ELSEVIER
Journal of Empirical Finance 9 (2002) 271–285
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/econbase

Asymmetric information and price discovery in the FX market: does Tokyo know more about the yen?☆

Vicentiu Covrig [a], Michael Melvin [b,*]

[a] Division of Banking and Finance, Nanyang Technological University, Singapore 639798
[b] Department of Economics, Arizona State University, Tempe, AZ 85287-3806, USA

Accepted 23 November 2001

Abstract

We identify a period with a high concentration of informed yen/dollar traders active in Tokyo. Findings include:

1. Japanese quotes lead the rest of the market when the informed are active. At other times, two-way causality is observed in quotes.
2. The contribution of Japanese quotes to yen/dollar price discovery relative to quotes of the rest-of-the-world is 5 to 12 percentage points higher when the informed are active compared to when they are absent.

Private information can, at times, be important, yet "normal" times are characterized as periods where public information results in a high contemporaneous correlation across quotes.

JEL classification: F31; G15
Keywords: Foreign exchange; Microstructure; Yen

☆ Helpful comments on an earlier draft were received from Jonathan Batten, Dan Bernhardt, JinWan Cho, Lutz Kilian, Mark Lundin, Lilian Ng, an anonymous referee and seminar participants at the City University of Hong Kong, Claremont, the Federal Reserve Bank of St. Louis, Nanyang Technological University, Oxford University, the University of Strathclyde, the HFDFII conference, the Econometric Society winter meetings and the Georgia Tech International Finance conference. We assume responsibility for any remaining errors.
* Corresponding author. Tel.: +1-480-965-6860; fax: +1-480-965-0748.
E-mail address: mmelvin@asu.edu (M. Melvin).

PII: S0927-5398(01)00056-1

1. Introduction

If we could identify periods where a financial market experienced a concentration of informed trading activity and then compared this to other periods where such informed trading was less pronounced, what would we expect to find? Theory suggests we should observe important differences. Specifically, microstructure models of financial markets suggest that market dynamics will be affected by the presence or absence of informed traders (see O'Hara, 1995 for a review). Empirical tests of such models are lacking due to the obvious problem of identifying an "informed" trader. We attempt to address this problem by focusing on a particular period of Asian yen/dollar trading when a compression of the time available for trading likely results in a high concentration of informed trading. By exploiting the information in yen/dollar quotes in this market interval, we have evidence supporting the hypothesis that there is, indeed, a high concentration of informed trading activity where the following empirical regularities are found. (1) The quotes of the informed traders tend to lead the quotes of the uninformed. (2) The contribution to price discovery of the quotes coming from the region where the informed are concentrated (Japan) is higher during times that are rich in informed trading than in other times.

The particular period of interest was identified in a recent paper by Ito et al. (1998), who studied a change in trading regime that occurred in the Tokyo foreign exchange market on December 22, 1994. On this day, Japan lifted a restriction against trading the yen over the lunch break from 12:00 to 1:30 (Tokyo time). Ito et al. analyzed foreign exchange market activity during Asian business hours for the periods before and after the lunchtime trading and found evidence consistent with the presence of significant private information in the foreign exchange market. A particularly important piece of evidence in their study was that there is a U-shape of volatility of yen/dollar exchange rate returns in the Tokyo morning when the lunchtime closing occurs that disappears once lunchtime trading begins. This is consistent with models of private information where the informed traders must exploit their informational advantage prior to the market closing so that a compressed trading time results in a faster revelation of the private information. A criticism of the Ito et al. paper was that change in the U-shape of volatility could also have been related to inventory control behavior so that the information story was not necessarily the only interpretation of the empirical findings. The evidence in the present paper, based on more direct tests of the informational asymmetries story, offers new support for the role of private information in foreign exchange markets.

For our purposes, the late morning period in the pre-December 22, 1994 era (which will henceforth be identified as the before period) is viewed to be relatively rich in informed trading compared to the post-December 22, 1994 era (which will henceforth be identified as the after period) when trading occurs through the midday lunchtime so no Tokyo market closing exists. In addition, we hypothesize that the Japanese traders have superior information on the yen/dollar exchange rate relative to non-Japanese traders.

A natural question at this point is to inquire about the source of this informational advantage. While one cannot offer any hard evidence on this question, there is wide recognition of some potential sources. One important source of private information in the foreign exchange market is the customer order flow. For instance, large Japanese banks will receive the foreign exchange business of the largest Japanese corporations and this proprietary information confers a temporary informational advantage on Japanese banks relative to non-Japanese banks in making a market in the yen/dollar. Another potential source of informational advantage is early knowledge of important government data releases or policy actions. Peiers (1997) presents evidence that Deutschebank acts as a price leader in the mark/dollar market around periods of Bundesbank intervention. The inference is that Deutschebank receives a significant portion of the intervention-related currency trades of Germany's central bank and this confers a temporary information advantage. This is but one possible example of how a large and important bank may have temporary sources of private information. In similar fashion, it may be the case that some Japanese banks have superior ability to anticipate policy movements of the Japanese authorities.[1]

We will focus on the pre-lunch before period as a useful sample for testing implications of microstructure theories regarding market dynamics in periods rich in informed trading. Section 2 will discuss empirically testable implications of informed trader concentration as suggested by the existing theoretical literature and present results for one hypothesis: If Japanese traders are believed to have superior information at certain times, then we would expect quotes from Japan to lead quotes from other areas during periods rich in informed trading. Section 2 provides evidence that Japanese quotes lead the rest of the market during the late morning before period. In other periods, two-way causality is observed. Section 3 provides another test of Japanese informed trader presence—is yen/dollar price discovery concentrated in Japan when the informed are active? It is shown that the relative share of information generated by Japanese quotes is higher in the period when the informed Japanese traders are believed to be active than in other periods. Finally, Section 4 provides a summary and conclusions. The battery of tests applied to the yen/dollar data go well beyond the investigation of the U-shape of volatility provided by Ito, Lyons, and Melvin. Such a U-shape might not only reflect private information but could also be a result of inventory rebalancing at the open and close of trading as discussed in Brock and Kleidon (1992) so that the volatility is simply volume induced rather than associated with privately informed trading. However, the following sections will present new evidence of price leadership and price discovery suggesting that Japanese banks are, indeed, better informed than the rest of the market. These new results cannot be explained by simple inventory rebalancing stories.

[1] An interesting and entertaining story told by one experienced trader is found in Niederhoffer's (1997) fable of "The Old Trader and the Yen." He asserts the informational advantage of the Japanese banks when he describes them in the following passage: "At 7:00 p.m. in New York, it is morning in Asia. . .. The banks will wait for their customers to sell the dollar to hedge their export earnings, and then they will buy it. . . They drink sake each night with their friends from school and the ministries, and they learn what is going to be announced and where they will be buying and selling." Whether truth or fiction, if non-Japanese traders believe that Japanese banks have superior information, then the dynamics investigated in the following sections will be expected. Recently, a *Wall Street Journal* article discussed Japanese government investigation of corruption at the Bank of Japan and stated that "Traders have long accused central bank officials of leaking the *tankan* survey, a closely watched quarterly poll of business sentiment that can often suggest when the bank will change interest rates" (Sapsford, 1998, p. A12).

2. Implications of informed trader concentration for Japanese and non-Japanese bank quotes

Due to the lunchtime shutdown of trading in Tokyo prior to December 22, 1994, we expect that there was an exogenous concentration of informed trading prior to the 1.5-h Tokyo market closure. What sort of market dynamics should result from this clustering of informed trading? The literature associated with papers such as Kyle (1985, 1989), Admati and Pfleiderer (1988) and Subrahmanyam (1991) suggests testable hypotheses related to informed trader presence. An Admati and Pfleiderer type model can be used to motivate the intradaily U-shape of volatility. More relevant to the empirical work in the current study is Subrahmanyam's model which indicates that the informational efficiency of price depends upon the number of informed traders. In the context of the problem under study, a greater concentration of informed traders prior to lunch in the pre-December 22, 1994 period should be associated with a faster speed of adjustment of the yen/dollar exchange rate to full-information levels.[2] The intuition behind this result is that the greater the number of informed traders, the more aggressively they trade. An increase in trading aggressiveness causes private information to be revealed in price faster. Therefore, prices are more informative (and will converge more quickly to full information levels), the greater the number of informed traders in the market.

During times of intense informed trader activity, the informed traders are revealing information to the rest of the market regarding unobservable fundamentals. In the case of the trading break during the Tokyo lunchtime, we believe that there should be a clustering of informed traders prior to the market closure as traders seek to exploit any informational advantages prior to the closure. This clustering of the informed has implications for quote dynamics between Japanese and non-Japanese banks.

As discussed in Section 1, it is commonly believed that Japanese banks have superior information regarding the yen. The source of this advantage is, presumably, the close links between banks and large non-bank corporate customers in Japan, who fill their currency needs through the Japanese banks. There may also be a closer relationship between Japanese governmental agencies and Japanese banks than foreign banks. Without being able to identify the source of informational advantage, if the rest of the market believes that Japanese banks are, at times, better informed on the yen/dollar exchange rate, we should expect to observe periods when the Japanese bank quotes *lead* the quotes of banks in other locations.

When the lunchtime closing occurred in Tokyo (the before period), we expect that the late morning pre-lunch period was a period of informed trading concentration, while the same time of day in the after period contained no such clustering of the informed. Therefore, if Japanese banks are believed to have superior information relative to non-Japanese banks, we hypothesize that the before pre-lunch period when informed trader

[2] None of the theoretical papers cited are aimed at the foreign exchange market. The varieties of models following Kyle assumes the one market-maker can observe total order flow. However, the foreign exchange market is a multiple dealer market with low transparency where no one market-maker can observe the market-wide order flow. As an appropriate theoretical model for currencies is still missing from the literature, we do not want to push any particular theory descendent from Kyle as being an accurate model of currency trading.

clustering occurs in Japan, should exhibit Japanese bank quotes causing non-Japanese quotes, but this one-way causality should not appear in the after pre-lunch period. This is, essentially, a test of whether the rest of the market believes that the informed yen traders are concentrated among the Japanese banks.

To test this hypothesis, we examine the evidence regarding causality between Japanese and non-Japanese quotes over the late morning (10:30 am–12:00 noon) in Tokyo for both the before and after periods and compare this evidence to other times of day. The data we use are based on quotes appearing on the Reuters FXFX screen during the period from November 1994 to January 1995.[3] We focus on this short period around the change in trading regime in order to concentrate on the effect of this one event rather than contaminate the sample with extraneous events that may enter in a longer time frame.

Since quotes from Japanese and non-Japanese banks arrive at different points in time and have a different number of observations in the sample we have a nonsynchronous quoting problem. Any approach to this problem requires some potentially controversial assumptions. Harris et al. (1995) chose to use observations only when all markets have quotes, so their sample is defined by the least-frequent quote source. De Jong et al. (1998) explicitly enter missing values for intervals when no new quote appears for a market and then develop an estimator which explicitly accounts for these missing observations. We will utilize two approaches. First, we will employ the methodology of De Jong et al. to utilize all quotes in inferring causality. Second, we will filter the data for quote revisions that would be noticed in the noise of intradaily trading and use these big quote revisions to investigate causality.

2.1. A model of nonsynchronous quoting and causality

The approach employed explicitly takes into account the irregular spacing of the data and uses all observations to investigate lead–lag relationships among Japanese and non-Japanese banks. There are two approaches to deal with irregularly spaced data: one which assumes that the most recent quote remains valid if the quote is not updated and another which explicitly puts a missing value at the intervals during which there is no new quote. We use the latter method with an estimator that takes these missing observations explicitly into account. This approach follows the contributions of Lo and MacKinlay (1990), De Jong and Nijman (1997) and De Jong et al. (1998). The method allows estimates of the cross-covariances between return series when the quotes are irregularly spaced in time. An important feature of the method is that time is measured as calendar (real) time and we assume that there exists a regular underlying unobserved discrete time process from which

[3] The data were obtained from Olsen and Associates Research Institute for Applied Economics in Zurich. The sample periods for analysis are before: November 24 to December 21 and after: January 4 to January 31. The holiday period is omitted from the early after sample as many traders take vacation and market activity is abnormally low. An outlier in the data on December 1, 1994 at 2:26:56 was discovered that had a reasonable bid quote but an ask quote a full yen higher than the bid. Rather than take the average of the bid and ask for this outlier, we use the bid quote in the work below.

observed quotes are drawn. The ultimate goal is to estimate the cross-correlations between these unobserved returns.

The first step in the derivation is to write the difference between two observed quotes as a sum of the price changes of the underlying unobservable quote process:

$$q_{t_{i+1}} - q_{t_i} = \sum_{t=t_i+1}^{t_{i+1}} \Delta q_t \tag{1}$$

where t_i denotes the clock time index of the ith observation. The unobserved quotes are sampled at the clock time frequency and the observed quotes are indexed by the quote time.

The expectation of the cross product of observed returns in the two markets can be written as:

$$E(y_{ij}) = E\left[\left(q^{\mathrm{J}}_{t_{i+1}} - q^{\mathrm{J}}_{t_i}\right)\left(q^{\mathrm{N}}_{t_{j+1}} - q^{\mathrm{N}}_{t_j}\right)\right] \tag{2}$$

One can prove that the expectation of this cross product, conditional on the observed quote times, is a linear combination of the cross covariances γ_k of the underlying returns:

$$E(y_{ij}) = \sum_{t=t_i+1}^{t_{i+1}} \sum_{k=t_j+1}^{t_{j+1}} \gamma_{t-k} \tag{3}$$

where k is the order of the cross covariances and

$$\gamma_k = \mathrm{cov}(r^{\mathrm{J}}_t, r^{\mathrm{NJ}}_{t-k}) \quad \text{where } r_t = q_t - q_{t-1} \tag{4}$$

Let $\chi_{ij}(k)$ be the number of times that γ_k appears in the right hand side of Eq. (3) and we can derive the following expression:

$$\chi_{ij}(k) = \max[0, \min(t_{i+1}, t_{j+1} + k) - \max(t_i, t_j + k)] \tag{5}$$

An important property of Eq. (5) is that $\chi_{ij}(k)$ are functions of the quote times t only, and not of the observed quotes. Substituting Eq. (5) on the right hand side of Eq. (3), we obtain:

$$E(y_{ij} \mid \chi_{ij}) = \sum_{k=-K}^{K} \chi_{ij}(k)\gamma_k \tag{6}$$

where we assume that all covariances of higher order than K are zero.

We can consider Eq. (6) as a regression with the unknown covariances γ_k as parameters and the coefficients $\chi_{ij}(k)$ as explanatory variables. In vector notation, the regression is:

$$y_{ij} = \chi'_{ij}\gamma + e_{ij} \tag{7}$$

Table 1
Wald tests for lead–lag correlations between Japanese and non-Japanese quotes

	Before	After
Japanese quotes lead Non-Japanese	16.2 (0.01)	12.6 (0.03)
Non-Japanese quotes lead Japanese	8.3 (0.14)	15.7 (0.01)

This table reports Wald tests for the joint significance of leads and lags of cross-correlations between Japanese and non-Japanese quotes on the yen/dollar for 20 days before and after the start of lunchtime trading in Tokyo. All data are used by employing an estimator that explicitly takes into account nonsynchronous quoting so that cross-covariances are estimated for quotes that are irregularly spaced in time. Before refers to the late morning period when lunchtime trading was prohibited and a concentration of informed Japanese traders is expected. After refers to the late-morning period after the lunchtime trading prohibition was lifted and informed trading is expected to be diffused throughout the day. We cannot reject the hypothesis that Japanese quotes lead non-Japanese in the before period. In the after period, the evidence suggests two-way causality. The table entries are *F*-statistics for the hypothesis that the lag or lead covariances are jointly equal to zero. This statistic has an asymptotic distribution of $\chi^2(5)$. *P*-values are in parentheses.

The regression (Eq. (7)) can be estimated by OLS. In principle, the regression equation involves all possible cross products in the observed quote changes. However, we can confine ourselves only to the adjacent quotes and we can ignore the observations that are more than K apart. De Jong and Nijman proved that the estimators of γ are consistent and they also provide a method to estimate the asymptotic standard errors of the estimators.

Estimates of the cross-correlations between Japanese and non-Japanese quotes are found as the estimated γ covariances from Eq. (7) divided by the standard deviation of the Japanese and non-Japanese returns over the same interval. Negative lags ($k>0$) represent the correlation between the current Japanese quote and non-Japanese quotes k periods ahead—so this tests whether Japanese quotes lead non-Japanese. Similarly, positive lags ($k<0$) test whether Japanese quotes lag non-Japanese. Wald tests for the joint significance of the leads and lags are reported in Table 1. The Wald tests send a clear signal: there is two-way causality in the after period, but in the before period there is one-way causality as Japanese quotes lead non-Japanese and the reverse causality is statistically insignificant.

2.2. Large quote revisions and causality

The evidence presented in the prior section is persuasive. Additional evidence is provided in this section via causality tests. A strong test of causality would be provided by filtering the quotes to identify relatively large changes that should be noticed in the noise of intradaily quoting. Specifically, we split the quotes into those originating in Japan and those originating elsewhere and filter the quote returns to establish a minimum threshold for size. We test for causality with the following regression:

$$r_t^d = a + br_{t-1}^i + cr_{t-1}^d + e_t \quad (8)$$

where d and i index the origin of quotes as either Japan or non-Japan and r is the exchange rate return. Therefore, r_t^d is the return from consecutive quotes of region d at time t. When *d* is Japan and *i* is non-Japan, we test the null hypothesis that b=0, or non-Japanese returns do not cause Japanese returns. Similarly, when *d* is non-Japan and *i* is Japan, we test the null that Japanese returns do not cause non-Japanese returns. A lagged dependent variable is added to the equation to account for autocorrelation existing in the dependent variable series and guard against the independent variable return simply capturing the time-dependency that exists in high-frequency data. The model is estimated using a Newey–West adjustment.

The returns series are constructed from tick-by-tick data on the yen/dollar and are defined as the change in the log of the midpoint of the bid and ask quotes multiplied by 10,000 (so that the units are basis points). The data are constructed as follows.

Table 2
Causality tests of Japanese and non-Japanese bank quotes

	a	*b*	*c*	q(12)	Observations
(A) Non-Japanese dependent variable, Japanese independent variable					
Before 9:00–10:30	0.185 (0.23)	0.180 (0.00)	−0.258 (0.03)	11.7 (0.47)	270
Before 10:30–12:00	−0.432 (0.03)	0.157 (0.00)	−0.301 (0.00)	10.6 (0.56)	134
Before 1:30–3:00	0.211 (0.23)	0.088 (0.01)	−0.215 (0.02)	13.3 (0.28)	220
After 9:00–10:30	−0.130 (0.17)	0.112 (0.00)	−0.273 (0.00)	10.8 (0.54)	599
After 10:30–12:00	−0.006 (0.96)	0.219 (0.00)	−0.377 (0.00)	13.5 (0.33)	410
After 1:30–3:00	0.105 (0.26)	0.142 (0.00)	−0.241 (0.00)	17.3 (0.10)	462
(B) Japanese dependent variable, non-Japanese independent variable					
Before 9:00–10:30	0.002 (0.99)	0.097 (0.01)	−0.376 (0.00)	12.9 (0.30)	233
Before 10:30–12:00	0.062 (0.67)	0.041 (0.21)	−0.267 (0.00)	17.6 (0.13)	158
Before 1:30–3:00	−0.170 (0.19)	0.135 (0.00)	−0.267 (0.00)	10.0 (0.62)	212
After 9:00–10:30	−0.037 (0.71)	0.149 (0.00)	−0.386 (0.00)	15.9 (0.14)	512
After 10:30–12:00	−0.091 (0.42)	0.130 (0.00)	−0.214 (0.00)	9.2 (0.69)	364
After 1:30–3:00	0.068 (0.56)	0.099 (0.00)	−0.299 (0.00)	10.5 (0.57)	359

Returns, r, are created from tick-by-tick yen/dollar quotes for 20 days before and after the start of lunchtime trading in Tokyo. First, the change in the log of the midpoint of the bid and ask multiplied by 10,000 is created to have units of basis points. Then, in order to identify causality emanating from large quote revisions that should be noticeable in the noise of intradaily quoting, the data are filtered so that only returns greater than 2.5 basis points are retained. This gives the independent variable series, r_t^i. Dependent variable observations, r_t^d, are matched by taking the corresponding dependent variable quote return following the independent variable observations. Causality is then tested by estimating the following equation:

$$r_t^d = a + b r_{t-1}^i + c r_{t-1}^d + e_t$$

When *d* is Japan (non-Japan) and *i* is non-Japan (Japan), we test the null hypothesis that b=0 or non-Japanese (Japanese) returns do not cause Japanese (non-Japanese) returns. The lagged dependent variable is included to account for autocorrelation existing in the dependent variable series and guard against the independent variable return capturing the time-dependency that exists in high-frequency data. Coefficient estimates and related *t*-statistics are reported in the table along with *q*-statistics on residual autocorrelation and the number of observations resulting from the filtering procedure for each time period. Moving average terms were estimated for the following four models where the *q*-statistics indicated autocorrelation problems: Part A—before 1:30–3:00, ma(6); after 1:30–3:00, ma(1). Part B—before 9:00–10:30, ma(10); after 9:00–10:30, ma(1). Estimation incorporates a Newey–West adjustment for heteroskedasticity. The results indicate two-way causality for every period except the late-morning before when Japanese quotes cause non-Japanese. This is consistent with the presence of a high concentration of informed Japanese traders in this period. *p*-values in parentheses.

(i) We begin by filtering the independent variable tick-by-tick returns so that only returns exceeding 2.5 basis points are kept.[4] This filter is intended to reduce the noise in the series and sharpen the test so that independent variable observations are measuring sizable exchange rate changes that would be noticed in the noise of intradaily trading.

(ii) For each observation on the independent variable return, we construct the corresponding dependent variable return by taking the change in the log of the midpoint observations based on the dependent variable quotes following the independent variable returns.[5]

The idea is to examine whether sizable changes in the quotes of one group have a significant effect on the quotes of the other group.[6]

Causality tests are performed separately for the before and after samples. Results are given in Table 2. Note that there is two-way causality for every period except the late morning before. In this period, Japanese quotes are seen to significantly cause non-Japanese quotes in Table 2A. However, in Table 2B, non-Japanese quotes do not significantly cause Japanese quotes for the same period. In every other period, both before and after, two-way causality exists. We interpret these results as supportive of a market-wide perception of a concentration of Japanese informed traders prior to the lunch break in the before period. In this particular period, the actions of the Japanese traders attract the attention of the rest of the market and generate the price leadership behavior found in the data.

Note that the finding of two-way causality is consistent with a market driven by public information arrival. Public information shared by all would result in a general price adjustment across the market that may be spread over a short time. This brief period of quotes adjusting to reflect the new information would generate the two-way causality found in most periods. This is not to say that important private information is not also present during these periods, but rather the private information is not as dominant a factor in “normal” times. The evidence suggests that the Tokyo lunchtime closing created an “abnormal” period of activity where private information revealed by the Tokyo traders dominated market dynamics in the period just prior to the lunch break in trading.

Whether we use only the observations involving large returns, as just reported, or all the data, as in Section 2.1, we find the same result. In the before period presumed to be rich in informed trading, there is one-way causality where Japanese quotes lead non-Japanese. In the after period, two-way causality exists regardless of which methodology is employed.

3. Price discovery in Japan and elsewhere

Having documented that non-Japanese quotes tended to lag behind Japanese quotes for the yen/dollar in the pre-lunch before period, we now want to conduct an analysis

[4] We choose 2.5 basis points on the basis of this equaling approximately 2 standard deviations for the morning data.

[5] For the before (after) sample, the average time between dependent and independent variable quotes was 33 (24) s.

[6] While we later discuss implications of cointegration between Japanese and non-Japanese quotes when observations are sampled at 1-min intervals, we point out that such cointegration issues are not relevant here since observations have been discarded through the filtering for large returns.

following Hasbrouck (1995) on the contribution of Japanese and non-Japanese quotes to yen/dollar price discovery. We begin by assuming that both Japanese (q_j) and non-Japanese (q_n) quotes contain a random-walk component (are $I(1)$) and the quote returns, i.e., the differences of logs of the q, are covariance stationary. These returns may be written in vector moving average form:

$$r_t = \Psi(L)e_t \tag{9}$$

where $\Psi(L)$ is a polynomial in the lag operator and e_t is a zero-mean vector of serially uncorrelated innovations with covariance matrix Ω. While each quote series is nonstationary, the difference between the two series is stationary or the quotes are cointegrated of order 1. Then $\beta' q_t$, where $\beta' =[1 -1]$, is a stationary series, implying that $\beta' \Psi(1)=0$, where $\Psi(1)$ is the sum of the moving-average coefficients. Similar to the Stock and Watson (1988) common trends representation of a cointegrated system, the rows of $\Psi(1)$ are identical. We denote ψ as the common row vector in $\Psi(1)$, and it can be shown that the elements of ψ sum to unity.

As shown in Hasbrouck, the system may be written in error correction form as:

$$r_t = \alpha(\beta' q_{t-1} - E\beta' q_t) + \Gamma_1 r_{t-1} + \Gamma_2 r_{t-2} + \ldots + \Gamma_{k-1} r_{t-k+1} + e_t \tag{10}$$

when there is a nonstationary autoregressive representation of order k for the quotes.[7] The $E\beta' q_t$ term reflects systematic differences in the two quote series. In order to estimate this model, we follow Hasbrouck in estimating this term by calculating the mean value over the sample so that ($\beta' q_{t-1} - E\beta' q_t$) is a "demeaned" quote vector. Once this vector is created, the model is then estimable by linear least squares.

If the error covariance matrix Ω is diagonal, then it is possible to directly calculate the contribution of the innovations in one market, say Japan, to the total variance. This may be thought of as the "information share" of the Japanese quotes. If only private information was relevant, we might expect such a structure. However, the presence of public information, shared equally by all, should lead to the innovations being contemporaneously correlated across markets. In this case, Ω is not diagonal. Hasbrouck recommends a procedure to bound the information shares of each region through a triangularization of Ω. First, let the innovations in the two regions' quotes be given by $e_t=Fz_t$, where F is the Cholesky factorization of Ω and z is a vector of random variables with $Ez_t=0$ and $\mathrm{Var}(z_t)=I$ (an identity matrix). Then the market share of the innovation attributable to a particular z_j is found as:

$$S_j = ([\psi F]_j)^2 / (\psi \Omega \psi') \tag{11}$$

where $[\psi F]_j$ is the jth element of the row matrix ψF. By permuting ψ and Ω, we get an upper (lower) bound on the information share of Japanese (non-Japanese) quotes by placing Japan

[7] De Jong et al. (2001) propose an alternative estimation strategy based upon a model developed in Harvey (1989). Their model imposes cointegration directly and offers a more parsimonious lag structure than the Hasbrouck method. This can often be important in high-frequency applications where long lags are sometimes needed and, consequently, parameter instability may result. However, in our case, we find that inference regarding information shares is unaffected after four lags, so the order of lags is not particularly high and inference is robust with regard to lag structure.

first. When non-Japanese quotes are placed first, we get an upper (lower) bound on non-Japanese (Japanese) quotes. The factorization imposes a structure that maximizes the information share on the first quote and minimizes the information share of the last quote. As Hasbrouck states, using the Cholesky factorization in this manner is analogous to allocating the explained variance in a multiple regression by considering the incremental improvement in the R^2 that occurs when variables are added in sequence. The information share is a measure of who moves first in adjusting quotes. The magnitude of the difference between the upper and lower bound reflects the importance of contemporaneous correlation between Japanese and non-Japanese quotes. With a diagonal Ω there is no such problem and unique values of the information shares may be estimated.

We estimate the information shares for the late-morning before and after periods. The quotes are sampled at 1-min frequency, where the prior quote to each minute is used. No overnight returns are included. The specific functional form of Eq. (10) uses four lags of

Table 3
Information shares of japanese and non-japanese quotes

(A) Japan and non-Japan quotes

Quote origin	Bounds for information shares (%)					
	Before			After		
	Upper	Lower	Mean	Upper	Lower	Mean
Japan	56	41	49	64	30	47
Non-Japan	58	44	51	70	36	53
t-statistic for difference of before and after means: 10.9						

(B) Japan and Hong Kong quotes

Quote origin	Bounds for information shares (%)					
	Before			After		
	Upper	Lower	Mean	Upper	Lower	Mean
Japan	59	54	56	65	40	52
Hong Kong	46	41	44	60	35	48
t-statistic for difference of before and after means: 22.3						

The information shares are based upon the innovations from a vector moving average model of quote returns on the yen/dollar exchange rate from banks in Japan and elsewhere during the late morning in Asia. The sample is split into 20 days prior to the beginning of lunchtime trading in Asia (the before period) and 20 days after the start of lunchtime trading (the after period). Specifically, the information shares are given by the equation:

$$S_j = ([\psi F]_j)^2 / (\psi \Omega \psi')$$

where Ω is the error covariance matrix of the residuals from the quote returns model, ψ is the common row vector of coefficients from the moving average model and F is the Cholesky factorization of Ω. Permuting ψ and Ω, we get an upper (lower) bound on the information share of Japanese (non-Japanese) quotes by placing Japan first. When non-Japanese quotes are placed first, we get an upper (lower) bound on non-Japanese (Japanese) quotes. The distributions and related *t*-statistics are generated via bootstrap methods. The information share of Japanese quotes relative to non-Japanese or Hong Kong quotes is higher in the before period, when the prohibition of lunchtime trading leads to a concentration of informed Japanese trading prior to lunch, than in the after period when lunchtime trading is possible so there is no midday market closing inducing a clustering of informed, late-morning trading.

returns of each region. We experimented with alternative lags running from 2 to 12 and found that generally after four lags, the information shares are essentially unchanged. Table 3A. contains the estimated bounds for the Japanese and non-Japanese information shares. We disaggregate all quotes into those originating in Japan and those originating elsewhere. Comparing the information share of Japanese quotes to those of non-Japanese quotes, we see in the before period that the Japanese information share is 96% of the non-Japanese (49%/51%) at the mean, 97% at the upper bound, and 93% at the lower bound. In the after period, the information share of Japanese quotes falls to 89% of non-Japanese at the mean, 91% at the upper bound and 83% at the lower bound. These results indicate that the before period of heavy Japanese informed-trader concentration has a larger share of information for Japanese quotes than in the after period when the concentration of informed traders is absent. The distribution of the information shares was generated via a bootstrap on the estimates. Following Sapp (1999), we employ the method suggested by Li and Maddala (1997).[8] The associated t-statistic on the difference of means indicates that there is a statistically significant drop in the Japanese information share in the after period compared to the before. This result is consistent with the hypothesis that the rest of the market is looking to Japan for information during the before period as the concentration of the informed Japanese traders tend to lead the market.

The fact that the information share of Japan is estimated to be slightly below that of the rest of the market is an artifact of the comparison of only one country's quotes, Japan, against all the rest of the quotes in the market. The methodology is sensitive to the frequency of quote updating and will be biased in favor of large information shares for a region with more frequent quoting. This is, of course, due to the fact that the more frequently quotes are posted, the smaller the number of 1-min observations with zero returns entering into the calculations. For this reason, we believe the relative information shares between Japanese and non-Japanese quotes are more useful to examine than the absolute estimates of the information shares of a single region.

To illustrate this quote frequency sensitivity, we estimated information shares of Japan versus Hong Kong, with results reported in Table 3B. We now find that Japan's information share in the before period is 59% at the upper bound and 54% at the lower bound while Hong Kong's is 46% at the upper bound and 41% at the lower bound. The less frequent quoting of Hong Kong compared to the entire non-Japanese region results in a higher information share for Japan relative to Hong Kong than Japan relative to all non-Japanese quotes. Calculating the shares of information coming from Japan relative to Hong Kong, we find values of 128% at the upper bound and 131% at the lower bound in the before period. For the after period, the shares are 108% and 114%. Again we bootstrap distributions to generate a t-statistic for the difference of the mean information shares and find a statistically significant drop in the relative share of information coming from Japanese quotes in the after period compared to the before. In both cases of Table 3 we find that the relative share of information generated by Japanese quotes drops in the after period relative to the before. Once the informed trader concentration ends, the importance of Japanese quotes relative to the rest of the market falls.

[8] Sapp provides a good discussion of the bootstrap method in his Appendix.

4. Summary and conclusions

Microstructure models of financial markets have shown that we should expect markets to behave differently when there are many informed traders present than in times of informed trader absence. Empirical tests of the implications of these models have been hindered by the difficulty of identifying the informed traders. We have identified a period in the foreign exchange market where there should be a high concentration of informed trading. Following Ito et al., we look to the late-morning period of yen/dollar trading in Tokyo prior to December 22, 1994 (the before period) as a period where there is a concentration of informed Japanese traders relative to other times. This concentration of the informed arises from the lunchtime closing of yen/dollar trading that occurred in Tokyo until December 22, 1994. After that date (the after period), Tokyo trading at lunchtime was permitted.

We focus on the late-morning in Asian trading, hypothesizing that there should be different market dynamics in the before period compared to the after period that would allow tests of the implications of market microstructure models. To this end, we investigate hypotheses consistent with an informed trader concentration before that would disappear after. If the market believes that Japanese traders are better informed regarding the yen, then any time there is a perceived concentration of the Japanese informed traders, we should expect to see quotes from Japan leading quotes of other regions. We tested this hypothesis by performing causality tests of Japanese versus non-Japanese quotes. In general, there is two-way causality. This is consistent with a market driven by public information arrival shared equally by all so that following the arrival of news, there is a brief period of quotes adjusting across the market to the new information. The general finding of two-way causality is also consistent with a market where private information is always important, but the informed behave strategically and reveal their information slowly over time so that in normal periods one cannot detect a dominant presence of the informed. The Tokyo lunch closing compressed the time available for the informed Japanese traders to exploit their informational advantage so that the period prior to the lunch break finds the informed traders forced to reveal their information more quickly. As a result, in this one period—the late-morning before—we find one-way causality where Japanese quotes lead non-Japanese quotes. In the period leading up to the Tokyo lunchtime closing, the clustering of the informed Japanese traders seems to have created an environment where the rest of the market believed that the Japanese possessed superior information and were acting on that information so that Japanese quote leadership is observed.

A final examination of the implications of informed trader clustering was conducted by estimating the contribution of different trading centers to price discovery. Specifically, we estimate the shares of information that may be attributed to Japanese quotes versus non-Japanese quotes. Given the contemporaneous correlation between Japanese and non-Japanese quotes, we are unable to estimate unique values of the information shares but, instead, estimate upper and lower bounds for each region. We find that in the late-morning before period the information share of Japanese quotes is between 97% and 93% of the information share of the rest of the market. However, in the late-morning after period, the information share of the Japanese quotes is between 91% and 83% of the rest of the market, depending on whether we evaluate at the upper or lower bound of estimates. We find qualitatively similar results comparing Japan with Hong Kong. This drop in the information share of Japanese quotes in the after period is consistent with the dissipation of the informed trader clustering existing in the before period. It appears that the Japanese traders' contributions to yen/dollar price discovery are at a peak when the market perceives these traders to have private information.

Taken as a whole, we have provided evidence on the implications of informed trader clustering in the foreign exchange market. While we cannot make any claims to generality of findings, it would be surprising if such implications did not extend to other financial markets. Returning to the question incorporated in the title: "Does Tokyo Know More About the Yen?", it appears the answer is a qualified "yes." At times, Japanese traders appear to have superior information—or, at least, the rest of the market believes that they do. However, the normal state of the market is one where Japan does not dominate in setting yen/dollar quotes. This does not mean that private information is unimportant in "normal" times, but rather that the informed are able to slowly reveal their information. In this case, it may not be unreasonable to think of the market as largely being driven by public information.

References

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

Brock, W.A., Kleidon, A.W., 1992. Periodic market closure and trading volume. Journal of Economic Dynamics and Control 16, 451–489.

De Jong, F., Nijman, T., 1997. High frequency analysis of lead–lag relationships between financial markets. Journal of Empirical Finance 4, 259–277.

De Jong, F., Mahieu, R., Schotman, P., 1998. Price discovery in the foreign exchange market: an empirical analysis of the yen/dmark rate. Journal of International Money and Finance 17, 5–28.

De Jong, F., Mahieu, R., Schotman, P., van Leeuwen, I., 2001. Price discovery on foreign exchange markets with differentially informed traders. Working Paper, Limburg Institute of Financial Economics.

Harris, F.H.deB., McInish, T.H., Shoesmith, G.L., Wood, R.A., 1995. Cointegration, error correction, and price discovery on informationally linked security markets. Journal of Financial and Quantitative Analysis 30, 563–579.

Harvey, A.C., 1989. Forecasting, Structural Time Series Models and the Kalman Filter Cambridge Univ. Press, Cambridge.

Hasbrouck, J., 1995. One security, many markets: determining the contributions to price discovery. Journal of Finance 50, 1175–1199.

Ito, T., Lyons, R.K., Melvin, M., 1998. Is there private information in the FX market? The Tokyo experiment. Journal of Finance 53, 1111–1130.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.

Kyle, A.S., 1989. Informed speculation with imperfect competition. Review of Economic Studies 56, 317–356.

Li, H., Maddala, G.S., 1997. Bootstrapping cointegrating regressions. Journal of Econometrics 80, 297–318.

Lo, A., MacKinlay, A.C., 1990. An econometric analysis of infrequent trading. Journal of Econometrics 45, 181–211.

Niederhoffer, V., 1997. The Education of a Speculator Wiley, New York.

O'Hara, M., 1995. Market Microstructure Theory Blackwell, Cambridge, MA.

Peiers, B., 1997. Informed traders, intervention, and price leadership: a deeper view of the microstructure of the foreign exchange market. Journal of Finance 52, 1589–1614.

Sapp, S., 1999. Price leadership in the spot foreign exchange market. Working Paper, University of Western Ontario.

Sapsford, J., 1998. Japan's corruption scandal hits central bank. The Wall Street Journal 12, A12 March.

Stock, J.H., Watson, M.W., 1988. Testing for common trends. Journal of the American Statistical Association 83, 1097–1107.

Subrahmanyam, A., 1991. Risk aversion, market liquidity, and price efficiency. The Review of Financial Studies 4, 417–441.