Order execution dynamics in a global FX spot market

Edward Howorka, Eugene Nagirner, Anatoly B. Schmidt

Business Risk, ICAP Electronic Broking
One Upper Pond, Bldg F, Parsippany NJ 07054 USA

Expected execution times for the limit EUR/USD orders in the global FX spot market are presented. The execution time estimates are given for different order sizes and for varying distances between the order price and the market best price. These results are used in a simple value-at-risk based theory for trading large amounts by slicing them into smaller orders. Finally, the current discussion of the long-memory order flows in equity markets is expanded into the global FX market.
.

1. Introduction

Modern inter-bank spot foreign exchange transactions are primarily conducted via two electronic broking systems, EBS and Reuters. EBS dominates the EUR/USD and USD/JPY exchange. In 2006, the daily transacted volume in the EBS market has exceeded 120 billion USD. As a result, EUR/USD and USD/JPY rates posted on the EBS trading screens have become the reference prices quoted by dealers to their customers worldwide [1].

Present empirical research of the high-frequency FX markets is overwhelmingly based on the Reuters indicative rates [2]. The disadvantages of such rates, as compared to the "firm" rates at which the inter-bank currency exchange is conducted, are well documented [3]. Thus, in recent years, several high-frequency FX market studies based on the EBS data have been reported [1, 4 – 7].

Understanding order execution dynamics is especially important for the back-testing of trading strategies. Indeed, in the known publications, it is assumed that when a trading model generates a signal to trade, the order is executed instantly at the current market price (see e.g. [8 - 10] and references therein). While this assumption may be sufficient for analysis of low-frequency trading (e.g. for daily returns), it is not valid in the inter-bank high-frequency market where only limit orders are accepted and the bid-ask spread plays an important role in the realized P/L.

This report is structured as follows. A description of EBS market specifics is given in the next section. The work's main empirical findings – the expected execution times for EUR/USD orders of various sizes and distances from market best prices – are listed in Section 3. Furthermore in Section 4, we offer a simple value-at-risk based theory for trading large amounts by slicing them into smaller orders. Finally, we expand the discussion of the long-memory order flow processes in equity markets [11, 12] into the global FX market (Section 5).

2. The EBS FX spot market

Several specifics of the EBS system are relevant for this work. Firstly, the system only accepts the limit orders. In other words, traders must always specify the price at which they are willing to execute an order. The EBS system distinguishes between two types of orders: *quotes* and *hits*. Quotes stay in the order book until they are filled or canceled; hits are automatically canceled if they have no matching counterparty when they reach the market. Hence, a hit is always a taker order while a quote may be either a maker or a taker order. Namely, when one quote matches another, the maker quote is that quote which arrived in the market earlier while the taker order is that which arrived later.

An important feature of the EBS market is that trading can only be conducted between those counterparties that have bilateral credit. Every EBS customer establishes a credit limit with all other EBS customers and can change it at any time. This implies that the *EBS best prices* (highest bid and lowest offer) may or may not be available to an EBS customer, depending on whether this customer has bilateral credit with the maker(s) of the best prices.

Orders in the EBS market are submitted in units of millions of the base currency (the first currency in the name of the currency pair, e.g., USD in USD/JPY and EUR in EUR/USD).

3. Order execution times

There may be different motivations behind the canceling of unfilled or partially filled orders. It may be a strategy that avoids trading at best prices, or trading setup that uses a default amount for all orders regardless of the target amount. It may be a firm belief that price will not revert within the accepted time horizon (e.g. due to some news) or may merely be the trader's lack of patience. In this work, we have chosen to consider only those orders that were fully executed.

For statistical analysis, we use EUR/USD quotes submitted within the six-week interval of Monday, 24-Apr-2006 through Friday, 2-Jun-2006 during the most liquid weekday time period of 7:00 – 17:00 GMT.

Let P, BB, and BO be the order price, the best bid, and the best offer, respectively. We define the bid-side distance from best price as $D = BB - P$ and the offer-side distance as $D = P -$ BO. The values of D are measured in pips (0.0001 USD). A typical weekly distribution of the number of orders at different distances from best price is presented in Fig.1 for each quote size studied. This distribution is clearly skewed to positive values of D as limit orders are more often placed inside the order book in expectation of a favorable price move. We find that our data sample is sufficient for reliable statistical analysis for orders placed at $|D| \leq 2$. The sample size for $|D| \leq 2$ is given in Table 1.

Order size, millions EUR	**1**	**5**	**10**	**20**
Total order count	1882388	176965	46542	8539
Completely filled	584323	85464	17845	2641
Percent of completely filled	31.0	48.3	38.3	30.9

Table 1. The data sample for $|D| \leq 2$ used in this work: EUR/USD orders, weekdays, 7:00 – 17:00 GMT, 24-Apr-2006 to 2-Jun-2006.

Percentage of completely filled orders is lowest for both the smallest (1M) and highest (20M) order sizes considered in this work. The latter minimum is not surprising due to the price volatility, however the minimum at 1M is not as obvious. Traders in the EBS market primarily submit orders for two reasons. Firstly, they execute customer (e.g. importer/exporter) orders. In this case, timely execution is imperative, even at the expense of bid/ask spread or more. Secondly, traders look for various arbitrage opportunities. We think that the latter type of trades is done with relatively small orders and high level of risk aversion.

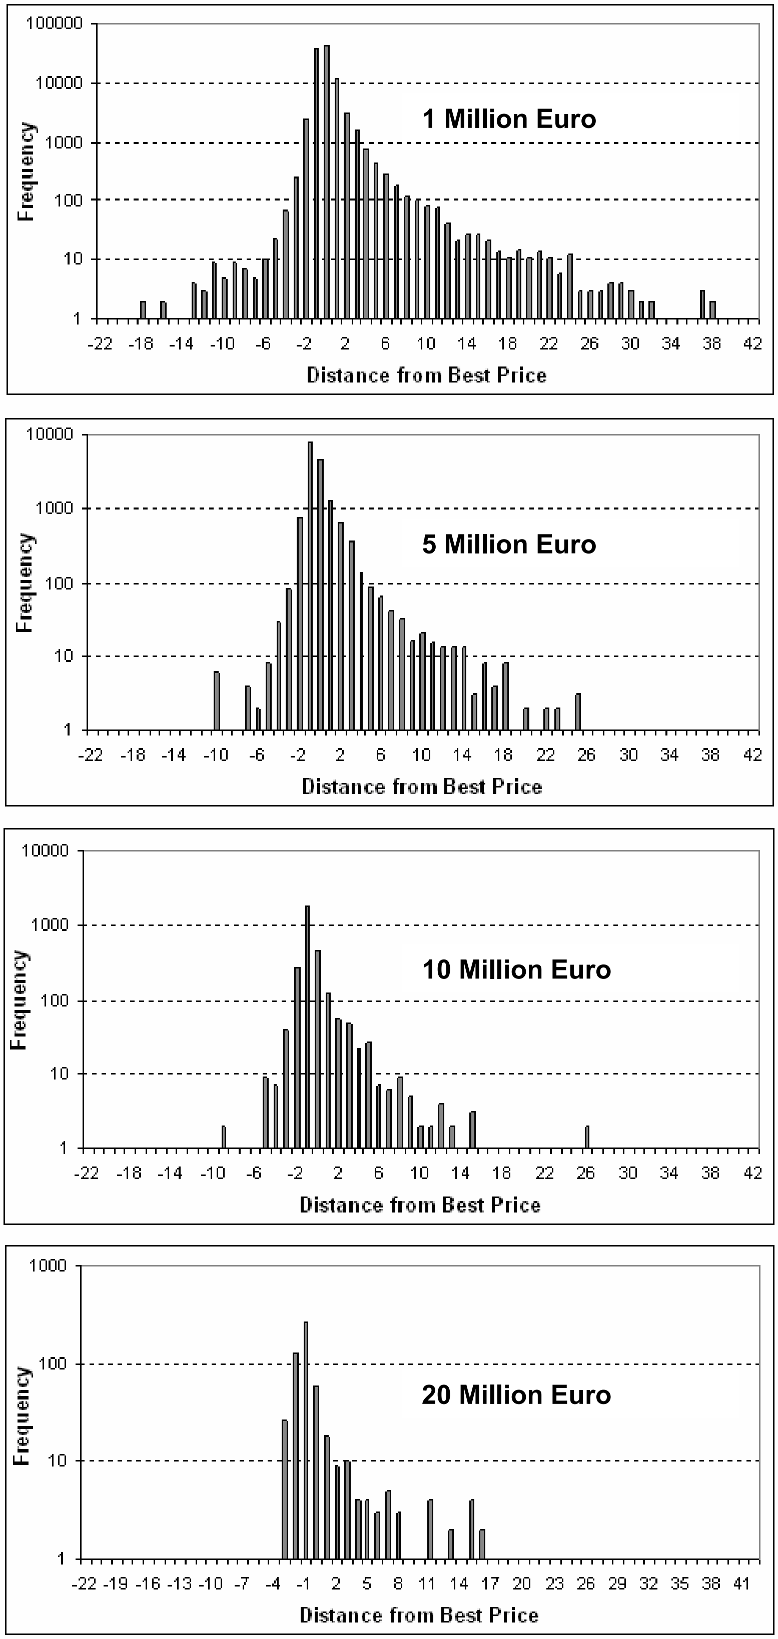

Fig. 1: Price distribution for EUR/USD orders. Data for 7:00 – 17:00 GMT, week ending 28-Apr-2006.

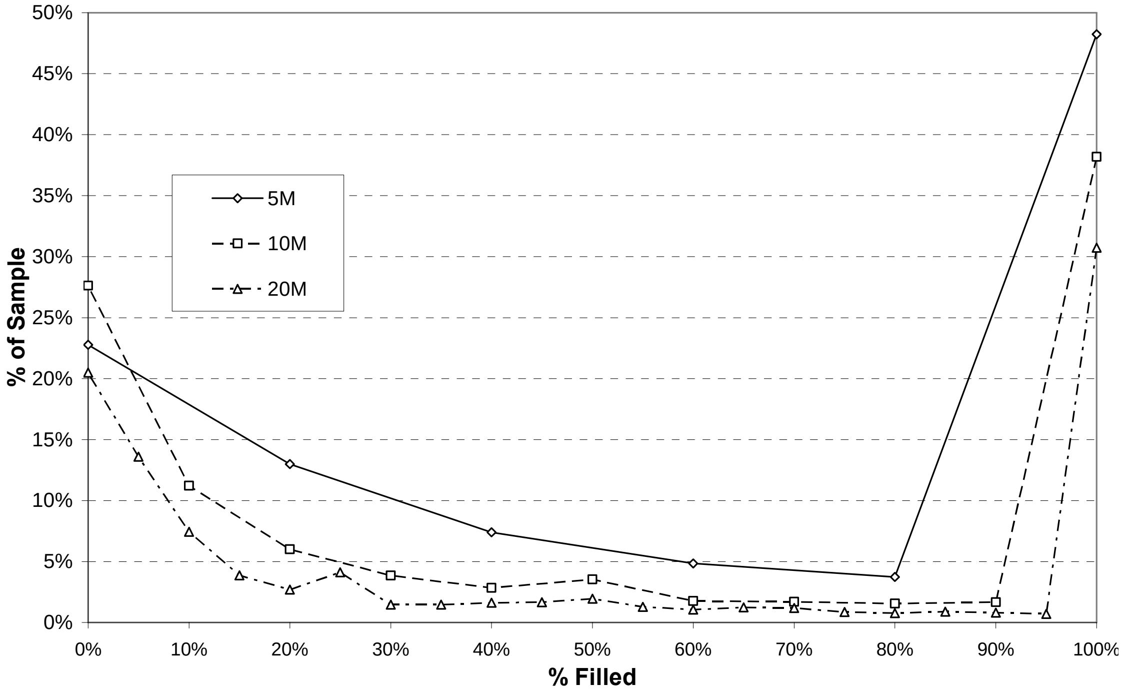

Fig.2: Percentage of quotes with different fill levels submitted within $|D| \leq 2$ (data for EUR/USD, 7:00 - 17:00 GMT, 4/24/06 - 6/02/06).

The distribution of quotes with different fill levels has two pronounced maxima at 0% and 100% (see Fig.2). This distribution shows that traders have limited patience while waiting for their first match. Yet once filling has started, they are inclined to wait until the order is completely filled.

The expected execution times *T(V, D)* for the EUR/USD orders as a function of order sizes, *V*, and distances from best price, *D*, are given in Table 2. We calculate *T*(*V*, *D*) as the arithmetic average of execution times for all quotes of size *V* and distance *D*. Since the minimal order size is $V = 1$ and the EUR/USD bid/offer spread during liquid hours rarely exceeds two pips, one may expect that *T*(1, -2) is very close to the system's processing time of an order match. On the other hand, *T*(1, -1) is close to the filling time of a 1M order at the top of the order book. We use the notion of *execution waiting time* to characterize the expected time for an order to reach the top of the order book after joining the best price order queue. We estimate this execution waiting time as the difference [*T*(1, 0) – *T*(1, -1)] $\approx$ 12 sec. The execution waiting time does not depend on the quote size; however, we have chosen 1M quotes for its estimation due to their prevalence in our data sample. It should be noted also that in the EBS market, an order may be taken before it reaches the top of the order book if the taker has no bilateral credit with makers closer to the top best quote.

T(*V*, *D*)	**Distance from Best Price (*D*), pips**				
Order size (*V*), millions Euro	**-2**	**-1**	**0**	**1**	**2**
1	0.4	1.3	13.5	44.1	108.3
5	0.6	2.8	16.6	58.9	104.6
10	1.0	4.7	18.5	77.0	164.6
20	1.8	6.4	21.4	124.9	136.8

Table 2. Expected execution time, *T(V, D),* (in seconds) for EUR/USD orders. Data for weekdays, 7:00 – 17:00 GMT, 24-Apr-2006 to 2-Jun-2006.

The expected execution times, *T(V, D),* in Table 2 relate to a EUR/USD one-second return volatility of about 0.5 pips for the entire investigated time interval. It should be noted that the order execution time is very sensitive to the market volatility. An example in Fig. 3 illustrates a sharp drop in the execution waiting time on both sides of the market when price jumps and volatility increases. Our estimates for shorter sub-intervals of time further show that execution waiting time grows linearly with inverse return volatility (see Fig. 4).

The results in Table 2 demonstrate a trade-off between the order execution speed and the trade P/L. Indeed, if a trader is willing to incur losses, he can execute a trade within one second by submitting an order with $D < 0$. On the other hand if a trader has enough patience, he has a chance to gain from submitting an order with $D > 0$.

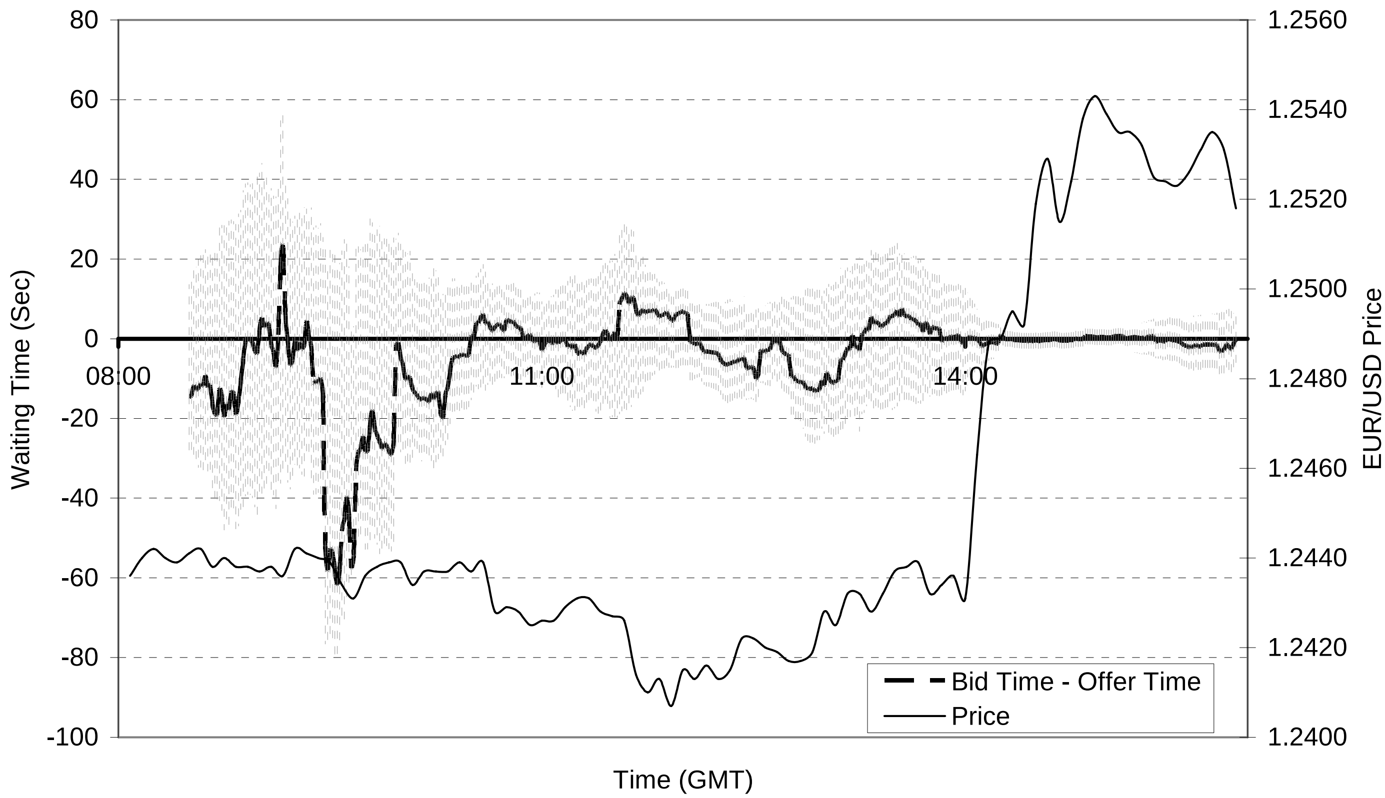

Fig. 3: An example of the change in execution time associated with a sharp price change on 27-Apr-2006. The dashed line is the difference between the bid-side and the offer-side execution times. The error bars stemming up (down) are the offer (bid) times.

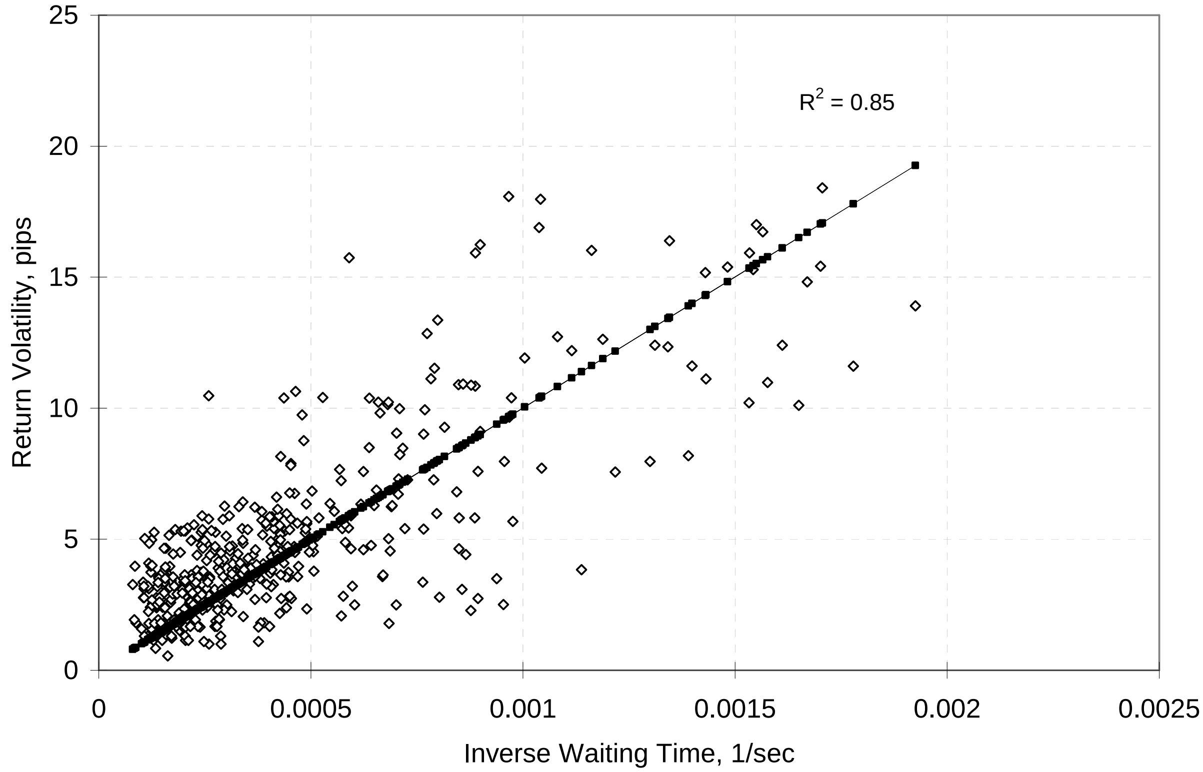

Fig. 4: Relation between the execution time and return volatility. Data for EUR/USD, week ending 28-Apr-2006. The calculations were performed for 5-minute returns and volatility was calculated for every 30-min window.

4. Algorithmic trading of large limit orders

Of course, our estimates of the order execution time do not imply that every order *is* executed within its estimated time interval. A trader who locks his capital in a limit order is exposed to market risk as price may move in the adverse direction and may not revert within an accepted time horizon. To reduce market risk, one has to submit an order that is more aggressive than those already in the order book (i.e. an order with $D < 0$). This, however, incurs immediate losses.

Let us introduce a loss function $L_1(V, D)$ for an order of size V placed at distance D from the best price:

$$L_1(V, D, \lambda) = 100\, V \cdot \left[\lambda\sigma \sqrt{T(V, D)} - D\right] \qquad (1)$$

The first term within the brackets of (1) is an estimate of potential losses due to return volatility, σ (similar to value-at-risk estimates often used in risk management); *T(V, D)* is the expected execution time and λ is the risk aversion coefficient. The second term is the order P/L based on the order price distance to the current market best price. While the bracketed value is measured in pips, V is given in millions of the base currency and the scaling factor outside the brackets, $100V$, presents the loss function in the local currency units (i.e. USD in the case of EUR/USD).

$L_1(V, D, \lambda)$	λ=1					λ=2				
	D					D				
V	**-2**	**-1**	**0**	**1**	**2**	**-2**	**-1**	**0**	**1**	**2**
1	2.31	1.53	1.72	2.12	2.89	2.62	2.05	3.45	5.24	7.78
5	2.36	1.79	1.92	2.61	2.81	2.72	2.58	3.83	6.21	7.61
10	2.46	2.02	2.02	3.12	4.03	2.92	3.03	4.04	7.25	10.06
20	2.64	2.19	2.18	4.25	3.50	3.27	3.39	4.35	9.51	9.00

Table 3. The loss function (1) for $\sigma = 0.47$. Data for $T(V, D)$ is taken from Table 2.

From the example in Table 3 we see that for a given V and λ, there may be an optimal value of D that minimizes the value of the loss function.

This approach can be applied to the algorithm of trading a large amount N by consecutively placed n small orders of amount V $(N = nV)$ such that after one small order is completely filled, the next one is immediately submitted. Optimal slicing of a large equity order into smaller orders is widely discussed in literature (see [13] and references therein). In equity markets, however, the main focus is placed on analysis of market order impact on price. Indeed, a large market order can wipe out a significant part of the order book and hence move price in the adverse direction. However, a large limit order can only absorb that part of the order book which has the order price (or a better one). In this case, potential losses are determined by the distance between the order price and the market best price.

Within our approach, the n^{th} order is locked during the time it takes to execute (n - 1) previous orders as well as the n^{th} order itself. Therefore the potential loss, L_n, for the n^{th} order is:

$$L_n(V, G, \lambda) = 100\, V \cdot \left[\lambda\sigma\sqrt{nT(V, D)} - D\right] \tag{2}$$

As a result, the loss function for the total amount, N, is the sum of the individual loss functions L_1 through L_n:

$$\begin{aligned} L_{(N=nV)}(V, G, \lambda) &= 100\, V\left[(\sqrt{1} + \sqrt{2} + \ldots + \sqrt{n})\lambda\sigma\sqrt{T(V, D)} - nD\right] \\ &= 100\, V\left[\lambda\sigma \cdot \sum_{k=1}^{n}\sqrt{kT(V, D)} - nD\right] \end{aligned} \tag{3}$$

Expression (3) can answer the question whether for a chosen value of λ it is preferable to trade N = 100 with, say, $n = V = 10$, or with $n = 20$ and $V = 5$, etc. Our estimates with $T(V, D)$ taken from Table 2 show that, if $\lambda = 1$, the loss function monotonically decreases with growing V (see Table 3). However, for $\lambda = 2$ and $D = -2$, there is a marginal minimum at $V = 10$.

L_N	λ=1		λ=2	
V	**D=-1**	**D=0**	**D=-1**	**D=-2**
1	46932	120754	83864	63649
5	35416	61611	60832	43012
10	33843	47341	57685	41461
20	30847	38037	51693	42225

Table 4. The loss function calculated using from $T(V, D)$ data taken from Table 2. The optimal values for D were chosen according to the results in Table 3.

5. Long-memory processes in the global FX market

Long-range autocorrelations have been reported for signed order flows in equity markets in several studies (see [11, 12] for recent analysis and references therein). Note that the *signed* order flow is proportional to the *difference* between the buy order volume and the sell order volume. In particular, autocorrelations of the signed order flow with values of approximately 0.05 or higher may last up to

100 5-min intervals (see Fig.4 in [11]). These autocorrelations are explained by order slicing, which is widely used to prevent price impact due to the placing large market orders. In known research of FX markets, weak autocorrelations of the signed deal flow (decaying below 0.05 within about 5 minutes) were described for the interval of 1999 – 2003 [5]. Our results for 2006 confirm such fast decay (see Fig. 5). However, autocorrelations of the signed quote flow decay significantly slower. Namely they remain above 0.1 for at least ten 5-min intervals. The discrepancy between the fast decay of deal autocorrelations and the slower decay of quote autocorrelations can possibly be explained by practically non-existent autocorrelations in the signed hit flow. Indeed, while the quote (maker) order autocorrelations might determine the memory range in the deal dynamics, the uncorrelated hit (taker) orders weaken this memory. We think that a shorter quote order memory in the FX market (versus the order memory in the equity markets) is related to less frequent usage of the order slicing due to restrictions on the minimal order size.

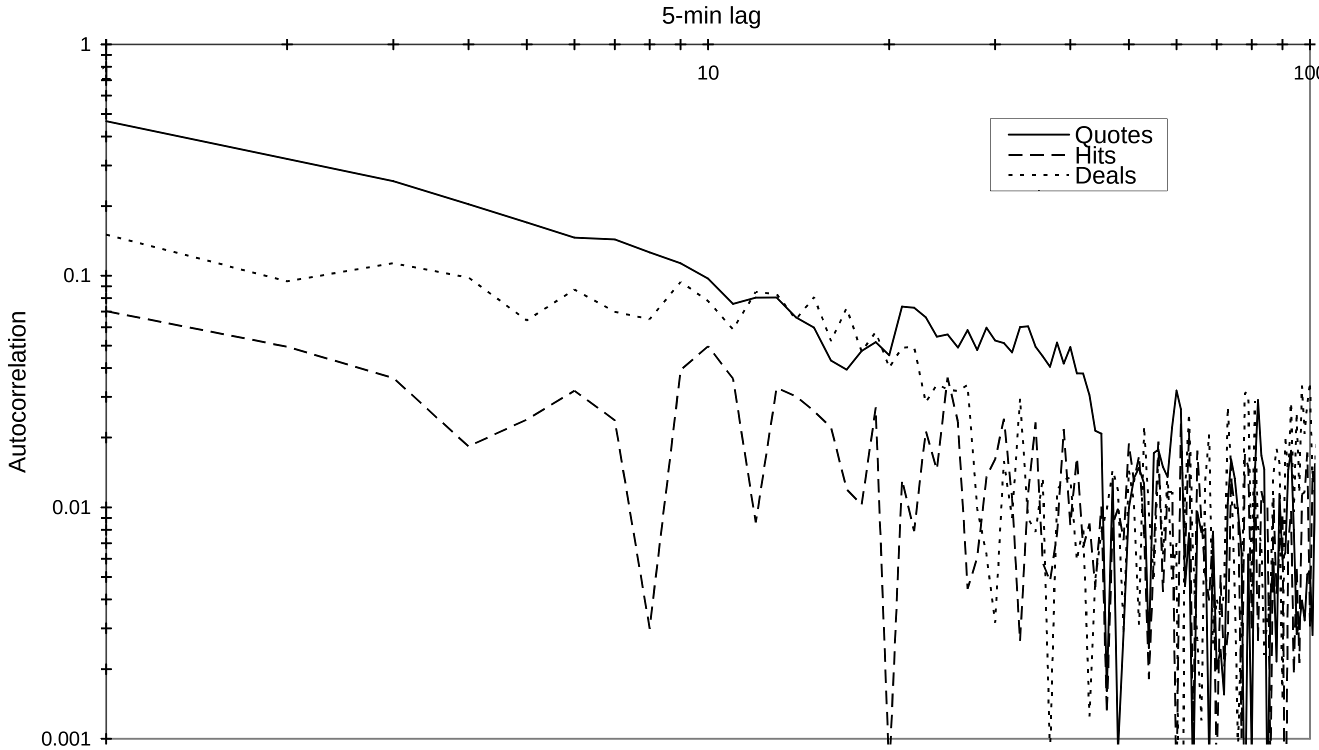

Fig. 5: Autocorrelations in the signed deal flow and the signed order flow.
EUR/USD data for 24-Apr-2006 to 2-Jun-2006.

References

1. Chaboud, A.P., S.V. Chernenko, E. Howorka, R.S. Krishnasami, D. Liu, and J.H. Wright, The High-Frequency Effects of US Macroeconomic Data Releases on Prices and Trading Activity in the Global Interdealer Foreign Exchange Market, *International Finance Discussion Papers,* 823 (2004).
2. Dacorogna, M.M., R. Gencay, U. Muller, R.B. Olsen, and O.V. Pictet, An Introduction to High-Frequency Finance. *Academic Press*, 2001.
3. Goodhart, C.A.O. and M. O'Hara, High frequency data in financial markets: Issues and applications, *Journal of Empirical Finance*, 4, 73-114 (1997).
4. Killeen, W.P., R.K. Lyons, and M.J. Moore, Fixed versus flexible: Lessons from EMS Order Flow, *NBER Working Paper* N8491, 2001.
5. Berger, D.W., A.P. Chaboud, S.V. Chernenko, E. Howorka, R.S. Krishnasami, D. Liu, and J.H. Wright, Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data, *International Finance Discussion Papers,* 830 (2005).
6. Howorka, E. and A.B. Schmidt, Dynamics of the top of the order book in a global FX spot market, In: *Computational Finance and Its Applications*, pp. 257-266, WIT Press (2006).
7. Ito, T., and Y. Hashimoto, Intra-day Seasonality in Activities of the Foreign Exchange Markets: Evidence from the Electronic Broking System, *Journal of the Japanese and International Economies* (2006, in press).
8. Dunis, C.L., and M. Williams, Applications of advanced regression analysis for trading and investment, *Applied Quantitative Methods for Trading and Investment*, pp. 1 – 40, Wiley (2003).
9. Osler, C., Currency Orders and Exchange-Rate Dynamics: Explaining the Success of Technical Analysis, *Journal of Finance* 58, 1791-1819 (2003).
10. James, J., FX trading models – how are they doing?, *Quantitative Finance* 5, N5, 425-431 (2005).
11. Farmer, J. D., A. Gerig, F. Lillo, S. Mike, Market efficiency and the long-memory of supply and demand: is price impact variable and permanent or fixed and temporary?, *Quantitative Finance,* 6, N2, 107 – 112 (2006).

12. Bouchaud, J.-P., J. Kockelkoren, M. Potters, Random walks, liquidity molasses and critical response in financial markets, *Quantitative Finance*, 6, N2, 115-123 (2006).

13. Kissell, R. and M. Glantz, Optimal trading strategies, *AMACOM* (2003).