

Lechner, Sandra; Nolte, Ingmar

Working Paper

Customer trading in the foreign exchange market empirical evidence from an internet trading platform

CoFE Discussion Paper, No. 07/03

Provided in Cooperation with:

University of Konstanz, Center of Finance and Econometrics (CoFE)

Suggested Citation: Lechner, Sandra; Nolte, Ingmar (2007) : Customer trading in the foreign exchange market empirical evidence from an internet trading platform, CoFE Discussion Paper, No. 07/03, University of Konstanz, Center of Finance and Econometrics (CoFE), Konstanz, <http://nbn-resolving.de/urn:nbn:de:bsz:352-opus-32402>

This Version is available at:

<http://hdl.handle.net/10419/32187>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Customer Trading in the Foreign Exchange Market
Empirical Evidence from an Internet Trading Platform

Sandra Lechner*

Ingmar Nolte[†]

University of Konstanz,
CMS

University of Konstanz,
CoFE

This Version: March 15, 2007

*Department of Economics, Box D124, University of Konstanz, 78457 Konstanz, Germany.
Phone +49-7531-88-3214, Fax -4450, email: Sandra.Lechner@uni-konstanz.de

[†]Department of Economics, Box D124, University of Konstanz, 78457 Konstanz, Germany.
Phone +49-7531-88-3753, Fax -4450, email: Ingmar.Nolte@uni-konstanz.de. The work is supported in part by the European Community's Human Potential Program under contract HPRN-CT-2002-00232, Microstructure of Financial Markets in Europe; and by the Fritz Thyssen Foundation through the project 'Dealer-Behavior and Price-Dynamics on the Foreign Exchange Market'. The authors want to thank Richard Olsen and Olsen Financial Technologies for helpful comments and for providing us with the data.

Abstract

This paper analyzes the relationship between currency price changes and their expectations. Currency price change expectations are derived with the help of different order flow measures, from the trading behavior of investors on OANDA FXTrade, which is an internet trading platform in the foreign exchange market. We investigate whether forecasts of intra-day price changes on different sampling frequencies can be improved with the information contained in the flow of our investors' orders. Moreover, we verify several hypotheses on the trading behavior and the preference structure of our investors by investigating how past price changes affect future order flow.

JEL classification: G10, F31, C32

Keywords: Customer Dataset, Order Flow, Price Changes, Foreign Exchange Market

1 Introduction

This paper analyzes the inter-temporal relationship between currency price changes and their expectations on intra-day frequencies. Currency price change expectation are approximated through different order flow measures, which reflect the trading behavior of market participants. The information content of price change expectations is evaluated for the prediction of future currency price changes, and, the influence of historical prices changes on trading decisions is investigated.

The way how information and expectations are aggregated by order flow is central in understanding the microstructure of the foreign exchange (FX) market, which is a highly decentralized market with low transparency. Information on the interpretation of specific news events, risk preferences, hedging demands, central bank interventions, and most important private information are therefore widely dispersed and disaggregated among agents. Traditionally, order flow measures are used to aggregate these dispersed information into one single figure.

Our analyses are based on a customer data-set from a FX internet trading platform, OANDA FXTrade, which contains detailed information on traders' characteristics and currency positions. Relating order flow to traders' characteristics provides valuable insights into the dynamics of order flow and price changes.

Most of the existing studies (e.g. Evans & Lyons (2002a,b), Rime (2003) and Daniéls-son, Payne & Luo (2002)) focus on agents in the interbank market and consider the relationship between prices and order flow obtained either from direct (e.g. Reuters Dealing 2000-1) or brokeraged (e.g. Reuters Dealing 2000-2, EBS) interdealer trading. The studies of Osler (2002) and Marsh & O'Rourke (2004) use a data-set on customer trades collected by the Royal Bank of Scotland. They investigate how customer-trading-order-flow, which is the primary source of private information for a player in the interbank market, is related to currency prices.

In these studies order flow is usually measured by the standard net order flow measure of Lyons (1995), who suggests aggregating all the dispersed information into one single measure: the difference between the number of buyer- and seller initiated trades over a certain sampling frequency. The study by Daniéls-son et al. (2002) underpins the central role of order flow in explaining exchange rate dynamics. It provides evidence that on intra-day aggregation levels, exchange rates are out-of-sample predictable. They propose simple models which outperform random walk forecasts using additional information on order flow and refute the Meese & Rogoff (1983a,b) findings.

We, in contrast, consider traders from an internet trading platform, which are mainly retail investors usually having no private information in terms of observing own customer order flow, and we investigate first whether their price expectations and their trading behavior are helpful to predict future currency prices. This approach can be justified by recognizing that for OANDA FXTrade itself, the actions of their customers (traders) create valuable private information, which can be incorporated into OANDA FXTrades' hedging and trading strategies on the primary market. Furthermore, even in the absence of private information (customer order flow) for the group of OANDA FXTrade traders, this group forms expectations based on different information sources (e.g., technical analysis, public news) and own trading experience, which might be pooled together and extracted in the correct way be helpful in explaining future currency price changes.

Second, we analyze whether and how the OANDA FXTrade investors are influenced by past currency prices. Considering the literature on market microstructure and behavioral finance, we derive four hypotheses about the relationship between price changes and order flow. The validity of our hypotheses is investigated with forecasting studies and out-of-sample prediction criterions. Applying the modified Diebold-Mariano test of Harvey, Leybourne & Newbold (1997), we test whether forecasting models for intra-day price changes (order flow) incorporating additional information on order flow (price changes) provide better forecasts than corresponding benchmark models, which contain information on historical prices changes (order flow) only. We, in contrast to Daniélsson et al. (2002), use $AR(p)$ specifications instead of random walk models as benchmark models, since on intra-day frequencies the price change processes, as well as the order flow processes, are subject to specific intraday autoregressive structures, such as bid-ask bounce (Roll (1984)) effects.

The paper is organized as follows. In Section 2, we briefly describe the foreign exchange market. Section 3 explains in detail the trading mechanism and the different order types on the OANDA FXTrade platform. Section 4 describes the data-set. We motivate our empirical study and we formulate the economic hypotheses in Section 5. Section 6 contains the empirical results and the verification of the hypotheses, while Section 7 concludes.

2 A Brief Description of the Foreign Exchange Market

The FX-market is generally characterized by a high degree of decentralization, low-transparency, and 24h trading. According to the Triennial Bank for International Settlements' (BIS) Report (2004), the nine most active trading centers in 2004 in the FX spot market are London (31.3%), New York (19.2%), Tokyo (8.3%), Singapore (5.2%), Frankfurt (4.9%), Hong Kong (4.2%), Sydney (3.4%), Zürich (3.3%), and Paris (2.7%) accounting for a total turnover of 82.5%. The three most actively traded currency pairs are *USD/EUR* (28%), *USD/JPY* (17%), and *USD/GBP* (14%). In 2004, the total average daily turnover amounts to 1,773 bn\$, which is proportioned into spot (35%), forward (12%), and swap (53%) markets.

Several groups of agents trade and interact with each other in the FX market. First of all, there is the group of non-financial customers (17% of the FX spot market trading), which mainly consists of corporations (exporters, importers). With the establishment of internet trading platforms, the access to the FX market has been facilitated for retail investors and retail customers. We consider them as the second group of agents, which alternatively may be treated as a subset of group one. There are no reliable figures on their share of the total turnover in the FX market, but it is considered to be very small. However, due to the increasing number of internet trading platforms and a move from traditional bi-lateral to electronic trading, their share is expected to continue increasing. Investors trading on internet trading platforms in the FX market may be associated with retail investors trading different instruments such as stocks and options on other markets through discount brokers. One important characteristic of these two groups is that they do not have access to the interbank (interdealer) market, and they have, if at all, only very limited information on each other. A third group of agents in the FX market consists of financial institutions without access to the interbank market (34% of the FX spot market trading), such as smaller commercial banks, security houses, mutual funds, pension and hedge funds, and insurance companies.

These three groups have basically two channels to settle a transaction: i) directly via dealer-banks that offer a bid-ask spread, which is mainly driven by order handling costs (the smaller and the more unconventional the order (size), the higher the bid-ask spread), ii) they can trade with each other via internet trading platforms, which try to offer a small (close or equal to interbank) bid-ask spread independently of the

trade size to attract customers. Therein, the market of internet trading platforms itself is divided into two groups: a) platforms which are established by banks or consortiums of banks, such as FXConnect or Currenex, and b) non-bank trading platforms such as Deal4Free or OANDA FXTrade, which is the source of our dataset. Usually, these internet trading platforms are at least partially organized as so called crossing networks, since there is too little trading to have a (arbitrage free) price discovery. Bid and ask quotes of crossing networks are based completely, or in addition to their own limit order book on other trading channels, e.g. electronic brokers like Reuters Dealing 3000-2 or EBS. The quoted prices are then either a simple put through of the external data-feed or forecasted prices based on the recent history of the data-feed. Besides offering interbank spreads, internet trading platforms have the advantage that depending on the platform, customers may have (limited) access to the limit order book and the history of the trades and quotes. Therefore, the transparency is higher than in the direct bi-lateral customer-to-dealerbank trading. The last group of agents consists of dealer-banks (48% of the FX spot market trading), which trade in general as the counterparty with members of the first three groups or with each other in the interdealer market. Trading in the interdealer market is usually done in two different ways: either directly (by telephone and via Reuters Dealing 3000-1), or indirectly via brokers (voice brokers and electronic brokers).¹

3 OANDA FXTrade in Detail

OANDA FXTrade is a fully virtual marketplace for trading currencies via the internet, without limits on the trade size, and with 24 hours trading time 7 days per week. This platform is a market making system that executes orders using the exchange rate prevalent in the market (determined either by their limit order book or by forecasted prices relying on an external data-feed). OANDA FXTrade offers immediate settlement of trades and tight spreads as low as 2 to 3 pips on all transaction sizes. Given various boundary conditions, as for example sufficient margin requirements, orders are always executed. The OANDA FXTrade platform is based on the concept of margin trading, this means that the trader can enter into positions larger than his funds. The platform requires a minimum initial margin of 2% on positions in the major currency pairs and 4% in all other currency pairs, which correspond to

¹For a more detailed description of the FX market we refer to Lyons (2001) and Rime (2003).

a leverage² of 50:1 and 25:1 respectively. In other words, for each dollar margin available the trader can make a 50 (25) dollar trade. The trader receives a margin call when the net asset value (i.e. the current value of all open positions plus the value of the remaining deposited funds) becomes half the margin requirement. Thus, if the trader does not have sufficient margin to cover twice the losses on an open position, a margin call order is used to close automatically all open positions using the prevalent market rates at this time.

Market orders (buy or sell) are executed immediately and affect existing open positions. Limit orders are maintained in the system for up to one month. The server manages the limit order book, the current exchange rates, and the current market orders to match existing limit orders. The limit order can therefore be matched either against a market order, or against a bid or an ask price obtained from the external data-feed. The legal counterparty of a trade, however, is always OANDA FXTrade. Stop-loss orders and take-profit orders are special limit orders in the sense that they can be set for existing open positions. They can be specified directly while entering a market or a limit order, but they can also be specified later for existing open positions. Stop-loss and take-profit orders are automatically erased from the system whenever a position is closed due to further trading activity. Take-profit (TP) orders are typically set to close an existing position after a certain profit has been realized. Stop-loss (SL) orders, in contrast, specify that the position should be closed after the realization of a certain loss to avoid further losses. Table 1 overviews the transactions and further activities of the traders on OANDA FXTrade, which are recorded in an activity record file.

²A leverage of 50:1 is the maximum offered by OANDA FXTrade.

Buy/Sell* market open (close)	Immediately executed to open or close a position in a specific currency pair.
Buy/Sell limit order	The trader posted a buy or sell limit order to the system, which is then pending.
Buy/Sell limit order executed open (close)	Pending limit order is executed to open or close a certain position.
Buy/Sell take-profit close	Closes an open position by buying or selling the currency pair when the exchange rate reaches a predetermined level, in order to make a profit.
Buy/Sell stop-loss close	Closes an open position by buying or selling the currency pair when the exchange rate reaches a predetermined level in order to avoid further losses.
Buy/Sell margin call close	Closes automatically all open positions using the prevalent market rates at the closing time. This happens if the trader has not sufficient margin to cover two times the losses of all open positions.
Change order	Change of a pending limit order (limits for take-profit or stop-loss, the value of the upper or lower bounds, the quote as well as the number of units).
Change stop-loss or take profit on open trade	Change stop-loss or take-profit limit on an open position.
Cancel order by hand	Cancel a pending limit order by hand.
Cancel order: insufficient funds	Automatically recorded when the trader has not enough funds to open a new position.
Cancel order: bound violation	Market order or limit order is cancelled because the applied exchange rate is not located inside the specified bounds.
Order expired	A pending limit order is expired.

Table 1: Activity record entries of OANDA FXTrade.

*On the OANDA FXTrade platform, buying *EUR/USD* means that you are buying the base currency (*EUR*) and selling the quote currency (*USD*), whereas selling *EUR/USD* means that you are selling the base currency (*EUR*) and buying the quote currency (*USD*). Recorded units always refer to the base currency.

4 Description of the Data-Set

The dataset that is used in our analysis is constructed from the trading activity record of OANDA FXTrade from 1st October 2003 to 14th May 2004 (227 days). This record contains for 30 currency pairs all trading activity on a second by second basis and allows us to distinguish between the transaction types listed in Table 1. In

addition, depending on the order type, we get information on the transaction prices (market orders, limit orders executed, stop-loss, take profit, margin call), on the bid and ask quotes (limit orders pending), on the adjoint transaction units, and on the limits of stop-loss and take-profit orders.

We focus in our analysis on the most actively traded currency pair *EUR/USD*, which accounted for nearly 39 % of all records with an average interrecord-duration of 8.5 seconds. For this currency pair, 13.5% of all transactions have a transaction-volume between 1€ (min.) and 100€, whereas only 1% of the transactions have a transaction-volume that ranges between 50,000€ and 1,000,000€ (max.). The average transaction-volume per trade is 26,546 € and the average number of different traders per day is 744.

Using only price determining orders (market orders, limit orders, limit orders executed, stop-loss, take profit, margin call), we construct from this data-set equidistant *EUR/USD* price series on 12 frequencies (1 min, 2 min, 5 min, 10 min, 15 min, 20 min, 25 min, 30 min, 45 min, 1 hour, 2 hours, 4 hours). Throughout the paper, we refer to these price series as the “OANDA based” prices series. In addition, we construct from a series of mid-quotes from the interbank market, available on an 1 min aggregation level, the corresponding price series for the remaining 11 frequencies applying the previous tick aggregation technique. These prices are denoted in the following as “interbank” prices. Since we can only observe mid-quotes from the interbank market, we decided to use on the 1 min frequency mid-transaction prices from OANDA, to avoid, because of bid-ask bounce effects, any kind of unfair comparison due to increasing forecastability when transaction prices aggregated with a previous tick aggregation technique are used. The mid-quotes from the interbank market are provided by Olsen Financial Technologies and represent tradeable quotes stemming from different electronic brokerage systems including Reuters Dealing 3000 and EBS. These mid-quote series do not coincide with the bid and ask quotes on OANDA FXTrade. The bid and ask quotes on OANDA FXTrade are generated by an proprietary forecasting algorithm based on an external data-feed from which also includes tradeable quotes from Reuters Dealing 3000 and EBS.

5 Motivation and Economic Hypotheses

Considering the economic literature, an everlasting extensively studied topic is the relationship between expectations and price formations. The literature can be traced back to the seminal works of Muth (1961), ?, and ?, where theories of rational, implicit, and adaptive expectations are introduced. The empirical verification of these hypotheses faces the major problem of measuring expectations reliably. For example, in the analyses of firm and expert surveys (e.g. Carlson & Parkin (1975), Nerlove (1983) and Pesaran (1987)), survey responses serve as a proxy of the participants' expectations on the future development of macroeconomic or financial price series. In addition to the information contained in the history of the underlying series itself, the responses, usually in aggregated form, are used to predict the underlying series for a medium term horizon (few months). The information obtained from the survey is treated, in that sense, as private or insider information which yields the (improved) forecastability of the underlying series.

For the short term prediction (up to one day) of asset price series (e.g. stocks, exchange rates and commodities), a different methodology can be applied to measure the expectations of the market participants, which may then serve as insider information as well. The only assumption which is required is that market participants reveal their expectations through their trading behavior. Therefore, the flow of the orders of market participants can be considered to contain information on their conditional expectations of the future development of the prices of those assets they are trading. One theoretical foundation is given in the portfolio allocation model of Evans & Lyons (2002a,b), where exchange rate movements are explained by changes of previous customer order flow, that represent changes in an underlying portfolio. In their model, there are two different markets, the customer-dealer market and the dealer-dealer (interbank) market. Dealers trading in the interbank market learn about order flow in the customer-dealer market and this customer order flow will predict currency price changes and order flow in the interbank market. Another foundation can be based on the argument of Sarno & Taylor (2001), who consider order flow as a proxy for macroeconomic fundamentals. Thus, changes in currency prices are driven by changes in macroeconomic variables, which are revealed to market participants in the form of news announcements, for example. Both models require that the market participants interpret information, either on portfolio changes or on changes of macroeconomic fundamentals, in the correct way, that they adjust their expectation on the future prices in the light of this information, and that they there-

fore place their orders accordingly. The forecasting study of Danielsson et al. (2002) shows that exchange rates are, contrary to Meese & Rogoff (1983a,b), out-of-sample predictable and outperform random walk forecasts using additional information on order flow. The multi-faceted literature on inventory and/or asymmetric information based models for security markets (Demsetz (1968), Ho & Stoll (1981), Kyle (1985), Foster & Viswanathan (1990), Easley & O'Hara (1992), Biais, Hillion & Spatt (1995), ?) provides further theoretical foundations that (bid and ask) prices can be explained by previous order flow. The common idea in virtually every market microstructure model, including the before mentioned, is that market participants react to previous actions (order flow) of other market participants resulting in impacts on current or future prices. In a pure inventory model based market, market makers adjust bid and ask prices according to their current inventory, which is naturally a consequence of orders executed previously. In a fully electronic order book market without market makers, traders react to actions of other traders, which are usually displayed to them (partially) through the limit order book. However, in all of these models the key determinant, which finally decides on the success or failure of the model is that expectations are interpreted, measured, and modelled appropriately.

Our analysis is concentrated on data from the FX market. Therein we focus on a very special segment, namely an internet trading platform, OANDA FXTrade, where most of the traders are retail investors or members of the group of non-financial customers. Most of the research on order flow and currencies focuses on the interbank market (e.g. Bjønnes & Rime (2003), Evans & Lyons (2002a,b), Payne (2003)) and the papers by Marsh & O'Rourke (2004) and Osler (2002) deal with customer orders observed by the Royal Bank of Scotland (dealer bank). However, to our knowledge there has been no analysis of customer data obtained from an internet trading platform, yet.

In the FX market customer order flow (trading between a dealer bank and their non-interbank market customers) is the most fruitful source of private information for a dealer bank. Their customers are usually large companies, commercial banks, security houses, mutual funds, hedge funds, and insurance companies, who want to settle transactions of sizes which are often several times higher than the standardized order sizes in the interbank market. In line with the portfolio allocation model of Evans & Lyons (2002a,b), these customer orders are the primary source to identify dispersed information and they consequently induce interdealer orders (e.g. "hot potato" trading, inventory control) that affect the currency price.

Therefore, one can argue that order flow from our internet trading platform **does not** contain any helpful information to predict future prices, since our traders submit only orders of small size, which do not affect the interdealer market. Stated differently, traders on our internet trading platform are noise traders. On the other hand, even our traders form expectations on the future development of the currency price that they reveal through their trading activity and which represents private information for OANDA FXTrade itself. Thus, one can assume that order flow from our internet trading platform **does** contain information that is helpful to predict future prices. Therefore we can derive our first hypothesis:

Hypothesis H1:

Information on the order flow on OANDA FXTrade is helpful in predicting future currency prices.

For the empirical proof of this hypothesis an important question arises: how should order flow be measured exactly? Lyons (1995) introduces the standard definition of an aggregated net order flow measure as the difference between buyer initiated and seller initiated trades (within a given period), or stated differently, as the cumulative sum of signed orders where buyer (seller) initiated orders get positive (negative) signs. Focusing on the initiating party of a trade, this definition aims to capture very recent changes in the expectations of future prices that may arise because of new (private) information. For example, an executed buy limit order is treated as a seller initiated trade since it has to be merged with a sell market order. Therefore the expectation of the seller is treated to be more important than the expectation of the buyer, who might not have the latest information. The standard order flow measure is very well suited in predicting future prices when the interbank market is considered, as demonstrated by Daniélsson et al. (2002).

Let us now consider trades on OANDA FXTrade where bid and ask prices depend on an external data-feed. A buy limit order (bid) is therefore usually matched against the ask price of OANDA FXTrade, which is a function of the prices in the interbank market. For the simplicity of the argument let us assume that prices from the primary market are put through one-to-one to OANDA FXTrade, so that the ask price process on OANDA FXTrade is the same as the one on the interbank market. The lower ask price (crossing limit sell order at the best ask) which is matched against the OANDA buy limit order is therefore generated by selling pressure in the primary market shortly before, for example a large sell market order, consuming the previ-

ous best bids in the interbank market and causing an adjustment also of the ask quotes to lower prices. Thus, measuring order flow on OANDA FXTrade with the standard net order flow measure yields a mixture of price expectations from traders on OANDA FXTrade (mainly through market orders) and price expectations from the interbank market (mainly through executed limit orders).

An alternative to the standard net order flow measure is to consider a measure that solely aggregates the price expectations on OANDA FXTrade. In Table 2, we summarize definitions of the standard and the alternative order flow measure, which we denoted as “OANDA order flow” measure. Therein, we list the different entries of the OANDA FXTrade activity record with the corresponding occurrence share in column one and two. Column three depicts the signs for the standard order flow, and column four the signs for the OANDA order flow measure. Buy market orders, irrespectively of whether they are submitted to open or close a position, get positive signs in both order flow measure since the traders on OANDA FXTrade initiate these trades or believe that the price will go up, respectively. Correspondingly, the symmetric sell market orders get negative signs.

In the standard order flow measure, submitted (pending) limit orders are not considered, since they are not yet executed, which means that there is no initiating party yet. However, they are taken into account in the OANDA order flow measure since the trader, who submits a limit order, expressed his beliefs that the price will go up (buy, positive sign) or down (sell, negative sign).

Executed buy limit orders are treated as seller initiated in the standard order flow measure (see the discussion above) and thus award negative signs, whereas they get positive signs in the OANDA order flow measure, since the submitter still believes that the price will go up. Otherwise he would have cancelled the order before execution. Executed sell limit orders are treated analogically. For the OANDA order flow measure limit orders are counted twice now, once at their submission time and once at their execution time. However, since they are usually counted at two different times this does not create a problem, because we still measure beliefs of the investors that might have been up-dated in between.

Buy take-profit orders (close) are buy limit orders that get negative signs in the standard order flow measure. In the OANDA order flow measure, they get positive signs, because the trader believes that the price will further fall. A buy take-profit order (close) can only be executed if the trader has a short position in a currency pair (short position in the base currency). Sell take-profit orders get the analog

signs.

Buy stop-loss orders (close) get negative signs in both measures. In the standard order flow measure the explanation is that it is a special buy limit order. In the OANDA order flow measure the explanation is that the trader believes that the price will further fall. Again, sell stop-loss orders are treated correspondingly. Buy margin call orders (close) are not used in the standard order flow measure. On the one hand, one can argue that they should get positive signs since they are buy market orders. On the other hand, one can argue that they are not motivated by new information and that the traders are proven by the price process of the primary market to have wrong expectations on the price. Therefore they should get negative signs. Anyway, due to their scarce occurrence (0.12% and 0.17%) they do not play an important role. However, in the OANDA order flow measure they are counted, since although the traders are proven to have wrong expectations about the price, they still believe that the price will go down (up) in the case of a buy (sell) margin call order.

Transaction Record	Percentages	Standard Order Flow Signs	OANDA Order Flow Signs
Buy market (open)	13.10	+	+
Sell market (open)	10.61	-	-
Buy market (close)	8.27	+	+
Sell market (close)	10.27	-	-
Limit order: Buy	5.41	not used	+
Limit order: Sell	4.76	not used	-
Buy limit order executed (open)	3.22	-	+
Sell limit order executed (open)	2.92	+	-
Buy limit order executed (close)	0.46	-	+
Sell limit order executed (close)	0.46	+	-
Buy take-profit (close)	3.14	-	+
Sell take-profit (close)	3.49	+	-
Buy stop-loss (close)	2.18	-	-
Sell stop-loss (close)	2.55	+	+
Buy margin call (close)	0.12	not used	-
Sell margin call (close)	0.17	not used	+
Change order	3.01	not used	not used
Change stop-loss or take-profit	22.36	not used	not used
Cancel order by hand	2.41	not used	not used
Cancel order: insufficient funds	0.28	not used	not used
Cancel order: bound violation	0.20	not used	not used
Order expired	0.65	not used	not used

Table 2: Col. 1 states the record entries, col. 2 gives the corresponding percentages, col. 3 contains the signs for the construction of the standard net order flow measure and col. 4 contains the signs for the construction of the OANDA order flow measure.

Given these two order flow measures, we can refine Hypothesis H1 with respect to the measuring of the order flow:

Hypothesis H1.1:

Order flow that contains information on price expectations from the interbank market and OANDA FXTrade (standard order flow measure) is helpful in predicting future currency prices.

Hypothesis H1.2:

Order flow that contains information on price expectations from OANDA FXTrade solely (OANDA order flow measure) is helpful in predicting future currency prices.

We verify these hypotheses by testing the in-sample fit and the out-of-sample fore-

casting performance of the following regressions:

$$\Delta y_t^h = c + \beta_{x1}x_{t-1}^k + \beta_{y1}\Delta y_{t-1}^h + \dots + \beta_{xp}x_{t-p}^k + \beta_{yp}\Delta y_{t-p}^h + \varepsilon_t,$$

where Δy_t^h denotes the currency price change from $t - 1$ to t , x_t^k the value of the order flow measure at t , and ε_t the error term. p defines the number of lags used in the regression. $k \in \{\text{SOF}, \text{OOF}\}$ denotes for x_t^k whether the standard order flow measure, using information from the interbank market ($k = \text{SOF}$), or the OANDA order flow measure, using information from OANDA FXTrade solely ($k = \text{OOF}$), is used. For the price change Δy_t^h , h distinguishes whether price changes from the interbank market ($h = \text{IP}$) or price changes from OANDA FXTrade ($h = \text{OP}$) are used.

For reasons of comparison, we also investigate the performance of a purely data driven order flow measure which is not based on any theoretical motivation of how expectations of future prices should be measured. Since in both order flow measures above buy and sell orders are treated symmetrically (opposite signs), we compute the change of the order flow for every transaction category. For example, we compute the order flow of the market order (open) category as the difference between the number of buy market orders (open) and sell market orders (open) over the sampling period. Thus we obtain eight category specific order flow measure which are summarized in Table 3.

Category	Description
1	Limit orders
2	Limit orders executed (open)
3	Limit orders executed (close)
4	Market orders (open)
5	Market orders (close)
6	Stop-loss orders (close)
7	Take-profit orders (close)
8	Margin call orders (close)

Table 3: Col. 1 states number of the category and col. 2 gives the category description.

The corresponding regression takes the following form:

$$\Delta y_t^h = c + \sum_{k=1}^8 \beta_{k1}x_{t-1}^k + \beta_{y1}\Delta y_{t-1}^h + \dots + \sum_{k=1}^8 \beta_{kp}x_{t-p}^k + \beta_{yp}\Delta y_{t-p}^h + \varepsilon_t,$$

where x_t^k denotes the order flow in the associated category $k = 1, \dots, 8$ at time t . Again Δy_t^h with $h \in \{\text{IP}, \text{OP}\}$ denotes the interbank or the OANDA FXTrade price

change and p the selected number of lags.

With the hypotheses derived above, the causal relationship from order flow to price changes is investigated. The survey study of Taylor & Allen (1992) however, shows that at least 90% of the London based dealers rely, in addition to private and fundamental information, on technical analyses information to design their trading strategies. This is a typical example that price changes or certain patterns in the price process cause reactions of market participants, and therewithin order flow. Another example of causality from prices to order flow is the study by Osler (2002), where it is analyzed whether executions of special limit orders (stop-loss and take-profit) contribute to self-reinforcing price movements. The idea behind this investigation is that there are local downward or upward trends in the price process, which are accelerated by the execution of stop-loss orders, which generate positive feedback trading, and are decelerated by the execution of take-profit orders, which generate negative feedback trading. For the illustration of the argument, let us assume that the price is decreasing, which in the first case may cause an execution of a sell stop-loss order and induces further selling pressure, which leads to further executions of sell stop-loss orders. Thus, we get an accelerated downward moving price process (price cascades). In the second case, a downward moving price may cause an execution of a buy take-profit order, which does not induce further selling pressure and therefore no execution of either further stop-loss or take profit orders, which yields a decelerated downward movement or even an upward moving price process.

The OANDA FXTrade activity data-set is well suited to investigate how traders react to specific patterns in the price process. In the light of the order flow measures introduced above, we can analyze whether information on price changes is helpful to predict future order flow, which constitutes our second hypothesis:

Hypothesis H2:

The price process contains information that is helpful in predicting future order flow.

We verify this hypothesis again with respect to the price process obtained from the external data-feed and the price process obtained from OANDA FXTrade directly. Since the traders on OANDA FXTrade are usually only affected by the FXTrade price process, we expect that it should have more power in explaining future order flow than the external data-feed price process. Moreover, we use the standard and

the OANDA order flow measures, as well as the category based order flows (Table 3), to investigate this hypothesis. Considering the influence of the price processes on the category based order flows more precisely, we can investigate whether we observe self-reinforcing price movements in the sense of Osler (2002) on OANDA FXTrade as well:

Hypothesis H3:

Executed stop-loss orders contribute to self-reinforcing price movements, whereas executed take-profit orders impede self-reinforcing price movements.

Two analyses are conducted to investigate this hypothesis with the help of category based order flow: i) given that hypothesis H3 is correct then based on their own histories, order flow in the stop-loss order category should be better predictable than order flow in the take-profit order category, ii) if stop-loss orders induce self-reinforcing price movements and take-profit order not, then information on the price process itself should be more valuable, in addition to their own histories, for predicting take-profit orders' order flow than for predicting stop-loss orders' order flow. Furthermore, the category specific order flow measures allow insights into several aspects of our traders' preference structure. Thereby, we are able to exploit the information whether trades are executed to open or close a certain position. Following the argument of Glosten (1994) and Seppi (1997) that impatient traders tend to submit market orders, whereas patient traders tend to submit limit orders, we can analyze the trader structure on OANDA FXTrade. Moreover, we can refine this analysis by investigating whether the patience or impatience pattern depends on the trader's current investment status, whether they already hold a (profitable or unprofitable) position in a specific currency pair or not. The extreme position in this respect is to claim that the degree of (im-)patience depends on the history of the price process but not on the investment status of the traders. An impatient trader would react to a price change buy submitting a (buy or sell) market order independently of the position already obtained, simply for the reason that it is a good opportunity to trade. In this case we should observe that the order flow in the market order (**open**) category is as good predictable (based on the information contained in the price process) as the order flow in the market order (**close**) category.

The prospect theory of Kahneman & Tversky (1979) however, relies on the assumption that people (traders) evaluate the outcome of a gamble (investment strategy)

with respect to a certain previously obtained reference point. The zero profit point thereby serves as the natural reference point³ and traders react differently in the case that their trading strategy has already generated profits or losses. Since being not invested in a certain currency pair trivially corresponds to the zero profit point one should expect that the degree of traders' (im-)patience depends on his investment status. With slightly different words, one could claim that there is a kind of monitoring effect in the sense that traders react to information more quickly when they fear to lose something – which is certainly true when they already hold a position – than in the case where they plan to invest into a position. Following this explanation, we should observe that the order flow in the market order (**close**) is better predictable (based on the information contained in the price process) than order flow in the market order (**open**) category. Still assuming that traders react and transact in the absence of private information or fundamental information based on the information contained in the history of the price process, we can formulate our last hypothesis:

Hypothesis H4:

Traders' degree of (im-)patience depends on their actual investment status.

Another reason to concentrate on the analysis and to develop models for the analysis of customer activity data-sets such as OANDA FXTrade is that customers have basically two possibilities to trade: either by trading with a dealer-bank or by trading via an electronic (internet) platform. As pointed out by Lyons (2002), there is recently a shift in the interdealer market from direct trading towards electronic brokerage trading. One argument to explain this shift is that there is more transparency on electronic brokerage systems. In the customer market segment, one can expect the same shift from dealer-bank trading towards internet platform trading, since these platforms are also more transparent and try to offer small (interbank) spreads to all of their customers.

³For a discussion on reference points we refer to Odean (1998), Weber & Camerer (1998) and Barberis & Thaler (2003).

6 Empirical Findings

6.1 Descriptive Analysis

In Figure 1 we show the diurnal seasonality function of the standard and the OANDA order flow measures, computed by a Nadaraya-Watson kernel regression with a Gaussian kernel and optimal bandwidth selection according to Silverman's (1986) rule on a 10 min aggregation level, where the time scale is measured in Eastern Standard Time (EST). The first observation that should be made is that there is a kind of diurnal seasonality pattern, which is much more pronounced for the OANDA order flow measure than for the standard one. However, both seasonality patterns correspond to standard market activity⁴: we observe a positive peak at 3 o'clock, when the European traders enter the market, and a negative peak around 5 o'clock which corresponds to lunch time in Europe. We see a strong upward recovery between 6-9 o'clock, which is the afternoon trading in Europe and the trading start in the US. The decline after 9 o'clock can be explained by European traders leaving the market successively and the positive peak around 11 o'clock corresponds to the market phase where the US traders are most active. The declining trading activity of the US traders from 12 o'clock onwards results in a negative peak around 17 o'clock. The recovery of the trading activity thereafter, with a peak at 19 o'clock, is due to Asian investors entering the market. Please, keep in mind that we postulate the same seasonality pattern for every weekday, because on a daily frequency we can only analyze 163 observations.




Figure 1: Diurnally seasonality in the standard (1st column) and the refined (2nd column) net order flow measure, computed on a 10 min aggregation level.

Figure 2 depicts the empirical bivariate autocorrelation functions up to 20 lags between price changes and order flow for a frequency of 1 minute.

⁴See Andersen & Bollerslev (1997) and Dacorogna, Gençay, Müller, Olsen & Pictet (2001).

There are four main panels, each divided into 2 by 2 subordinated panels. The upper left main panel displays the (empirical bivariate) autocorrelation function of OANDA based price changes and standard order flow; the upper right main panel displays the autocorrelation function of OANDA based price changes and the OANDA order flow measure; the lower left main panel displays the autocorrelation function of interbank price changes and standard order flow; and the lower right main panel displays the autocorrelation function of interbank price changes and the OANDA order flow measure. For each main panel, the upper left subordinated panel depicts the autocorrelation function of the particular order flow measure, the lower right panel depicts the autocorrelation function for price changes. For these two, we plot lag 1 up to lag 20. The lower left subordinated panel depicts the cross-correlation function of lagged order flow with price changes, and the upper right panel depicts the cross-correlation function of lagged price changes with order flow. For these two, we plot lag 0 up to lag 19. The value at lag 0 is in both cross-panels the same and represents the contemporaneous correlation between the particular order flow and price changes.

The analysis of the bivariate autocorrelation functions allows us to shed light on the dynamic interaction of the particular order flow and price change series and it enables us to verify some of the hypothesis stated in the previous section from a descriptive point of view. The following observations are worth pointing out:

- For both order flow measures we observe in the lower left subordinated panels significant cross-correlation coefficients, which show that future (OANDA based and interbank) price changes are driven by current order flow, which support hypotheses H1, H1.1 and H1.2 that order flow is helpful in predicting future currency prices. We observe that only the first order cross-correlation coefficients are significantly positive between current OANDA order flow and future price changes of both price series. In the case of current standard order flow and future interbank prices changes only the first order cross-correlation coefficient is significantly positive again, whereas in the case of current standard order flow and future OANDA based prices changes, the first three cross-correlation coefficients are significantly different from zero. Thereby, the positivity of the first order cross-correlation coefficients is partially compensated by the negativity the second and the third. Note, in all four cases the first order cross-correlation coefficients are always positive but higher when the interbank instead of the OANDA based prices are involved.

- For both order flow measures we observe in the upper right subordinated panels significant cross-correlation coefficients, which show that future order flow is driven by current price changes. This observation supports from a descriptive point of view hypothesis H2: that investors update their beliefs and place their orders based on the past development of the price process. However this effect seems to be a short term effect, since the cross-correlation coefficients between future order flow and current price changes are significant only up to 3 lags with the OANDA based prices, and up to 5 lags with the interbank prices. Furthermore, the correlation coefficients for the standard order flow measure are larger than those for the OANDA order flow measure. This means that the standard order flow measure has not only a higher contemporaneous correlation with price changes, but is also influenced more severely by past price changes than the OANDA order flow.
- In the upper left subordinated panels, we observe the autocorrelation function of the order flow measures themselves. For the standard order flow measure, we get a very clear slowly declining pattern of the autocorrelation function, whereas for the OANDA order flow measure, only the first, the third, the fourth and the twelfth autocorrelation coefficients are significantly different from zero, generating an unsystematic pattern for the autocorrelation function. Relating order flow to the process of price expectation updates, we observe a persistent updating process when information from the interbank market is incorporated (standard order flow) and a process with an irregular updating pattern in the case where only the information from OANDA market is used.
- In the lower right subordinated panels, we observe the autocorrelation function of the price changes themselves. The price changes are positively first order auto-correlated, which is partially compensated by negative auto-correlation coefficients of order 2 to 5 for the OANDA based price change series. Thus, we observe a kind of short term positive feedback trading pattern for both price processes. Due to the fact that we consider mid-quotes on a 1 minute frequency, we cannot observe the traditional bid-ask bounds effect.




Figure 2: Empirical bivariate autocorrelation function of price changes and order flow for an aggregation level of 1 min. There are four main panels, each divided into 2 times 2 subordinated panels. The upper left main panel displays the (empirical bivariate) autocorrelation function of OANDA based price changes (OP) and standard order flow (SOF), the upper right main panel displays the autocorrelation function of OANDA based price changes and the OANDA order flow measure (OOF), the lower left main panel displays the autocorrelation function of interbank price changes (IP) and standard order flow and the lower right main panel displays the autocorrelation function of interbank price changes and the OANDA order flow measure. For each group, the upper left panel depicts the autocorrelation function (lag: 1–20) of the particular order flow measure, the lower right panel depicts the autocorrelation function (lag: 1–20) for price changes. The lower left panel depicts the cross-correlation function (lag: 0–19) of lagged order flow with price changes and the upper right panel depicts the cross-correlation function (lag: 0–19) of lagged price changes with order flow. The dotted lines mark the approximate 99% confidence bounds, computed as $\frac{\pm 2.58}{\sqrt{T}}$, where T denotes the particular number of observations.

These descriptive analysis shows that the dynamic properties of the OANDA based price series differ from those of the interbank price series. To understand these differences, we decide to have a closer look at the relationship between these price series by means of a bivariate vector error correction (VEC) model. This investigation is again based on the price series sampled at a 1 minute frequency. Since the OANDA based price series are derived from an external data-feed of the interbank market, we expect that our interbank price series and the OANDA based price series are co-integrated. This hypothesis is verified with the ? Co-integration test which indicates one co-integrating equation even at the 1% significance level. Since our interbank price series is not the external data-feed interbank price series on which the OANDA prices are based, we cannot expect to figure out how the OANDA price series is derived from the interbank price series. The VEC model can be formulated in the following way. Let $y_t = (y_t^{\text{OP}}, y_t^{\text{IP}})'$ denote the vector of OANDA and interbank prices at time t for $t = 1, \dots, T$. Let $\beta = (\beta_{\text{OP}}, \beta_{\text{IP}})'$ with $\beta_{\text{OP}} = 1$ denote the coefficient vector of the co-integrating equation which is assumed to take the following form:

$$z_t = c + \beta' y_t,$$

where z_t denotes the co-integrating error and the associated VEC model is given by

$$\Delta y_t = \lambda z_{t-1} + \Phi_p(L) \Delta y_{t-1} + \varepsilon_t,$$

with $\lambda = (\lambda_{\text{OP}}, \lambda_{\text{IP}})'$ denoting the adjustment coefficients. $\Phi_p(L)$ denotes the lag-polynomial of order p consisting of matrices $\Phi^{(i)} = \begin{pmatrix} \phi_{11}^{(i)} & \phi_{12}^{(i)} \\ \phi_{21}^{(i)} & \phi_{22}^{(i)} \end{pmatrix}$ where $i = 1, \dots, p$. ε_t is assumed to be an independent bivariate normally distributed error term process with zero mean. The estimation results are summarized in Table 4, where the number of lags $p = 6$ are chosen according to the Schwarz Information Criterion (SIC). The most important observation that should be made is that both adjustment coefficients $(\lambda_{\text{OP}}, \lambda_{\text{IP}})$ are significantly different from zero and have opposite signs, implying that after a shock in the co-integrating error z_t , both price series respond to this shock aiming to get back to their equilibrium relationship. This means, from an economic point of view, that we do not observe a lead-lag relationship between our two price series on a 1 min aggregation level, which might have been expected since the interbank data-feed might have caused the price process on OANDA FXTrade. The non-existence of a lead-lag relationship, however, can be explained by the fact that we compare mid-quotes and mid-transaction prices on a 1 min frequency, in

which the lead-lag structure might already be aggregated away since we do not know how exactly and more importantly on which frequency the OANDA price process relies on the interbank market data-feed.

Parameter	Estimates	Standard Deviation
c	0.00037	0.00011
β_{OP}	1.0000	
β_{IP}	-1.00031	0.00009
λ	(-0.05284 0.035260)	(0.00201 0.00175)
$\Phi^{(1)}$	$\begin{pmatrix} -0.11204 & 0.30832 \\ 0.33177 & -0.13975 \end{pmatrix}$	$\begin{pmatrix} 0.00345 & 0.00384 \\ 0.00301 & 0.00335 \end{pmatrix}$
$\Phi^{(2)}$	$\begin{pmatrix} -0.18967 & 0.13512 \\ 0.11037 & -0.14589 \end{pmatrix}$	$\begin{pmatrix} 0.00371 & 0.00404 \\ 0.00323 & 0.00352 \end{pmatrix}$
$\Phi^{(3)}$	$\begin{pmatrix} -0.11552 & 0.10699 \\ 0.09116 & -0.08426 \end{pmatrix}$	$\begin{pmatrix} 0.00375 & 0.00407 \\ 0.00327 & 0.00355 \end{pmatrix}$
$\Phi^{(4)}$	$\begin{pmatrix} -0.08784 & 0.07205 \\ 0.05672 & -0.05572 \end{pmatrix}$	$\begin{pmatrix} 0.00370 & 0.00401 \\ 0.00323 & 0.00349 \end{pmatrix}$
$\Phi^{(5)}$	$\begin{pmatrix} -0.06554 & 0.05206 \\ 0.02927 & -0.03138 \end{pmatrix}$	$\begin{pmatrix} 0.00354 & 0.00384 \\ 0.00309 & 0.00335 \end{pmatrix}$
$\Phi^{(6)}$	$\begin{pmatrix} -0.03921 & 0.03255 \\ 0.01740 & -0.00773 \end{pmatrix}$	$\begin{pmatrix} 0.00319 & 0.00337 \\ 0.00278 & 0.00294 \end{pmatrix}$

Table 4: Estimation Results of the VEC Model. The parameters estimates for the co-integrating equation are given in the upper part of the table and the parameters of the associated VEC model in the lower part.

6.2 Verification of the Economic Hypotheses

Although the descriptive analysis already provides some insights into the dynamic relationship between order flow and price changes allowing for a first idea on the validity of the hypotheses raised in Chapter 5, we investigate them now in detail with the help of forecasting analyses. The hypotheses H1, H1.1 and H1.2, being concerned with the causality direction from order flow to price changes are verified in the following way. We conduct a forecasting study that investigates whether prices are better predictable using information on the order flow in addition to the information already contained in the history of the price process itself than using information contained in the history of the price process solely (benchmark model). Based on these two forecasting models, we compute their Root-Mean-Squared-Prediction-Errors (RMSPE) and analyze whether the model incorporating information on order

flow provides significantly better forecasts than the benchmark model with the help of the modified Diebold-Mariano (mDM) test of Harvey et al. (1997). The forecasting study is performed on 12 intra-day sampling frequencies stated in the first column of Table 5. Since on an intra-day level, there exists a specific autoregressive structure in the price change processes, as shown in the descriptive analysis, we decided not to use a random-walk specification, like Daniélsson et al. (2002), as the benchmark model but the following AR(p) specification:

$$(1 - B_p^y(L)) \Delta y_t^h = c + \varepsilon_t, \quad (\text{BM-}h)$$

where $B_p^y(L)$ denotes the associated lag-polynomial specified as

$$B_p^y(L) = \beta_{y1}L + \dots + \beta_{yp}L^p,$$

with ε_t a white noise process. The forecasting study is implemented once for the interbank price change process ($h = \text{IP}$) and once for the OANDA based price change process ($h = \text{OP}$). The optimal lag length p is again chosen according to the SIC. In order to verify whether order flow containing information on price expectations from the interbank market and OANDA FXTrade is helpful in predicting future currency prices, (H1.1), we use the following forecasting model

$$(1 - B_p^y(L)) \Delta y_t^h = c + B_p^x(L)x_t^{\text{SOF}} + \varepsilon_t, \quad (\text{SOF-}h)$$

in which, in addition to the benchmark model, the history of the standard order flow measure (x_t^{SOF}) is included. $B_p^x(L)$ denotes the corresponding lag-polynomial.

In order to verify whether order flow that contains information on price expectations from OANDA FXTrade solely is helpful in predicting future currency prices, (H1.2), we use the following forecasting model

$$(1 - B_p^y(L)) \Delta y_t^h = c + B_p^x(L)x_t^{\text{OOF}} + \varepsilon_t, \quad (\text{OOF-}h)$$

in which, in addition to the benchmark model, the history of the OANDA order flow measure (x_t^{OOF}) is included. Furthermore, we use a more flexible forecasting specification in which we include the order flows of the eight trading categories (x_t^k , $k = 1, \dots, 8$) separately

$$(1 - B_p^y(L)) \Delta y_t^h = c + \sum_{k=1}^8 B_p^{xk}(L)x_t^k + \varepsilon_t. \quad (\text{CAT-}h)$$

Table 5 presents the results of the forecasting studies, where the different order flow measures are used in deseasonalized form to predict OANDA based (OP) and interbank (IP) price changes. The seasonality functions of the order flow measures

are assumed to be additive and estimated for each frequency separately with a Nadaraya-Watson kernel regression using a Gaussian kernel and an optimal bandwidth selection. The out-of-sample prediction horizon, though for different frequencies, corresponds to the week from 9th May 2004 to 14th May 2004. The model selection period covers the period from 1st October 2003 to 8th May 2004. The results presented in Table 5 are robust to reasonable selections of the in-sample and out-of-sample horizon and do not change the results qualitatively. Weekends and holidays are excluded from the analysis. The first cell entry in Table 5 is the RMSPE of the associated forecasting model. The second cell entry in parenthesis is the p-value from the mDM test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the corresponding benchmark model. Cell entries in bold are those where the RMSPE of the associated forecasting model is smaller than the RMSPE of the corresponding benchmark model.

First of all, we notice that the forecasting models incorporating the standard order flow measure (SOF- h) have a higher forecasting power, in terms of delivering smaller RMSPEs than the benchmark models' (bold cell entries), for both price series ($h \in \{OP, IP\}$) than models relying on the OANDA order flow measure (OOF- h) and models including all eight order flow categories (CAT- h) separately. Being better than the benchmark model is concentrated for the latter two models on short forecasting horizons (1 and 2 minutes), whereas the SOF- h models deliver better forecasts for basically all investigated frequencies. Thereby, we observe the tendency that the interbank price change process can be predicted slightly better than the OANDA based price change process. This observation is confirmed considering a 10% significance level for the mDM test where we observe that the SOF- h forecasting models deliver significantly better (than those of the benchmark models') 20 and 25 minutes forecasts for the OANDA based price change process, and better 1, 20 and 25 minutes forecasts for the interbank price change process. The OANDA based price change process, however, can be predicted on a 1 minute level by the OOF-OP and the CAT-OP forecasting models significantly better, even on a 1% significance level, than by the benchmark model. For the interbank price change process, the same observation is only valid for the CAT-IP forecasting model. Altogether, we can conclude that expectations of the market participants revealed through the order flow process help to predict future currency price changes. In particular, we observe that the standard order flow measure in which information from the interbank market, in addition to information from OANDA on price expectations, is included is most beneficial. However, we observe that for short term price change predictions infor-

mation stemming solely from OANDA measured through the OANDA order flow measure is useful as well. Since the conducted out-of-sample forecasting study can be considered as judging additional information contained in the order flow measures, in addition to the information contained in the price change process itself, in a very rigorous way we infer that hypotheses H1, H1.1, and H1.2 cannot be rejected. Considering the joint significance of the related coefficients in the regression outputs for the in-sample fit of the considered forecasting models yields the same result.⁵

Freq	BM-OP	SOF-OP	OOF-OP	CAT-OP	BM-IP	SOF-IP	OOF-IP	CAT-IP
1 min	0.2045	0.2044 (0.3154)	0.2021 (0.0003)	0.2007 (0.0000)	0.1653	0.1648 (0.0967)	0.1646 (0.1798)	0.1618 (0.0000)
2 min	0.3068	0.3070 (0.7853)	0.3067 (0.4251)	0.3065 (0.3102)	0.2721	0.2718 (0.1725)	0.2724 (0.8027)	0.2725 (0.7178)
5 min	0.4799	0.4800 (0.6287)	0.4805 (0.7588)	0.4819 (0.9088)	0.4541	0.4541 (0.4315)	0.4542 (0.8702)	0.4573 (0.9627)
10 min	0.7016	0.7017 (0.6746)	0.7079 (0.9777)	0.7051 (0.8908)	0.6824	0.6814 (0.1905)	0.6862 (0.9662)	0.6839 (0.7470)
15 min	0.9301	0.9307 (0.6357)	0.9390 (0.9953)	0.9315 (0.6718)	0.9281	0.9266 (0.1081)	0.9325 (0.9590)	0.9253 (0.2131)
20 min	0.9741	0.9668 (0.0288)	0.9842 (0.9592)	0.9799 (0.8674)	0.9665	0.9604 (0.0368)	0.9752 (0.9655)	0.9725 (0.8734)
25 min	1.0653	1.0607 (0.0262)	1.0953 (0.9983)	1.0812 (0.9940)	1.0794	1.0758 (0.0560)	1.1017 (0.9972)	1.0900 (0.9600)
30 min	1.2963	1.2922 (0.1141)	1.3124 (0.9921)	1.2899 (0.2485)	1.3059	1.3025 (0.1211)	1.3170 (0.9932)	1.2952 (0.1782)
45 min	1.4806	1.4796 (0.2532)	1.4989 (0.9904)	1.4928 (0.9696)	1.4908	1.4901 (0.2597)	1.5050 (0.9990)	1.4999 (0.9213)
1 hr	1.6610	1.6616 (0.5338)	1.6768 (0.9813)	1.6875 (0.9677)	1.6939	1.6947 (0.5540)	1.7069 (0.9750)	1.7177 (0.9523)
2 hr	2.2541	2.2479 (0.1848)	2.2966 (0.9256)	2.2638 (0.7321)	2.2904	2.2853 (0.1982)	2.3287 (0.9191)	2.2986 (0.7176)
4 hr	2.8072	2.7945 (0.1780)	3.1587 (0.9517)	3.2069 (0.9639)	2.8165	2.8078 (0.2983)	3.1496 (0.9499)	3.2021 (0.9630)

Table 5: Results for the price change out-of-sample prediction on different sampling frequencies (Freq). The out-of-sample prediction horizon, though for different frequencies, corresponds to the week from 9th May 2004 to 14th May 2004. The model selection period covers the period from 1st October 2003 to 8th May 2004. Weekends and holidays are excluded from the analysis. The first cell entry is the Root-Mean-Squared-Prediction Error (RMSPE) of the associated forecasting model. The second cell entry in parenthesis is the p-value from the modified Diebold-Mariano Test with the null hypothesis, that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the corresponding benchmark model. Cell entries in bold are those, where the RMSPE of the associated forecasting model is smaller than the RMSPE of the corresponding benchmark model.

We now investigate the economic hypotheses which consider the causal relationships from price changes to order flow with forecasting set-ups similar to those applied

⁵Due to space limitations and due to the more meaningful and powerful result of the out-of-sample study we decided to omit these regression outputs.

before. Based on the results of the descriptive analysis, we consider for the standard, the OANDA based and the category specific order flow measures benchmark models in which, based on $AR(p)$ specifications, only the histories of the order flow measures themselves serve to explain and to predict future order flows. These predictions are then compared, using the mDM test, to the predictions of those forecasting models in which, in addition to the information contained in the history of the order flows, the information contained in the history of the price change processes is incorporated. This proceeding enables us to figure out whether the additional information contained in the past prices is helpful to improve order flow forecasts significantly. The benchmark models are given by

$$(1 - B_p^x(L)) x_t^k = c + \varepsilon_t, \quad (\text{BM-}k)$$

where $B_p^x(L)$ denotes the associated lag-polynomial and ε_t a white noise process. The forecasting study is implemented for the standard order flow measure ($k = \text{SOF}$), for the OANDA order flow measure ($k = \text{OOF}$), and the eight category specific order flow measures ($k = 1, \dots, 8$) already listed in Table 3. The forecasting models containing, in addition, information on the history of the price change process are given by

$$(1 - B_p^x(L)) x_t^k = c + B_p^y(L) \Delta y_t^h + \varepsilon_t, \quad (h-k)$$

where $h \in \{\text{IP}, \text{OP}\}$ denotes whether the interbank or the OANDA based price change process is included.

The RMSPEs and the p-values of the mDM tests are presented for the models above in Tables 6 to 8. We observe that the general hypothesis H2, which claims that the information contained in the price process is helpful in predicting future order flow, cannot be rejected, since the information contained in the history of the price process in addition to the information contained in the order flow measures themselves is both helpful in predicting both aggregated order flow measures (standard and OANDA) and helpful in predicting the eight category specific order flow measures. This statement is based on the observations that i) for the prediction of the standard order flow measure, the RMSPEs are for 6 (7) forecasting horizons smaller than those of the benchmark model, when additional information on the OANDA based (interbank) price change process is incorporated in the forecasting models. Irrespectively of the choice of the price series we see that the 2 (2) RMSPEs for 1 and 2 minutes forecasting horizons are significantly smaller using a 5% significance level in the mDM test. ii) for the prediction of the OANDA order flow measure we

observe 9 (10) smaller RMSPEs in comparison to the benchmark model ones; 5 (5) of them are significantly smaller at the 5% level and even 6 (6) at the 10% level. iii) for the prediction of the category specific order flow we observe that over all eight categories 54 (55) RMSPEs are smaller than the benchmark models'; 23 (23) of them are significantly smaller at a 5% level and even 29 (29) at the 10% level. Interpreting these figures, two interesting observations can be derived: first, we basically do not observe any information advantages between the OANDA based price series and the interbank price series since the forecasting models incorporating these series generate very similar RMSPE patterns. Therefore we do not observe that the OANDA based price series influences the traders on OANDA FXTrade more severely than the interbank price series. This can be explained by the fact that both price series are co-integrated and therefore convey closely related dynamic patterns and that in the absence of macroeconomic news and private (customer order flow) information, which is certainly true for most traders on OANDA FXTrade, traders update their beliefs and place their orders based on interpretations of technical analysis patterns. The second observation, that the OANDA order flow measure is better predictable than the standard order flow measure, points in the same direction. Here, we observe that the information contained in the price process has more influence on the price expectation process of the OANDA market, represented by the OANDA order flow measure than on the price expectation process of the OANDA and the interbank market, represented by the standard order flow measure. This is intuitively clear, since information on the price process is more valuable on the OANDA market than on the interbank market where private information is available as well.

Let us now consider hypothesis H3 that executed stop-loss orders (category 6) contribute whereas executed take-profit orders (category 7) impede, self-reinforcing price movements. Table 8 shows that on all considered forecasting frequencies, the RMSPEs of the benchmark models BM-6 are smaller than those of benchmark models BM-7. These RMSPEs can be compared with each other since deseasonalized (standardized) order flows are considered. Furthermore, comparing the RMSPE pattern of the OP-6 (IP-6) forecasting models with those of the OP-7 (IP-7) forecasting models, we observe that the information contained in the histories of both price processes is more valuable in terms of significance on the 5% level and longer forecasting horizons, to predict take-profit order flow than to predict stop-loss order flow. Both observations support hypothesis H3, since first stop-loss order flow is better predictable than take-profit order flow (BM-6 vs. BM-7 RMSPEs), which is natural if stop-loss orders contribute to self-reinforcing price movements causing

a sequence of stop-loss order executions. The second observation is related to the fact that if stop-loss orders contribute to self-reinforcing price movements then the information on the direction of the price change process is already included in the historical stop-loss order flow, and, therefore is of less importance than in the case of historical take-profit order flow, which contains less information on the direction of the price change process.

Hypothesis H4, that the traders' degree of (im-)patience depends on their actual investment status, which can be related to the existence of a zero-profit reference point as a basis for investment decisions or to the existence of a monitoring effect, is clearly supported. Indeed, we see that the price process contributes more to the prediction of market order (close) order flow (OP-5, IP-5) than to the prediction of market order (open) order flow (OP-4, IP-4). In detail, we observe that no RMSPE of the OP-4 (IP-4) forecasting models are significantly smaller (10% level) than those of the benchmark models BM-4, but that 8 (8) RMSPEs are significantly smaller than those of the benchmark models BM-5 at the 10% and even 6 (6) at the 1% significance level of the mDM test for the OP-5 (IP-5) forecasting models. A similar, but not that pronounced, observation can be made for limit order executed (close) order flow (OP-3, IP-3) and limit order executed (open) order flow (OP-2, IP-2) as well. The reason this effect is not as clear as for market orders is that limit orders are posted to the system before, and their execution is later simply implied by the price process. Market orders, however, reflect changes in price preferences directly since they are executed immediately.

Freq	BM-SOF	OP-SOF	IP-SOF	BM-OOF	OP-OOF	IP-OOF
1 min	4.6688	4.6469 (0.0311)	4.6458 (0.0302)	6.7677	6.6609 (0.0000)	6.6569 (0.0000)
2 min	7.0391	6.9893 (0.0381)	6.9757 (0.0098)	11.4667	11.2540 (0.0000)	11.2477 (0.0000)
5 min	12.9287	12.9314 (0.5145)	12.9200 (0.4543)	23.9086	23.3968 (0.0142)	23.4035 (0.0161)
10 min	18.4695	18.4228 (0.3469)	18.4267 (0.3564)	41.4958	40.6636 (0.0244)	40.6788 (0.0265)
15 min	21.5305	21.4564 (0.2987)	21.4669 (0.3235)	59.8002	59.6020 (0.2952)	59.5754 (0.2578)
20 min	26.1122	26.3509 (0.9840)	26.3696 (0.9893)	64.3401	63.4812 (0.0682)	63.4739 (0.0578)
25 min	28.7964	28.8469 (0.6607)	28.8527 (0.6698)	89.4254	88.8717 (0.2446)	88.8282 (0.2335)
30 min	33.9792	34.1016 (0.6859)	34.1745 (0.7655)	97.7076	97.7158 (0.5050)	97.6498 (0.4647)
45 min	44.5986	44.9963 (0.7555)	45.0139 (0.7666)	130.3802	129.1925 (0.0340)	129.0866 (0.0244)
1 hr	51.0773	50.5123 (0.1395)	50.5101 (0.1368)	147.5702	147.1536 (0.3846)	147.2210 (0.4017)
2 hr	79.4658	78.3111 (0.1424)	78.1480 (0.1300)	268.4134	269.2610 (0.6156)	269.2955 (0.6307)
4 hr	134.0885	136.1753 (0.7735)	136.1745 (0.7748)	437.2603	455.4722 (0.9351)	454.7048 (0.9320)

Table 6: Results for the standard and the OANDA order flow measures out-of-sample predictions on different sampling frequencies (Freq). The out-of-sample prediction horizon, though for different frequencies, corresponds to the week from 9th May 2004 to 14th May 2004. The model selection period covers the period from 1st October 2003 to 8th May 2004. Weekends and holidays are excluded from the analysis. The first cell entry is the Root-Mean-Squared-Prediction Error (RMSPE) of the associated forecasting model. The second cell entry in parenthesis is the p-value from the modified Diebold-Mariano Test with the null hypothesis, that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the corresponding benchmark model. Cell entries in bold are those, where the RMSPE of the associated forecasting model is smaller than the RMSPE of the corresponding benchmark model.

Freq	BM-1	OP-1	IP-1	BM-2	OP-2	IP-2	BM-3	OP-3	IP-3	BM-4	OP-4	IP-4
1 min	2.1722	2.1746 (0.9393)	2.1745 (0.9316)	3.4387	3.4362 (0.0002)	3.4355 (0.0010)	1.3958	1.3947 (0.0431)	1.3945 (0.0370)	2.0807	2.0781 (0.2077)	2.0786 (0.2625)
2 min	3.7150	3.7108 (0.0597)	3.7110 (0.0581)	5.3822	5.3782 (0.0167)	5.3775 (0.0234)	2.5072	2.5010 (0.1153)	2.5007 (0.1240)	3.1875	3.1794 (0.1634)	3.1818 (0.2728)
5 min	8.8940	8.8793 (0.1358)	8.8802 (0.1571)	9.8877	9.8840 (0.1003)	9.8844 (0.1164)	3.2849	3.2780 (0.0269)	3.2775 (0.0328)	5.7850	5.8339 (0.9930)	5.8367 (0.9941)
10 min	15.9458	15.9969 (0.7190)	15.9977 (0.7216)	16.1545	16.1806 (0.8898)	16.1810 (0.8980)	5.8855	5.8738 (0.0660)	5.8725 (0.0706)	9.3110	9.4171 (0.9976)	9.4241 (0.9979)
15 min	25.0806	25.2198 (0.8366)	25.2205 (0.8383)	25.2300	25.2238 (0.2903)	25.2188 (0.1583)	7.2595	7.2549 (0.1641)	7.2540 (0.1307)	12.9731	13.1843 (0.9996)	13.1846 (0.9997)
20 min	29.6264	29.7877 (0.7363)	29.7904 (0.7375)	25.3197	25.3495 (0.9574)	25.3551 (0.9625)	8.5186	8.5056 (0.0588)	8.5051 (0.0432)	14.7481	15.0686 (0.9959)	15.0721 (0.9953)
25 min	37.6371	38.0371 (0.8715)	38.0364 (0.8705)	34.3930	34.4073 (0.8013)	34.3992 (0.6670)	9.6302	9.6246 (0.1972)	9.6237 (0.1870)	18.7257	19.1674 (0.9986)	19.1665 (0.9987)
30 min	40.7652	41.1303 (0.7949)	41.1314 (0.7955)	38.5794	38.5063 (0.2858)	38.5106 (0.2959)	10.9258	10.9404 (0.9664)	10.9407 (0.9585)	21.8613	22.2614 (0.9955)	22.2944 (0.9967)
45 min	57.4849	57.5383 (0.7639)	57.5525 (0.8184)	48.4388	48.4211 (0.4787)	48.4215 (0.4793)	13.5863	13.5960 (0.9160)	13.5957 (0.9031)	27.0817	27.6538 (0.9999)	27.6854 (0.9999)
1 hr	69.8865	69.9538 (0.5723)	69.9603 (0.5790)	56.2525	56.3085 (0.8690)	56.2995 (0.8256)	15.2791	15.2466 (0.3595)	15.2482 (0.3644)	30.7204	30.4386 (0.1273)	30.4557 (0.1340)
2 hr	114.5187	114.4591 (0.2353)	114.4610 (0.2385)	107.6508	105.0767 (0.0573)	105.0904 (0.0577)	21.9844	21.9295 (0.1749)	21.9309 (0.1993)	53.0909	52.7008 (0.1073)	52.6834 (0.1044)
4 hr	231.6298	219.7287 (0.0076)	219.7121 (0.0077)	140.9649	139.7789 (0.1268)	139.8010 (0.1325)	31.5516	31.4473 (0.1972)	31.4506 (0.2149)	94.0589	98.9859 (0.9258)	99.0657 (0.9281)

Table 7: Results of the out-of-sample predictions of the category specific order flow measures on different sampling frequencies (Freq). The out-of-sample prediction horizon, though for different frequencies, corresponds to the week from 9th May 2004 to 14th May 2004. The model selection period covers the period from 1st October 2003 to 8th May 2004. Weekends and holidays are excluded from the analysis. The first cell entry is the Root-Mean-Squared-Prediction Error (RMSPE) of the associated forecasting model. The second cell entry in parenthesis is the p-value from the modified Diebold-Mariano Test with the null hypothesis, that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the corresponding benchmark model. Cell entries in bold are those, where the RMSPE of the associated forecasting model is smaller than the RMSPE of the corresponding benchmark model.

Freq	BM-5	OP-5	IP-5	BM-6	OP-6	IP-6	BM-7	OP-7	IP-7	BM-8	OP-8	IP-8
1 min	2.8626	2.7728 (0.0000)	2.7590 (0.0000)	1.6361	1.6228 (0.0000)	1.6231 (0.0000)	5.0943	5.0860 (0.2597)	5.0784 (0.1156)	0.1816	0.1861 (0.9999)	0.1855 (0.9997)
2 min	4.5178	4.3253 (0.0000)	4.3080 (0.0000)	2.5981	2.5684 (0.0000)	2.5680 (0.0000)	9.1026	9.0491 (0.0037)	9.0437 (0.0036)	0.2810	0.2893 (0.9999)	0.2890 (0.9999)
5 min	9.2733	8.8477 (0.0000)	8.8296 (0.0000)	4.4767	4.4216 (0.0000)	4.4207 (0.0000)	17.2736	17.1766 (0.0177)	17.1859 (0.0337)	0.5061	0.5495 (0.9999)	0.5483 (0.9998)
10 min	15.6532	15.0850 (0.0000)	15.0845 (0.0000)	7.0464	6.9911 (0.0185)	6.9924 (0.0096)	27.9659	27.8622 (0.2326)	27.8771 (0.2792)	0.7886	0.8376 (0.9847)	0.8388 (0.9869)
15 min	21.4368	21.1433 (0.0088)	21.0996 (0.0037)	9.7075	9.7287 (0.7452)	9.7226 (0.6916)	43.5326	42.9314 (0.0457)	43.2969 (0.0356)	0.9112	0.9872 (0.9599)	0.9944 (0.9681)
20 min	23.3513	22.8615 (0.0078)	22.8368 (0.0034)	10.5435	10.5335 (0.3471)	10.5324 (0.3149)	42.9012	42.4084 (0.0415)	42.3916 (0.0379)	1.2009	1.2956 (0.9334)	1.2956 (0.9169)
25 min	35.4461	35.0990 (0.0908)	35.0472 (0.0828)	12.8666	12.8835 (0.6648)	12.8804 (0.6555)	58.6815	57.7769 (0.0152)	57.7330 (0.0170)	1.3506	1.4445 (0.8807)	1.4457 (0.8675)
30 min	36.4444	36.6709 (0.7798)	36.6503 (0.7600)	15.1755	15.2622 (0.8404)	15.2994 (0.9041)	67.6763	66.9273 (0.0183)	66.8867 (0.0179)	1.5422	1.7742 (0.9911)	1.7868 (0.9920)
45 min	47.9285	47.3010 (0.0715)	47.3018 (0.0708)	18.6784	18.6844 (0.8734)	18.6105 (0.1769)	83.6687	82.3866 (0.0252)	82.4131 (0.0242)	1.8783	2.1214 (0.9717)	2.1430 (0.9811)
1 hr	53.5001	53.1720 (0.1936)	53.1447 (0.1798)	20.7151	20.9582 (0.8162)	20.9333 (0.8825)	92.8033	94.1736 (0.9423)	94.1609 (0.9384)	2.3081	2.6650 (0.9995)	2.6716 (0.9997)
2 hr	109.3541	109.0568 (0.1880)	109.0589 (0.1663)	26.4431	26.1510 (0.2790)	26.1066 (0.2525)	169.3417	169.2100 (0.0854)	169.3232 (0.0874)	3.7456	3.4597 (0.2398)	3.4308 (0.2245)
4 hr	149.9129	150.1995 (0.6258)	150.1661 (0.6155)	35.5663	36.3020 (0.8966)	36.2503 (0.8758)	215.9333	214.6672 (0.1597)	214.7852 (0.1696)	4.1827	6.6095 (0.9457)	6.6356 (0.9497)

Table 8: Results of the out-of-sample predictions of the category specific order flow measures on different sampling frequencies (Freq). The out-of-sample prediction horizon, though for different frequencies, corresponds to the week from 9th May 2004 to 14th May 2004. The model selection period covers the period from 1st October 2003 to 8th May 2004. Weekends and holidays are excluded from the analysis. The first cell entry is the Root-Mean-Squared-Prediction Error (RMSPE) of the associated forecasting model. The second cell entry in parenthesis is the p-value from the modified Diebold-Mariano Test with the null hypothesis, that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the corresponding benchmark model. Cell entries in bold are those, where the RMSPE of the associated forecasting model is smaller than the RMSPE of the corresponding benchmark model.

7 Conclusion

We investigate the relationship between currency price changes and their expectations with the help of a customer data-set from OANDA FXTrade. We infer price expectations using order flow measures from the trading behavior on OANDA FXTrade. We distinguish between price expectations relying on information from the interbank market and OANDA FXTrade, which are measured with the standard order flow measure of Lyons (1995) and price expectations derived solely from OANDA FXTrade measured by the OANDA order flow measure. We conduct forecasting studies on 12 intra-day frequencies and we find that those forecasting models which incorporate information on order flow and price change provide significantly better forecasts than benchmark models, which use only information on past price changes through $AR(p)$ specifications. In comparison to the Root-Mean-Squared-Prediction Errors (RMSPE) of the benchmark specifications, forecasting models relying on the OANDA order flow measure provide smaller RMSPEs for 1 and 2 minutes horizons, and models relying on the standard order flow measure yield even smaller RMSPEs for horizons up to 4 hours.

Applying a similar forecasting set-up allows to investigate the influence of past price changes on these two order flow measures and on eight transaction category specific order flow measures. We find i) that the trading behavior, and therefore the price expectations, of our investors are affected by past currency price changes, ii) that stop-loss orders contribute to and take-profit impede self-reinforcing price movements, which supports the hypothesis of Osler (2002) using a different methodology, and, iii) that the degree of investors' impatience depends on his current investment status, which has several consequences for theoretical market microstructure models in which patient investors are assumed to submit limit orders and impatient investors are assumed to submit market orders.

References

- ANDERSEN, T. G. & T. BOLLERSLEV (1997): “Intraday Periodicity and Volatility Persistence in Financial Markets,” *Journal of Empirical Finance*, 4, 115–158.
- BANK FOR INTERNATIONAL SETTLEMENTS’ (BIS) REPORT (2004): “Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2004,” *Bank for International Settlements*.
- BARBERIS, N. & R. THALER (2003): “A Survey of Behavioral Finance,” *in*: CONSTANTINIDES, G., M. HARRIS, & R. STULZ, eds., “Handbook of the Economics of Finance,” Elsevier Science B.V., pp. 1052–1121.
- BIAIS, B., P. HILLION, & 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.
- BJØNNES, G. H. & D. RIME (2003): “Dealer Behavior and Trading Systems in Foreign Exchange Markets,” A Research Report from Stockholm Institute for Financial Research, No 17 – December 2003.
- CARLSON, J. A. & M. PARKIN (1975): “Inflation Expectations,” *Economica*, 42, 123–138.
- DACOROGNA, M. M., R. GENÇAY, U. MÜLLER, R. B. OLSEN, & O. V. PICTET (2001): *An Introduction to High-Frequency Finance*, Academic Press, London.
- DANÍELSSON, J., R. PAYNE, & J. LUO (2002): “Exchange Rate Determination and Inter-Market Order Flow Effects,” Working Paper, Financial Markets Group, London School of Economics.
- DEMSETZ, H. (1968): “The Cost of Transacting,” *Quarterly Journal of Economics*, 82, 33–53.
- EASLEY, D. & M. O’HARA (1992): “Time and Process of Security Price Adjustment,” *The Journal of Finance*, 47 (2), 577–605.
- EVANS, M. D. & R. K. LYONS (2002a): “Informational Integration and FX Trading,” *Journal of International Money and Finance*, 21, 807–831.
- (2002b): “Order Flow and Exchange Rate Dynamics,” *Journal of Political Economy*, 110, 170–180.

- FOSTER, D. F. & S. VISWANATHAN (1990): “A Theory of the Intraday Variations in Volumes, Variances and Trading Costs in Security Markets,” *Review of Financial Studies*, 3, 593–624.
- GLOSTEN, L. R. (1994): “Is the Electronic Open Limit Order Book Inevitable?” *The Journal of Finance*, 49 (4), 1127–1161.
- HARVEY, D., S. LEYBOURNE, & P. NEWBOLD (1997): “Testing the Equality of Prediction Mean Squared Errors,” *International Journal of Forecasting*, 13, 281–291.
- HO, T. & H. STOLL (1981): “Optimal Dealer Pricing Under Transactions and Return Uncertainty,” *Journal of Financial Economics*, 9, 47–73.
- KAHNEMAN, D. & A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–292.
- KYLE, A. (1985): “Continuous Auctions and Insider Trading,” *Econometrica*, 22, 477–498.
- LYONS, R. (1995): “Tests of Microstructural Hypothesis in the Foreign Exchange Market,” *Journal of Financial Economics*, 39, 321–351.
- LYONS, R. K. (2001): *The Microstructure Approach to Exchange Rates*, M.I.T. Press, Cambridge, MA.
- (2002): “The Future of the Foreign Exchange Market,” Brookings-Wharton Papers on Financial Services.
- MARSH, I. W. & C. O’ROURKE (2004): “Customer order flows in foreign exchange markets: Do they really contain information?” Working Paper, Cass Business School, London.
- MEESE, R. & K. ROGOFF (1983a): “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics*, 14, 3–24.
- (1983b): “The Out-of-sample Failure of Empirical Exchange Rate Models,” *in*: FRENKEL, J., ed., “Exchange Rate and International Macroeconomics,” vol. 14, NBER and University of Chicago Press, pp. 67–105.

- MUTH, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29 (3), 315–335.
- NERLOVE, M. (1983): "Expectations, Plans, and Realizations in Theory and Practice," *Econometrica*, 51 (5), 1251–1279.
- ODEAN, T. (1998): "Are investors reluctant to realize their losses?" *Journal of Finance*, 53, 1775–1798.
- OSLER, C. (2002): "Stop-Loss Orders and Price Cascades in Currency Markets," Staff Reports from Federal Reserve Bank of New York, No. 150.
- PAYNE, R. (2003): "Informed trade in spot foreign exchange markets: An empirical investigation," *Journal of International Economics*, 61, 307–329.
- PESARAN, M. H. (1987): *The Limits to Rational Expectations*, Basil Blackwell, Oxford.
- RIME, D. (2003): "New Electronic Trading Systems in the Foreign Exchange Markets," *New Economy Handbook*, 469–504.
- ROLL, R. (1984): "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *The Journal of Finance*, 39 (4), 1127–1139.
- SARNO, L. & M. P. TAYLOR (2001): "The Microstructure of the Foreign-Exchange Market: A Selective Survey of the Literature," Princeton Studies in International Economics, No. 89.
- SEPPI, D. J. (1997): "Liquidity Provision with Limit Orders and a Strategic Specialist," *Review of Financial Studies*, 10, 103–150.
- SILVERMAN, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, Chapman & Hall, London.
- TAYLOR, M. P. & H. ALLEN (1992): "The Use of Technical Analysis in the Foreign Exchange Market," *Journal of International Money and Finance*, 11, 304–314.
- WEBER, M. & C. F. CAMERER (1998): "The Disposition Effect in Securities Trading: An Experimental Analysis," *Journal of Economic Behavior & Organization*, 33, 167–184.