Real Trading Patterns and Prices in Spot Foreign Exchange Markets

Jón Daníelsson and Richard Payne*
Financial Markets Group, London School of Economics.

April 2000

Abstract

Most empirical research on FX market microstructure uses indicative quotes as proxies for firm, tradeable quotes. This paper examines the validity of the approximation implicit in this practice by comparing the characteristics of one week of indicative quote data on the DEM/USD with those of contemporaneous, transactions–based data drawn from Reuters D2000–2, an electronic FX broking system. The following comparative results emerge from a very high frequency analysis. Indicative quote returns are found to be significantly more volatile and more strongly autocorrelated than D2000–2 quote returns. Unlike bid–ask spreads on D2000–2, those in the indicative data contain virtually no information on market liquidity. In addition indicative quote returns lag firm quote returns by up to three minutes. However these anomalies disappear with aggregation. The statistical properties of the indicative quotes are very similar to those of firm quotes at a 5 minute sampling frequency, and are virtually indistinguishable from those of firm quotes and transaction prices when both are sampled once every 10 minutes.

*We would like to thank Casper De Vries, Charles Goodhart, Enrique Sentana, two anonymous referees plus seminar participants at LSE and CEMFI for helpful comments. All remaining errors are our own responsibility. Thanks also to Reuters Group PLC and Olsen and Associates for providing the data. Updated versions of the paper may be downloaded from http://fmg.lse.ac.uk/~jond/research and http://fmg.lse.ac.uk/~payne.
For correspondence: j.danielsson@lse.ac.uk and r.payne@lse.ac.uk, respectively

1 Introduction

In recent years there has been a large increase in the amount of research devoted to the microstructure of foreign exchange markets and the behaviour of high–frequency exchange rate data. While some of this research has been driven by theoretical and statistical advances, the main impetus has been the increased availability of tick–by–tick exchange rate data to researchers. This data has been derived primarily from the FXFX, and later EFX, pages of Reuters information systems and has been extensively employed in empirical studies. For example, FXFX midquotes are widely used as proxies for transaction prices in analysis of intra–day exchange rate volatility; (Baillie and Bollerslev (1991), Dacorogna, Müller, Nagler, Olsen, and Pictet (1993), Andersen and Bollerslev (1997) and Payne (1996)); triangular arbitrage relationships (de Jong, Mahieu, and Schotman (1998)); and intra–day technical trading rule performance (Curcio, Goodhart, Guillaume, and Payne (1997)). FXFX spreads have been used as measures of FX market liquidity in studies such as Bollerslev and Melvin (1994) and Hartmann (1996). Finally, FXFX quote frequency is used as a proxy for traded currency volumes in Bollerslev and Domowitz (1993) and Melvin and Yin (1996).

EFX data, however, has a number of shortcomings. First, and most importantly from a microstructure perspective, it contains no measure of traded currency volumes. This renders many interesting microstructure hypotheses untestable. Second, the bid and ask quotes derived from EFX screens are *indicative* rather than *firm*. This means that such quotes are not binding commitments to trade from the originator and hence they may not be accurate measures of tradeable exchange rates. Furthermore, while the EFX system gives a timestamp for the entry of a quote pair, no such timing is given for the exit of quotes. Hence there is no information on the effective lifetime of EFX quotes. Last, each EFX bid and ask quote pair is input by a single dealer. As such, these quotes are likely to reflect dealer specific characteristics (e.g. inventories or beliefs) and may be a poor representation of 'market quotes'.

These shortcomings raise concerns about the validity of EFX data as a proxy for both transaction prices and firm quotes. Below, we present an empirical comparison of the features of data derived from both indicative EFX quotes and firm data drawn from Reuters D2000–2 (an electronic FX broking system), using tick–by–tick observations for the same five days. As such, this work builds upon, extends, and clarifies the results in Goodhart, Ito, and Payne (1996) who compare the statistical features of EFX data and data from D2000–2. However, their analysis is subject to a number of limitations. First, their data covered only a single day. Hence, for example, they could not examine recurrent intra–daily patterns in the data and couldn't provide accurate information on how the statistical features of the data varied with time–of–day. The second main limitation of their study is that their D2000-2 data is not timestamped. Hence, in order to construct approximately contemporaneous data sets they match them by maximizing the correlation between EFX and D2000–2 midquotes. Clearly, this implies that Goodhart, Ito, and Payne (1996) cannot properly examine the lead–lag relationships between EFX and D2000–2 returns and volatility measures.

Our analysis is based on a newly available transactions–based data set on the DEM/USD drawn from Reuters D2000–2.[1] The data covers the week from October 6 to October 10, 1997 and, in addition, we have EFX data from the same week. We examine the following statistical features of series drawn from the 2 data sources;

1. Intra–day activity patterns.

2. Sample moments and dependence measures for returns.

3. Cross-correlations between firm and indicative returns.

4. The information content of each quotation series.

5. Volatility spillovers from D2000–2 to EFX returns and vice versa.

6. The manner in which the relationship between the data series depends on sampling frequency.

Our results indicate that EFX quotes are a poor proxy for firm quotes at very high sampling frequencies (20 seconds.) The EFX midquote is far more volatile and more strongly autocorrelated than its D2000–2 counterpart. Further, EFX spreads contain little/no information on the pace of the market. D2000–2 midquote returns lead those on EFX by 2 to 3 minutes and the proportion of all information entering the exchange rate via the D2000–2 price is above 90% (during normal trading hours.)

We surmise that EFX quotes tend to be used primarily as an indication of willingness to trade on a certain side of the market. Generally an initiating trader wishes to either buy or sell, but not both. To signal his/her intentions, a standard, large spread is used to make only one side of the quote a keen price. Then, as the quotes contain the name and location of the originating

[1]The data is available for academic research from the Financial Markets Group at the London School of Economics. See `http://fmg.lse.ac.uk`. To our knowledge, the data have not been used as yet in any published work.

institution, this trader can be easily contacted by potential counterparties. Such quote setting behaviour will imply precisely the statistical data features indicated above.

As one might expect, as sampling frequency is reduced, many of the discrepancies between the indicative and firm data disappear (with the exception of the differences in spread behaviour.) We find that at levels of aggregation of 5 minutes and above, EFX returns are a fairly good proxy for firm returns. At a sampling frequency of 10 minutes, the two return series are extremely similar.

The rest of the paper is set out as follows. The next Section contains a brief description of the nature of the spot FX markets and how the D2000–2 and EFX systems fit in to the market structure. Sections 2.1.2 and 2.1.1 give an overview of the two data sources and the structure of each data set. Section 2.2.1 provides our analysis of the activity patterns on EFX and D2000–2, mainly through a set of graphics. In Section 2.2.2 we describe the basic statistical nature of the two sets of midquote returns. Section 3 contains an analysis of the information content of each of the two return series and our bivariate GARCH results and our analysis of the effects of temporal aggregation on return characteristics is contained in Section 4. Section 5 provides conclusions and presents ideas for further work.

2 The Spot FX Market

2.1 Spot FX Market Structure

The spot FX market has grown tremendously in recent years. According to the BIS surveys of FX activity in 1995 and 1998 (Bank for International Settlements (1996), Bank for International Settlements (1998)), total spot volume in all exchange rates has risen annually by over 4% in recent years to a daily turnover figure of just above $590bn. Of this global turnover the most commonly traded currency was the US dollar which was a counterparty to 44% of transactions, and the second largest was the German mark with 15%. The largest trading center was London, with 32% of global activity. The market for spot DEM/USD (analyzed here) was the most one of the most liquid in the world, and that DEM/USD activity can be expected to center around London trading hours.

Transaction activity in spot FX is heavily fragmented and largely opaque. Order flow can be divided into a number of segments;

1. Customer–dealer (approx. 25% of volume)

2. Inter–dealer (approx. 75% of volume)
    (a) Direct inter–dealer trade
        i. Telephone based trade
        ii. D2000–1 based trade
    (b) Indirect inter–dealer trade
        i. Voice brokered trade
        ii. Electronically brokered trade (EBS and D2000–2)

Electronically brokered trade has been the major market innovation of the last few years and has grown tremendously in volume terms, mainly at the expense of voice brokered trade. One of the two main vendors of electronic brokerage is Reuters through its D2000–2 service.

2.1.1 EFX

One source of FX price information which is available to all market participants is the Reuters' EFX page. This screen provides a continuously updated sequence of bid and ask exchange rate quotation pairs (timestamped to the nearest second) from individual institutions whose names and locations are also displayed. EFX does not contain any data on traded volumes but gives an evolving picture of the quotes available from other dealers. These quotes, however, are not *firm* but *indicative* i.e. they do not present a binding commitment from the advertising institution to trade at these prices. It has been argued that reputation considerations would almost force those submitting EFX quotes to regard them as firm (see Bollerslev and Domowitz (1993), for example), but this is an assertion which has not been empirically validated. During our sample period there were 32,121 quotes entries on the EFX system.

2.1.2 The Reuters' D2000–2 Dealing System and Data

The Reuters D2000–2 data set consists of all entries to the D2000–2 system for the week of October 6 to October 10, 1997. The system display and its basic trading mechanism are as follows.

Reuters D2000–2 operates as an electronic limit order book with liquidity supply via limit order and liquidity demand via market order (and direct limit order crosses.) A subscriber to D2000–2 sees the following items on the trading screen, for up to 6 exchange rates;

- Best limit buy and sell prices

- The quantities available for trade at the best prices
- An indicator of the characteristics of the last trade

These data items are available to us on a tick–by–tick basis. Furthermore, the data set contains information not available to market participants, specifically every away limit order on the D2000–2 order book at every point in time.[2] Hence we can examine variations in liquidity supply which are unknown to those actually trading. The entry and exit times of all limit and market orders are supplied to the one hundredth of a second.

Limit orders are queued via price and then time priority. In general, market orders will hit the best outstanding limit order on a given side of the market. There are exceptions to this rule however. In a small number of cases in the data set a market order failed to complete because of a lack of bilateral credit arrangements between the counterparties. As a result, at some points in time, the initiator of a market order may find the best priced limit order unavailable to him/her implying that market orders may also trade outside the touch. This causes complications in processing the D2000–2 data set, but since such occurrences are rare, we feel that they will not affect the results in any way.

Trades also occur when bid and offer limit orders cross i.e. the book contains a bid limit order with price greater than or equal to the best outstanding limit sell.[3] These crosses occur automatically on the system and are straightforward to retrieve from the data set supplied. Again, in a few cases the best limit buy and sell do not cross for reasons outlined above, resulting in negative bid–ask spreads. Such observations are deleted before constructing the data used in this study.

For the five trading day period included in the data set, there were 130,535 system entries of one of the following types;

1. Limit buy entry.
2. Limit sell entry.
3. Take: a market buy order.
4. Hit: a market sell order.

In addition there are a few other entry types, most of which can be reclassified within one of the four preceding categories.

[2]By an away order we mean a limit sell with price above the current best or a limit buy with price below the best.

[3]The price of such a transaction is that of the limit order entered earliest i.e. the system treats the order entering latest similar to a market order.

2.1.3 Data Processing and Aggregation

The following basic D2000–2 variables were employed. First, the best bid and offer in the system at every observation point were used as our real quotes. From these, the midquote and bid–ask spread were constructed. Further, as measures of FX market activity we constructed time series of the number and aggregate size (in $m.) of all limit orders outstanding as well as traded volume. Finally, the actual transaction prices are directly observed, and are used here as well. Midquote returns for the EFX data were constructed along with spread and quote frequency variables.[4]

In order to convert both EFX and D2000–2 data from event to calendar time, the following procedure was used. For EFX quotes, the final observation pair in each calendar time interval was recorded. The D2000–2 quotes used are the best limit bid and offer prices outstanding at the end of each interval. For EFX quote frequency and D2000–2 transaction frequency/volume, the number (or, in the case of volume, aggregate $ quantity) of each event occurring in each interval was calculated. Finally, the number and total size of limit orders outstanding at the end of each calendar time interval was recorded.

It has been suggested to us that our sampling convention of taking the final EFX observation in each interval biases our results against the EFX data due to their single dealer nature. The suggested alternative was to attempt to combine information from current and past EFX quotes in order to gain a set of bids and offers which better approximate at the market quotes.[5] Our response to this is two–fold. First, most work on EFX has used the quotes as we do.[6] Second, implementation of this suggestion implies that one needs to use a necessarily ad hoc procedure for estimating the lifetimes of EFX quotes. There is no guarantee that such an ad hoc scheme will generate sensible output (e.g. non–negative spreads.)

2.2 Statistical Features of D2000–2 and EFX data

Initially, we sample the data with a 20 second frequency. This frequency was chosen to balance two factors. On one hand, one would like to sample very frequently in order to maximize the information content of the data i.e. not omitting large numbers of observations in busy intervals. Second, however,

[4]More information on the structure of the raw data set and the processing of the data is available from the authors on request.

[5]A very simple example would be to take the highest bid and lowest offer from the last 10 EFX quote pairs as the ‘market’ bid and offer at every observation time.

[6]One exception to this is Bollerslev and Domowitz (1993).

sampling very finely leads to observation intervals which are empty (when markets are quiet) especially for the EFX data.[7] The 20 second frequency represents a reasonable compromise between these two considerations. We investigate the impact of changing sampling frequency in Section 4.

2.2.1 Seasonal Patterns

Figure 1 clearly shows the existence of strong intra–day seasonal patterns in D2000–2 and EFX activity. Liquidity supply to D2000–2 is very light in the GMT evening and overnight period, while from 6 to 18 GMT the level of activity is very high. These hours correspond broadly to European and North American trading hours, reflecting the fact that we are examining the USD/DEM, the lack of impact of the D2000–2 system in Asia and the pre–eminence of London as FX trading center. EFX quote frequency correlates very strongly with D2000–2 transaction frequency. Both are high relative to their unconditional means in the periods covering 6 to 10 GMT and 12 to 16 GMT, while both measures are close to zero on average from 18 GMT to 6 GMT.

Certain differences between EFX and D2000–2 emerge in Figures 1(b) and 1(d). In the former, EFX quote frequency during the 10 to 12 GMT period stays relatively high while transaction frequency dips strongly on D2000–2. This is most likely due to the processing limits of the EFX system, i.e. the minimum time between quotes of 2 seconds, which will mask changes in activity at higher frequencies.[8] Figure 1(d) shows that there is essentially no intra–day seasonal pattern in the EFX bid–offer spread. D2000–2 spreads, on the other hand, vary widely across the GMT day. During European and North American trading hours D2000–2 spreads are an order of magnitude lower than those observed in the GMT evening and overnight period. As such, D2000–2 spreads seem to follow a similar intra–day pattern to the U–shape found for spreads observed on many major stock markets (see Foster and Viswanathan (1990) and Biais, Hillion, and Spatt (1995) for results from the NYSE and Paris Bourse.) Hence EFX spreads are much greater than those on D2000–2 during peak D2000–2 trading hours and much lower outside of that period. This implies that using EFX spreads as a measure of liquidity understates true market liquidity throughout the hours from 6 to 18 GMT and overstates liquidity from 18 to 6 GMT.

[7] For example, the minimum time between EFX quote arrivals was 2 seconds, a lower bound imposed by the processing constraints of the EFX system.

[8] The argument here is that the if the input to EFX is such that it is operating at maximum capacity, then marginal changes in quote inputs will not show up in the data which is displayed on screen.

Our analysis of activity patterns indicates that certain market activity statistics derived from the EFX data are likely to misrepresent actual FX market liquidity and trading activity. The EFX spread is particularly problematic. EFX quote frequency on the other hand, shares a fairly similar seasonal pattern to aggregate liquidity demand measures on D2000–2 aside from the interval around mid-day.

2.2.2 Characterizing D2000–2 and EFX Returns

A complementary picture emerges from a statistical analysis of the characteristics of the EFX mid–quote return, D2000–2 transaction price returns and D2000-2 midquote returns.[9]

Table 1 presents the first four moments of the three return series along with their first autocorrelation and a fifth order Box–Ljung statistic. Given the evidence of strong seasonal patterns in activity from Figure 1, the statistics are computed for seven non–overlapping subsamples of the entire data sample for each returns series.[10] It is immediately clear that EFX midquote returns are generally twice as volatile as D2000–2 returns. An exception to this result occurs for the overnight subsample where D2000–2 returns are much more volatile due to the effect of illiquidity. The volatility results are illustrated in Figure 2. There seems to be little useful comparative information in the reported return skew and kurtosis figures. As one might expect, all three return series share similar patterns in the sign of skew across subsamples. Also all return series exhibit excess kurtosis.[11]

Much prior work using EFX data has noted that returns contain a strong negative moving average component. Various explanations have been put forward for this phenomenon, e.g. the effect of idiosyncratic inventory positions of individual dealers, dealers working with different information sets and noise in the EFX data. Analysis of our data demonstrates that the EFX data used here displays negative first order autocorrelation also. Across the entire trading day the first return autocorrelation is -0.33. Further, the Box–Ljung statistics indicate that the hypothesis of a lack of up to fifth order

[9]Goodhart, Ito, and Payne (1996) present comparative statistical results similar to those below. However, they compare the properties of tick-by-tick data from D2000-2 and EFX, rather than data series with common calendar time sampling frequency. Hence, they results are not directly comparable with ours.

[10]Each subsample corresponds to a given two hour interval from every trading day. Hence, each subsample contains 4 breaks (between the last observation in that period on day k and the first on day $k+1$. When appropriate, these breaks are modelled in estimation using dummy variables.

[11]Formal tests for the existence of the fourth moment of returns, using a procedure proposed by Danielsson and de Vries (1997), indicate that it is unbounded in all cases.

autocorrelation can be rejected at the 5% level in all EFX return subsamples. A different picture emerges when we look at the autocorrelations for D2000–2 returns. These are generally between one third and one half the value for the corresponding EFX subsample. Further, D2000–2 return autocorrelation displays a marked, inverted U–shape over the trading day, a pattern which does not appear in the EFX autocorrelations. This implies that D2000-2 return autocorrelation is inversely related to D2000-2 trading activity.

These patterns can be explained with our hypothesis regarding EFX quote making behavior. Assume an agent who wishes to buy dollars. He submits a competitive bid to EFX and computes the ask by adding a large, standard spread. Another agent who, instead, wishes to sell dollars will enter a competitive offer and subtract the standard spread to gain his bid quote. Note that, even if both of these agents agree on the competitive bid and offer, the midquote submitted by the first will exceed that submitted by the second. Then, if agents wanting to trade at the bid and ask arrive randomly to EFX, quote returns will contain (inflated) negative first order autocorrelation. This is exactly the intuition used in the spread estimator of Roll (1984). Furthermore, this hypothesis delivers the result that EFX midquotes should be much more volatile than the 'true' midquotes

A preliminary look at the relationships between the market activity measures derived from the D2000–2 and EFX data is given in Table 2 in the form of contemporaneous cross–correlations. These results mostly confirm the insights from the seasonal patterns. D2000–2 spreads are strongly negatively correlated with D2000–2 liquidity supply measures. D2000–2 transaction intensity correlates far less well with spreads however. As expected, EFX spreads are extremely poorly correlated with all other series while EFX quote frequency is positively correlated with both D2000–2 liquidity and transaction activity. Hence, prior insights are corroborated in that EFX spreads seem to convey little information regarding underlying spot market activity while EFX quote frequency can be regarded as a fair proxy for volumes and liquidity supply.[12]

[12] Using half-hour aggregates of the relevant series, Goodhart, Ito, and Payne (1996) also report a positive correlation between D2000-2 market order frequency and EFX quote frequency.

3 The Inter–Relationship Between D2000–2 and EFX Quotes

The previous section provides a general comparison of the patterns in and sample statistics derived from D2000–2 and EFX activity measures and returns. In this section, we focus solely on the latter. Given the number of existing studies which use EFX returns as proxies for firm quote returns, a clear examination of their similarities and comovements would seem to be important. We analyze the dynamic relationship between the two quote return series using three techniques; return cross-correlations, a cointegrating VAR in returns and a bivariate GARCH specification. Hence we examine not only causality in returns but also return volatility.

3.1 Cross–Correlation Analysis

As a first pass at examining the lead–lag relationships in the two return series we calculated Box–Ljung Q–statistics for the cross correlations between D2000–2 returns and leads and lags of EFX midquote returns.[13] Figure 3 plots the cross correlations between D2000–2 and EFX returns calculated using data from the period 6 to 18 GMT. It is immediately clear that there is strong asymmetry in cross–correlations with those for negative lags (i.e. leads) larger and more significant than those for positive lags. The implication is that D2000–2 returns tend to lead EFX returns (i.e. EFX returns are predictable with D2000–2 returns) while the converse is not true. Given the 20 second sampling of the data, the estimated cross–correlations imply that EFX returns are predictable between 2 and 3 minutes ahead using D2000–2 returns. As the D2000–2 midquote is formed from firm quotes and the EFX midquote is indicative this is in line with intuition.

3.2 Cointegrating VAR analysis

The cross-correlation analysis in section 3.1 gives rough answers to questions of information transmission between EFX and D2000–2 midquote returns. It ignores, however, the fact that the two midquote series are likely to be cointegrated. This cointegration comes from the fact that we are essentially examining two sources of price information for the same asset. In this section

[13]Table 1 indicates that both D2000–2 and EFX returns contain MA(1) components. We filter the MA(1) structure from each sample before constructing the cross–correlations. This prewhitening of returns implies that the estimated cross–correlations can be treated as independent and asymptotically normal with variance T^{-1} under the null of return independence.

we present a bivariate vector error correction mechanism (ECM), suggested by Hasbrouck (1995), which takes account of the cointegration and allows one to conduct appropriate inference regarding the roles played by D2000–2 and EFX in price discovery.

3.2.1 VAR Model

Hasbrouck (1995) proposes a cointegrating vector autoregression (VAR) procedure to determine the relative information content of two different price series for the same financial asset. The basic empirical model is;

$$r_t = \mu + \alpha z_{t-1} + \beta(L) r_{t-1} + \epsilon_t, \quad E(\epsilon_t \epsilon_t') = \Omega \tag{1}$$

where $r_t = (r_t^{D2}, r_t^{EFX})'$, is a vector containing the D2000–2 and EFX midquote returns, $\mu = (\mu_1, \mu_2)'$, $\alpha = (\alpha_1, \alpha_2)'$, $\epsilon = (\epsilon_{1t}, \epsilon_{2t})'$ and $\beta(L)$ is a conformable polynomial in the lag operator. Equation (1) is just an error correction representation for the pair of return series where we assume that the difference between the two midquote variables (denoted $z_t = q_t^{D2} - q_t^{EFX}$) is $I(0)$. This assumption guarantees that the two price series cannot permanently diverge and is shown to be valid through a series of unit root tests on the midquotes and the difference between them.[14] We refer to this difference as the pricing error from now on. In equilibrium the EFX and D2000–2 midquotes are identical (i.e. $z_t = 0$.)

One measure of each system's contribution to price discovery can be gained through comparison of the α coefficients i.e. how each series reacts to pricing errors. Given the way in which we constructed z_t one would expect α_1 to be negative and α_2 positive, but asymmetry in their sizes will indicate one system reacting more to deviations from equilibrium than the other and hence an asymmetry in price discovery contributions (see also Harris, McInish, Shoesmith, and Wood (1995).) If one of the two quote series dominates completely in terms of information assimilation the α coefficient for that return series is zero.

To calculate the information share of each source of USD/DEM quotes explicitly we first compute the vector moving average representation of the system through inversion of equation (1).

$$r_t = \Phi(L)\epsilon_t \tag{2}$$

[14] ADF test statistics for both quote series exceed -2.5 and for both return series are less than -40. Hence we conclude that the quote series are I(1). The ADF test statistic for the difference between the quote series is -17.1, implying that the quote series are CI(1,1).

The long run impacts of the pair of innovations on both quote series are summarized in the value of $\Phi(1)$. Note that, due to the cointegration between quotes, $\Phi(1)$ must contain two identical rows. Denote a row of $\Phi(1)$ by ϕ.

By assuming that the VAR innovations are uncorrelated, the calculation of the information share of each price measure is straightforward. The information share for price measure j is calculated as;

$$S_j = \frac{\phi_j^2 \Omega_{jj}}{\phi \Omega \phi} \tag{3}$$

However, the assumption that innovations are uncorrelated is unlikely to be satisfied in practice. Hence we must modify the preceding estimator. In order to bound the information share we use a Choleski decomposition of the variance–covariance (VCV) matrix. An upper bound on the information share of system i is obtained from this decomposition when the innovation to equation i is represented in the first row of the VCV matrix. A lower bound on i's information share is obtained when system i is represented in the second row of the VCV. Denoting the Choleski factor of Ω by F, the information share of series j is then;

$$S_j = \frac{([\phi F]_j)^2}{\phi \Omega \phi} \tag{4}$$

where $[\phi F]_j$ is the j^{th} element of the given row vector. In the empirical results which follow, we present upper and lower bounds on the D2000–2 information share.

3.2.2 Empirical Estimates

The properties of the pricing error (z_t) are presented in Table 3, broken down across intervals of the trading day. Several facts are immediately apparent. First, the difference between the two midquote series is small on average. Second, examination of Table 3 in conjunction with Figure 1 shows that the variance of the pricing error covaries negatively with D2000–2 trading volume. In times of heavy D2000–2 activity D2000–2 and EFX midquotes stay closer together on average. Finally, the dependence in pricing errors, measured by the first order autocorrelation and a fifth order Box–Ljung statistic, is also inversely related to D2000–2 volume. Hence, when trading volume is large, the pricing error is less persistent, implying that deviations from equilibrium are removed more speedily. Also, this final result suggests that the ECM structure in equation (1) will not be stable across the trading day and we therefore estimate the ECM separately for each of our trading day subsamples.

Table 4 contains the key ECM parameters for all seven trading day subsamples plus the estimated upper and lower bounds on the D2000–2 information share for each. The crucial ECM parameters are α_1 and α_2 and we expect the former to be negative while the latter should be positive. For all of the trading day subsamples, the sign of α_2 is as expected i.e. when the D2000–2 midquote exceeds the EFX midquote, the EFX midquote adjusts upwards. In only 4 of the 6 subsamples does α_1 take the expected sign.

With regard to the relative size and significance of the parameters on the pricing error, a clear asymmetry is visible. For the subsamples covering 6 to 18 GMT, α_2 is always an order of magnitude larger than α_1 and is always highly significant. On the other hand, the lagged equilibrium error is only significant in 3 of 6 D2000–2 return regressions. Further, the R^2 for the EFX return equation is always an order of magnitude greater than that for the D2000–2 equation. These results imply that the majority of any dis–equilibrium in the system is removed through the adjustment of EFX quotes. D2000–2 quotes react very weakly to dis–equilibrium.

Results are very different for the overnight subsample (6pm to 6am.) In this case the parameters on the lagged equilibrium error have similar size and significance level. This is likely due to the lack of activity on both D2000–2 and EFX during this period.

A final point regarding the ECM results relates to those parameters which have been omitted from Table 4, specifically the coefficients on lagged EFX and D2000–2 returns in each equation.[15] For the D2000–2 return equation, neither lagged own returns nor lagged EFX returns are significant in general. On the other hand, both sets of lagged returns are significant in the EFX equation with lagged D2000–2 returns generally having a positive effect and lagged EFX returns a negative impact on current EFX returns.[16] Note that as equilibrium variation in the EFX quote occurs through the coefficient on z_{t-1} this predictability reflects inefficiency in the EFX quote setting process.

Finally, the last two columns of Table 4 presents the estimated lower and upper bounds on the D2000–2 information share for each subsample. It is clear that, for those hours of the day when D2000–2 is active (i.e. 6am to 6pm), it is the dominant location for price discovery with a minimal information share in excess of 85%. For the overnight period, however, the lack of D2000–2 activity implies that its share drops to around 35%. Within the effective trading day, the information share on D2000–2 follows an inverted

[15] Full results are available upon request from the authors.

[16] These results are also consistent with our autocorrelation analysis from Section 2.2.2.

U–shape, such that the information share is minimized in the early GMT morning and early GMT evening.

3.3 Bivariate GARCH Analysis

Below we examine the inter-relationships between D2000–2 and EFX return volatility, similar to our analysis of lead-lag relationships in returns, but using a bivariate GARCH specification. We use the residuals from the VAR specifications of section 3.2.2 in estimation so as to remove any covariation between the series which arises from the previously estimated lead–lag relationships in returns themselves.

In their most general forms, multivariate GARCH models are notoriously hard to estimate and, for practical purposes, a restricted version can often be used without detriment to the application. Hence, we follow prior work and employ the BEKK specification introduced in Engle and Kroner (1995). In a bivariate setting, the structure of the BEKK model is as follows;

$$(5) \qquad r_t = \Sigma_t \epsilon_t, \quad \Sigma_t = V'V + B' r_{t-1} r'_{t-1} B + A' \Sigma_{t-1} A$$

$r_t = (r_t^{D2}, r_t^{EFX})'$, i.e. the vector of D2000–2 and EFX midquote returns, ϵ_t is a 2×1 vector of NID(0,1) innovations and Σ_t is the conditional variance–covariance matrix at t.

This specification has the advantage of being both parsimonious while capturing the relevant dynamics, and ensuring positive definiteness of Σ_t. Each matrix, V, A, B is 2×2 and V is restricted to be upper triangular. Hence we have 11 free parameters in the model.

Estimated parameters from the BEKK specification for residual returns are given in Table 5.[17] From these parameters it is clear that both volatility series are strongly autocorrelated, a result which is in line with those from univariate GARCH specifications on intra–day FX data. The diagonal elements of A and B are all positive and strongly significant. An interesting feature of the results appears in the off diagonal elements. The upper right coefficients of A and B (i.e. a_{12} and b_{12}) are greater in magnitude and more significant than a_{21} and b_{21}. This implies that D2000–2 volatility affects EFX return volatility to a greater degree than in the converse direction. A further result which agrees with those from univariate volatility models (and

[17] Previous research on intra–day FX volatility has shown that the intra–day seasonal patterns in volatility can bias estimated volatility process coefficients. Hence in Table 5 we also present estimated coefficients from the bivariate GARCH using residual returns with deseasonalised volatility. Results on volatility spillovers from the two tables are qualitatively similar.

the prior kurtosis measures for returns) is that the coefficients in Table 5 imply that the unconditional variance–covariance matrix does not exist, a multivariate equivalent of IGARCH.[18]

Hence, these results imply that volatility spillovers between the two series, tend to be in the direction from real to indicative quotes.

4 Temporal Aggregation

The preceding analysis was based on a 20 second sampling of the data. However, most previous work on EFX has used much coarser sampling frequencies e.g. 5 or 10 minutes. The reason for that may be suspicion regarding the quality of EFX quotes at extremely high frequencies. Below we examine how the statistical properties of D2000–2 and EFX returns alter as sampling frequency is varied from 20 seconds to 10 minutes.[19] We consider the following comparative measures;

1. Return variance.
2. Return autocorrelation.
3. Cross-correlation of return series.
4. Q–statistics for lead and lag cross-correlations.
5. Pricing error autocorrelation.

We report the results from these investigations in Figure 4.

The autocorrelation evidence is presented in Figure 4(a). The main features of this plot are that the D2000–2 autocorrelation is significant only for frequencies less than or equal to 180 seconds and is always closer to zero than the EFX coefficient. In addition, the EFX coefficient is significant until we reach a sampling frequency of 5 minutes.

Figure 4(b) plots scaled variance measures for returns across sampling frequencies.[20] Here we see that (scaled) D2000–2 volatility is fairly constant under aggregation whilst (scaled) EFX volatility decreases dramatically. At a

[18] The unconditional covariance matrix is given by $\Omega = (I - [A \otimes A]' - [B \otimes B]')^{-1} vec(V'V)$.

[19] Throughout this section, data from the overnight period from 18GMT to 6GMT is omitted due to lack of activity on both systems.

[20] The variance measures are scaled by the degree of aggregation such that we plot the 20 second return variance, one third of the one minute return variance, one sixth of the 2 minute variance etc.

5 minute sampling frequency the difference between the two scaled variances is small, and this difference has almost disappeared at a 10 minute sampling frequency.

This evidence is reinforced in Figures 4(c) to 4(e) which show the contemporaneous cross–correlation of the two return series, lead and lag Q–statistics and pricing error autocorrelations across sampling frequencies. At very high frequencies (i.e. lower than or equal to 5 minutes) EFX returns are predictable with lagged D2000–2 returns and the pricing error is strongly autocorrelated. Such effects disappear with aggregation such that at a ten minute sampling the first order autocorrelation of the pricing error is less than 0.1 and the contemporaneous cross–correlation between D2000–2 and EFX returns exceeds 0.90.

The implication of this analysis is straightforward. The bias in EFX data is strongest at very high frequencies, but disappears with data aggregation. At a 10 minute sampling frequency, the indicative quotes have statistical properties much like those of firm quotes.

5 Conclusion

In recent years, a large literature on empirical FX microstructure has emerged. Most of this work has been conducted using data derived from indicative EFX quotes due to the fact that firm quotes and/or transaction prices have been unavailable at sufficiently high frequencies. Researchers have however remained skeptical of the accuracy of EFX quotes. Perhaps for this reason, EFX data have generally been employed at relatively coarse sampling frequencies, e.g. 5 or 10 minutes. Our analysis vindicates this practice. Based on a 20 second data sampling we derive the following results;

1. EFX midquote returns are excessively volatile when compared to returns in the midquote derived from the best D2000–2 limit buy and sell prices.

2. While EFX returns are strongly negatively autocorrelated, D2000–2 returns have much lower negative autocorrelation.

3. Bid–ask spreads in D2000–2 data exhibit the expected intra day U pattern, while EFX spreads do not.

4. Quote frequency on the EFX system and D2000–2 transaction frequency are fairly strongly positively correlated.

5. News is first incorporated into D2000–2 quotes, and impacts EFX quotes over the next 2 to 3 minutes. This implies that D2000–2 returns forecast EFX returns.

6. On average, over 90% of price relevant information enters the D2000-2 midquote first, and from there is impounded in the EFX midquote

7. D2000–2 volatility affects EFX volatility. The converse is not true.

These results can be explained by the following model of trader behavior. Most agents submitting a quote to EFX only really wish to deal on one side of the market, the sell side for example, and hence place an approximately competitive ask quote. To complete the quote a standard spread is subtracted to get the bid. This implies that spreads contain little or no information, quote returns will be excessively volatile and will display inflated first order autocorrelations. By this hypothesis, traders consider the EFX screens as a forum for advertising, signaling willingness to buy/sell, whilst trades are conducted directly possibly at different prices.

Our final results will be comforting to those that have used EFX return data with more coarse sampling frequencies. Measurement errors in EFX disappear with aggregation. Most anomalies in EFX returns have disappeared at a five minute sampling frequency and all are insignificant at a ten minute sampling.

References

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

Baillie, R., and T. Bollerslev (1991): "Intra-Day and Inter-Market Volatility in Foreign Exchange Markets," *Review of Economic Studies*, 58, 565–586.

Bank for International Settlements (1996): *Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 1995*. B.I.S., Basle.

——— (1998): *Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 1998: Preliminary Global Data*. B.I.S., Basle.

Biais, B., P. Hillion, and C. Spatt (1995): "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50, 1655–1689.

Bollerslev, T., and I. Domowitz (1993): "Trading Patterns and Prices in the Interbank Foreign Exchange Market," *Journal of Finance*, 48, 1421–1444.

Bollerslev, T., and M. Melvin (1994): "Bid-ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis," *Journal of International Economics*, 36, 355–372.

Curcio, R., C. Goodhart, D. Guillaume, and R. Payne (1997): "Do Technical Trading Rules Generate Profits? Conclusions from the Intra-Day Foreign Exchange Market," *International Journal of Finance and Economics*, forthcoming.

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

Danielsson, J., and C. G. de Vries (1997): "Beyond the Sample: Extreme Quantile and Probability Estimation," London School of Economics, FMG DP 298.

de Jong, F., R. Mahieu, and P. Schotman (1998): "Price Discovery in the Foreign Exchange Market," *Journal of International Money and Finance*, forthcoming.

Engle, R., and K. Kroner (1995): "Multivariate Simultaneous Generalized ARCH," *Econometric Theory*, 11(1), 122–150.

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

Goodhart, C., T. Ito, and R. Payne (1996): "One Day in June, 1993: A Study of the Workings of Reuters D2000-2 Electronic Foreign Exchange Trading System," in *The Microstructure of Foreign Exchange Markets*, ed. by J. Frankel, G. Galli, and A. Giovannini. University of Chicago Press.

Harris, F., T. McInish, G. Shoesmith, and R. Wood (1995): "Cointegration, Error Correction and Price Discovery on Informationally Linked Securities Markets," *Journal of Financial and Quantitative Analysis*, 30, 565–579.

Hartmann, P. (1996): "Trading Volumes and Transaction Costs in the Foreign Exchange Market: Evidence from Daily Dollar-Yen Spot Data," Discussion Paper 232, Financial Markets Group, London School of Economics.

Hasbrouck, J. (1995): "One Security, Many Markets: Determining the Contributions to Price Discovery," *Journal of Finance*, 50, 1175–1199.

Melvin, M., and X. Yin (1996): "Public Information Arrival, Exchange Rate Volatility and Quote Frequency," Working Paper, Arizona State University.

Payne, R. (1996): "Announcement Effects and Seasonality in the Intraday Foreign Exchange Market," Discussion Paper 238, Financial Markets Group, London School of Economics.

Roll, R. (1984): "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 39, 1127–1139.

Table 1: Summary Statistics for 20 Second Return Subsamples

Data	Mean	Var.	Skew	Kurt.	ρ_1	$Q(5)$
D2000–2 Midquote						
6-8 GMT	0.00011	0.00014	-0.04	10.66	-0.14	63.28
8-10 GMT	-0.00001	0.00008	0.28	19.48	-0.10	37.65
10-12 GMT	-0.00050	0.00019	-1.63	31.53	0.03	35.49
12-14 GMT	0.00011	0.00020	-0.82	20.66	-0.15	59.90
14-16 GMT	0.00028	0.00015	0.07	6.76	-0.19	76.75
16-18 GMT	-0.00024	0.00023	0.12	17.76	-0.33	205.65
18-6 GMT	0.00003	0.00075	-0.01	124.71	-0.26	742.01
EFX Midquote						
6-8 GMT	0.00006	0.00046	-0.52	14.29	-0.37	263.69
8-10 GMT	-0.00002	0.00034	0.06	6.46	-0.37	257.00
10-12 GMT	-0.00050	0.00055	-1.42	34.45	-0.27	144.19
12-14 GMT	0.00013	0.00053	-0.13	6.75	-0.34	216.55
14-16 GMT	0.00029	0.00052	0.11	10.57	-0.41	318.37
16-18 GMT	-0.00022	0.00055	-0.24	14.06	-0.25	122.17
18-6 GMT	0.00002	0.00032	0.13	33.49	-0.15	442.98
D2000–2 Prices						
6-8 GMT	0.00010	0.00014	-0.11	13.63	-0.06	17.24
8-10 GMT	-0.00001	0.00009	0.21	8.40	-0.09	38.39
10-12 GMT	-0.00050	0.00024	-1.79	34.26	0.02	23.29
12-14 GMT	0.00010	0.00026	-0.34	13.83	-0.19	83.21
14-16 GMT	0.00029	0.00017	0.31	9.35	-0.07	19.93
16-18 GMT	-0.00025	0.00012	-1.33	32.18	-0.06	22.78
18-6 GMT	0.00002	0.00037	2.03	727.06	-0.15	393.02

Notes: the first our columns of the table give the sample mean, variance, skewness and kurtosis of returns. The next column gives the first order return autocorrelation and the final column a fifth order Box-Ljung statistic for return dependence. The 5% critical value for the Box-Ljung statistic is 11.07.

Table 2: Cross-correlations of D2000-2 and EFX Activity Variables

Data	Correlation						
D2 spread	1	-	-	-	-	-	-
D2 orders	-0.50	1	-	-	-	-	-
D2 size	-0.47	0.95	1	-	-	-	-
EFX spread	0.05	-0.04	-0.05	1	-	-	-
EFX quotes	-0.26	0.43	0.37	-0.01	1	-	-
D2 deals	-0.13	0.29	0.26	-0.00	0.38	1	-
D2 volume	-0.11	0.27	0.24	-0.01	0.34	0.94	1

Notes: spreads are measured in percentage terms. Orders and size refer to the number of outstanding D2000-2 limit orders and their aggregate size respectively. EFX quotes refers to the number of EFX quotes posted in each interval. Deals is a count of the number of transactions in each interval and Volume is the aggregate transacted volume in a given interval.

Table 3: Properties of Price Error between D2000-2 and EFX

Subsample	$\bar{z}$	σ_z^2	ρ_{1z}	$Q(5)$
6-8 GMT	-0.00088	0.00040	0.36*	386.8*
8-10 GMT	0.00030	0.00030	0.34*	405.9*
10-12 GMT	0.00134	0.00074	0.56*	1256.1*
12-14 GMT	0.00213	0.00040	0.23*	172.9*
14-16 GMT	-0.00141	0.00044	0.31*	391.0*
16-18 GMT	-0.00478	0.00093	0.59*	1445.8*
18-6 GMT	-0.00646	0.0069	0.92*	37706.3*

Notes: z_t is the price error at t. ρ_{1z} is the first autocorrelation of the pricing error. $Q(5)$ is the 5th order Box-Ljung statistic for the pricing error. The 5% critical value for the Box-Ljung statistic is 11.07.

Table 4: Cointegrating VAR Results for D2000-2 and EFX Returns

Subsample	Lags	α_1	t(α_1)	R^2_{D2}	α_2	t(α_2)	R^2_{EFX}	D2$_{min}$	D2$_{max}$
6am to 8am	2	-0.05	-2.10	0.03	0.55	9.35	0.34	90.5	91.6
8am to 10am	1	-0.03	-1.77	0.02	0.56	16.26	0.33	94.7	95.8
10am to 12pm	1	0.09	1.65	0.03	0.48	3.93	0.34	93.0	99.8
12pm to 2pm	1	0.06	0.20	0.02	0.72	22.48	0.41	92.9	1.00
2pm to 4pm	1	-0.05	-3.05	0.04	0.51	11.82	0.34	89.9	90.6
4pm to 6pm	2	-0.05	-2.43	0.14	0.27	8.40	0.17	86.0	88.8
6pm to 6am	7	-0.03	-4.74	0.10	0.01	5.34	0.08	35.4	35.7

Notes: α_1 and α_2 are the coefficients on the lagged pricing error in the D2000-2 and EFX return equations respectively. The columns immediately following give the t-values for these coefficients. The R^2 for the D2000-2 and EFX return equations are also presented and, finally, we present the lower and upper bounds on the D2000-2 information share.

Table 5: Bivariate GARCH Results: 20 Second Data

	Raw Data		Deseasonalized Data	
Parameter	Coeff	t-value	Coeff	t-value
v_{11}	1.43×10^{-5}	38.37	0.18	19.87
v_{12}	1.33×10^{-5}	6.02	-0.05	-0.49
v_{22}	5.14×10^{-5}	52.95	0.15	2.34
a_{11}	0.96	1053.18	0.89	112.39
a_{12}	0.08	37.70	0.54	33.16
a_{21}	-0.02	-15.64	0.17	6.26
a_{22}	0.90	367.68	-0.89	-103.58
b_{11}	0.30	88.28	0.34	95.19
b_{12}	-0.14	-22.51	-0.05	-12.68
b_{21}	0.02	6.39	-0.01	-2.18
b_{22}	0.33	66.59	0.27	54.05

Notes: estimated coefficients from the bivariate GARCH specification from residual return data. estimates of the bivariate GARCH parameters from the residual return data with deseasonalised volatility.

Figure 1: Intra-day Seasonal Patterns

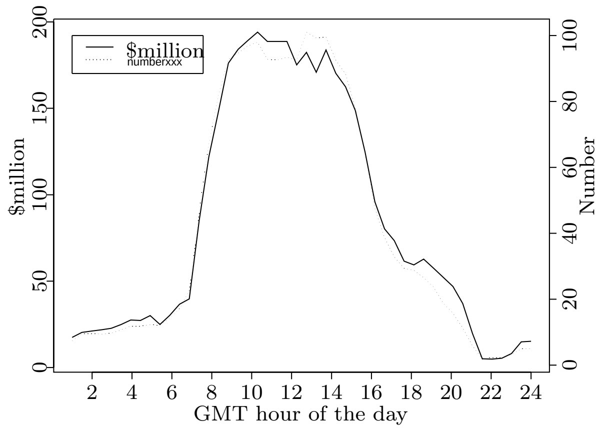

(a) Aggregate Size ($m) and Aggregate Number of D2000-2 Limit Orders Outstanding

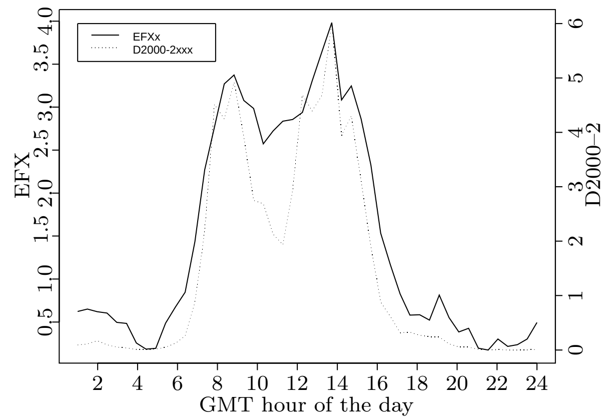

(b) EFX Quotation Frequency and D2000–2 Market Order Frequency

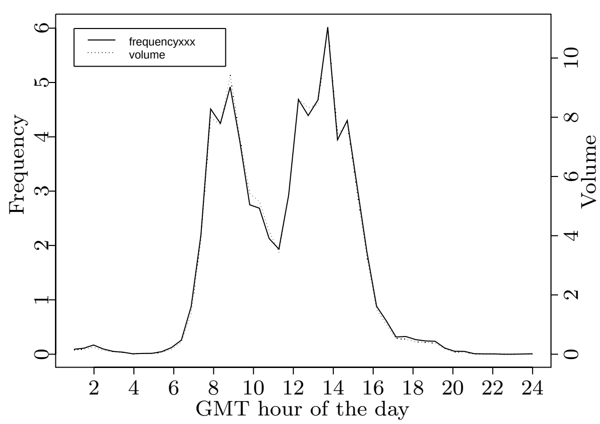

(c) D2000-2 Market Order frequency and volume ($m.)

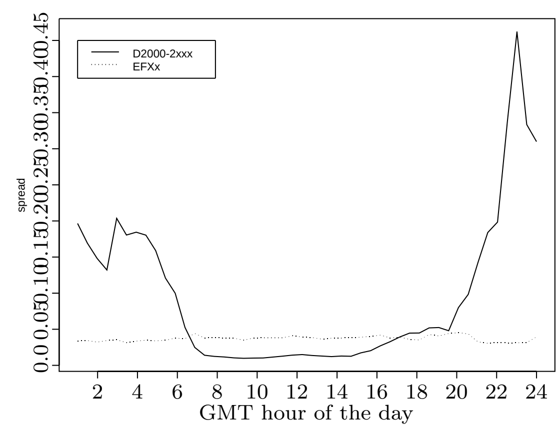

(d) D2000-2 and EFX Spread

Notes: These figures are based upon data sampled once every 20 seconds.

Figure 2: Intra-day Pattern in Return Variances

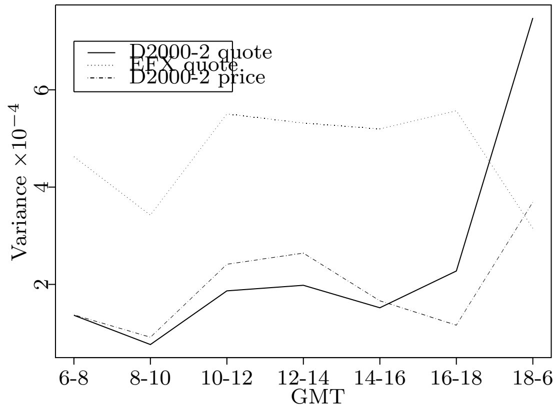

Figure 3: Cross Correlations Between D2000–2 and EFX Returns (with 95% Confidence Band)

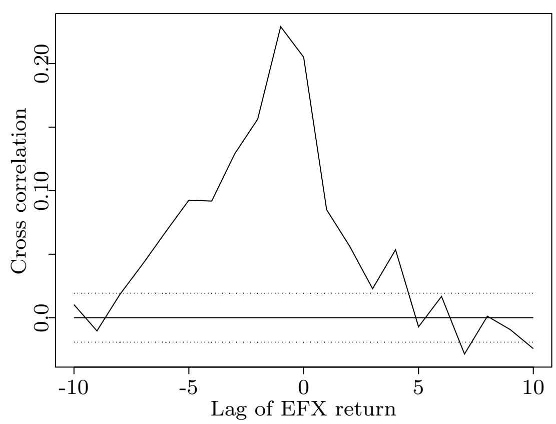

Notes: The graph plots the correlation between D2000–2 returns and lagged EFX returns. A negative lagged EFX return is an EFX lead.

Figure 4: EFX and D2000–2 Returns at Varying Sampling Frequencies

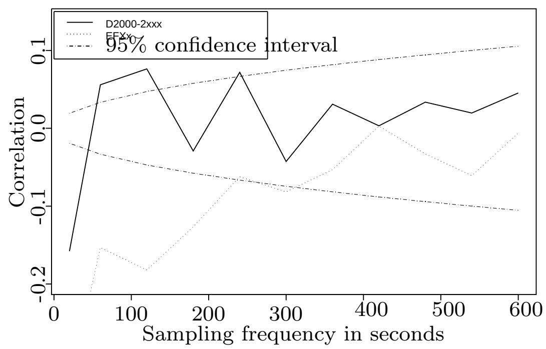

(a) First order autocorrelation

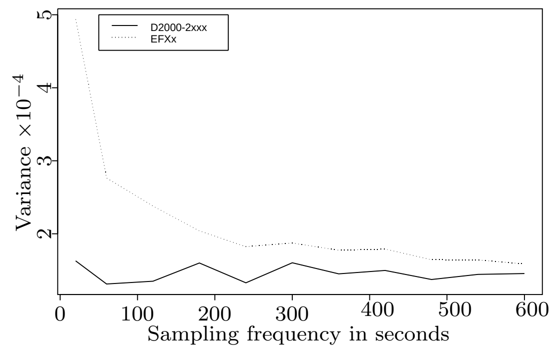

(b) Variance

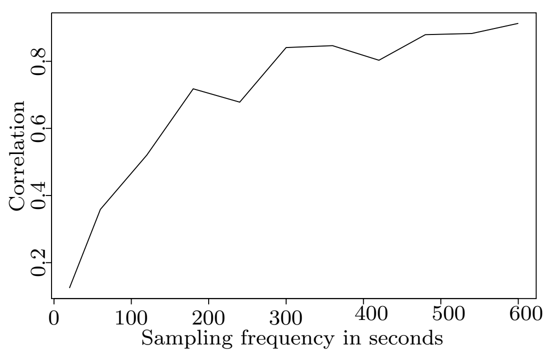

(c) Correlation between D2000–2 and EFX

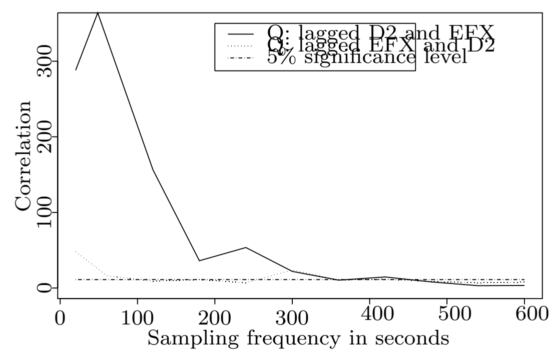

(d) Q statistics for cross correlations

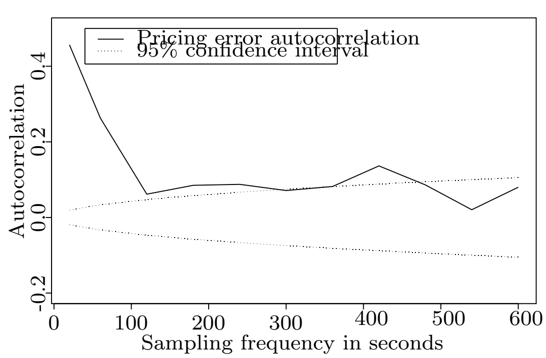

(e) Pricing error autocorrelation