

mational trading (any impact would have to come from a changing *share* of informational order flow). In this case, though private information is present, our test would have little power to detect it. Our results do not support this view either.

A third view also admits a role for private information, but in contrast to the second it attributes the drop in trading over the closed lunch hour to fewer informational trades. For example, the informed may delay trades because trading costs outside Tokyo are increased by illiquidity or unfamiliarity with counterparties. This reduction in the lunch-hour trading of the informed reduces lunch volatility. This third view is consistent with our results.

IV. Conclusions

The case for private information in the FX market rests mainly on four facts. The first fact is that opening trade causes mispricing's share in variance to fall. This is inconsistent with a null of no private information and no change in public information; under this null the only possible source of higher lunch variance is mispricing, but this would cause mispricing's share in variance to rise, not fall. Facts two through four strengthen the case considerably since they are predicted by the relevant asymmetric-information theory. Moreover, we are unaware of any model of mispricing that can account for them. Fact two is that opening trade causes the full-day volatility U-shape to flatten, a natural result of informative trades being reallocated to the lunch hour. Fact three is that opening trade causes the full-day U-shape to tilt upward, an implication of information whose private value is short-lived. Finally, fact four is that the lunch break produces a morning U-shape, and opening trade causes that morning U-shape to disappear.

We emphasize that any competing interpretation of these four facts must account for them without relying on private information. At present, no such symmetric-information model exists. Nevertheless, this remains an open area for future research, particularly in the context of multiple dealers.

Summarizing, this paper uses microstructure theory and the path of prices to learn something about the trading environment in FX. The evidence is incompatible with the common view that private information in FX does not exist. Relaxing the common view is eased by recognizing that private information might only predict price over short horizons; that is, it may not be "fundamental" in the sense of the taxonomy in the introduction.

REFERENCES

- Admati, Anat, 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, 1991, Volatility, efficiency, and trading: Evidence from the Japanese stock market, *Journal of Finance* 46, 1765–1790.
- Andersen, Torben, and Tim Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.

- Baillie, Richard, and Tim Bollerslev, 1991, Intra-day and inter-market volatility in foreign exchange rates, *Review of Economic Studies* 58, 565–585.
- Bank for International Settlements, 1996, *Central Bank Survey of Foreign Exchange and Derivatives Market Activity* (Bank for International Settlements, Basel).
- Barclay, Michael, Robert Litzenberger, and Jerold Warner, 1990, Private information, trading volume, and stock return variances, *Review of Financial Studies* 3, 233–253.
- Brock, William, and Allan Kleidon, 1992, Periodic market closure and trading volume: A model of intraday bids and asks, *Journal of Economic Dynamics and Control* 16, 451–489.
- Dacorogna, Michel, Ulrich Muller, Robert Nagler, Richard Olsen, and Olivier 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.
- Engle, Robert, Takatoshi Ito, and Wen-Ling Lin, 1990, Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market, *Econometrica* 58, 525–542.
- Fama, Eugene, 1965, The behavior of stock market prices, *Journal of Business* 38, 34–105.
- Fama, Eugene, and Kenneth French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246–273.
- Foster, Douglas, 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, Douglas, 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, Kenneth, and Richard Roll, 1986, Stock return variance: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 99–117.
- Goodhart, Charles, 1989, 'News' and the foreign exchange market, Working paper, Manchester Statistical Society.
- Goodhart, Charles, Takatoshi Ito, and Richard Payne, 1996, One day in June 1993: A study of the working of Reuters' Dealing 2000-2 electronic foreign exchange trading system, in Jeffrey Frankel et al., eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago, Ill.).
- Harris, Lawrence, 1986, A transaction data survey of weekly and intraday patterns in stock returns, *Journal of Financial Economics* 16, 99–117.
- Harris, Milton, and Artur Raviv, 1993, Differences of opinion make a horse race, *Review of Financial Studies* 6, 473–506.
- Harrison, Michael, and David Kreps, 1978, Speculative investor behavior in a stock market with heterogeneous expectations, *Quarterly Journal of Economics* 93, 323–336.
- Hasbrouck, Joel, 1988, Trades, quotes, inventories, and information, *Journal of Financial Economics* 22, 229–252.
- Hertzel, Michael, Coleman Kendall, and Peter Kretzmer, 1990, The volatility of asset returns during trading and nontrading hours: Some evidence from the foreign exchange markets, *Journal of International Money and Finance* 9, 335–343.
- Ho, Thomas, and Hans Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 1053–1074.
- Hogan, Ked, and Michael Melvin, 1994, Sources of meteor showers and heat waves in the foreign exchange market, *Journal of International Economics* 37, 239–247.
- Holden, Craig, and Avandhar Subrahmanyam, 1992, Long-lived private information and imperfect competition, *Journal of Finance* 47, 247–270.
- Hong, Harrison, and Jiang Wang, 1997, Trading and returns under periodic market closures, Working paper, MIT.
- Hsieh, David, and Allan Kleidon, 1996, Bid-ask spreads in foreign exchange markets: Implications for models of asymmetric information, in Jeffrey Frankel et al., eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago, Ill.).
- Ito, Takatoshi, and Wen-Ling 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.
- Kandel, Eugene, and Neil Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy* 103, 831–872.

- Kraus, Alan, and Maxwell Smith, 1994, Beliefs about beliefs, Working paper, University of British Columbia.
- Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Lyons, Richard, 1995, Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321–351.
- Lyons, Richard, 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275–298.
- Madhavan, Ananth, and Seymour Smidt, 1991, A Bayesian model of intraday specialist pricing, *Journal of Financial Economics* 30, 99–134.
- Madrigal, Vicente, 1996, Non-fundamental speculation, *Journal of Finance* 51, 553–578.
- Meese, Richard, 1986, Testing for bubbles in exchange markets: A case of sparkling rates? *Journal of Political Economy* 94, 345–373.
- O'Hara, Maureen, and George Oldfield, 1986, The microeconomics of market making, *Journal of Financial and Quantitative Analysis* 21, 361–376.
- Oldfield, George, and Richard Rogalski, 1980, A theory of common stock returns over trading and non-trading periods, *Journal of Finance* 35, 729–751.
- Peiers, Bettina, 1997, Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market, *Journal of Finance* 52, 1589–1614.
- Slezak, Steve, 1994, A theory of the dynamics of security returns around market closures, *Journal of Finance* 49, 1163–1211.
- Spiegel, Matthew, and Avanidhar Subrahmanyam, 1995, On intraday risk premia, *Journal of Finance* 50, 319–339.
- Summers, Lawrence, 1986, Does the stock market rationally reflect fundamental values?, *Journal of Finance* 41, 591–601.
- Varian, Hal, 1985, Divergence of opinion in complete markets, *Journal of Finance* 40, 309–317.
- Wood, Robert, Thomas McNish, and Keith Ord, 1985, An investigation of transaction data on NYSE stocks, *Journal of Finance* 40, 723–741.
- Zhou, Bin, 1996, High frequency data and volatility in foreign exchange rates, *Journal of Business and Economic Statistics* 14, 45–52.

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