

Nolte, Ingmar; Nolte, Sandra

Conference Paper

How Do Individual Investors Trade?

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2010: Ökonomie der Familie - Session: Trading, Information, and Market Microstructure, No. G18-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Nolte, Ingmar; Nolte, Sandra (2010) : How Do Individual Investors Trade?, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2010: Ökonomie der Familie - Session: Trading, Information, and Market Microstructure, No. G18-V1, Verein für Socialpolitik, Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/37375>

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.

How Do Individual Investors Trade?

Ingmar Nolte*

Sandra Nolte[†]

Warwick Business School,

Warwick Business School,

FERC, CoFE

FERC, CMS

This Version: February 17, 2010

*Warwick Business School, Finance Group, Coventry, CV4 7AL, United Kingdom. Phone +44-24765-72838, Fax -23779, email: Ingmar.Nolte@wbs.ac.uk.

[†]Warwick Business School, Finance Group, Coventry, CV4 7AL, United Kingdom. Phone +44-24765-22853, Fax -23779, email: Sandra.Nolte@wbs.ac.uk. The work is supported by the Fritz Thyssen Foundation through the project 'Measurement Error in Nonlinear Regression Models and Business Tendency Data'.

The authors want to thank Günter Franke, Richard Olsen, Carol Osler, Richard Payne, Mark Salmon, and Winfried Pohlmeier for helpful comments and constructive discussions. Special thanks go to Olsen Financial Technologies for providing us with the data. The usual disclaimer applies.

Abstract

This paper examines how high-frequency trading decisions of individual investors are influenced by past price changes. Specifically, we address the question as to whether decisions to open or close a position are different when investors already hold a position compared to when they don't. Based on a unique dataset from an electronic foreign exchange trading platform, OANDA FXTrade, we find that investors' future order flow is (significantly) driven by past price movements and that these predictive patterns last up to several hours. This observation clearly shows that for high-frequency trading, investors rely on previous price movements in making future investment decisions. We provide clear evidence that market and limit orders flows are much more predictable if those orders are submitted to close an existing position than if they are used to open one. We interpret this finding as evidence for the existence of a monitoring effect, which has implications for theoretical market microstructure models and behavioral finance phenomena, such as the endowment effect.

JEL classification: G10, F31, C32

Keywords: Trading Activity Dataset, Order Flow, Foreign Exchange Market, Monitoring Effect

1 Introduction

In this paper we study how investors trade on a high-frequency time scale. We investigate how information on the past price process is fed back into the investors' trading decisions. Specifically, we examine whether investors' decisions to open or close a position are different if they already hold a position compared to when they don't. We also investigate whether stop-loss orders contribute to self-reinforcing price movements and whether take-profit orders impede them (Osler (2005)).

Our study concentrates on a large set of individual investors over a period of 8 months who trade currencies on an electronic trading platform: OANDA FXTrade. Most of the traders are small retail investors without access to private information such as own customer order flow or to news networks such as Reuters or Bloomberg. Electronic trading platforms such as OANDA FXTrade have become very popular during the last decade, since they provide immediate access to trading a large range of securities such as currencies, stocks, and options for retail investors by bypassing traditional trading venues including banks and brokers.¹ It is not surprising, therefore, that consortiums of banks have started to set up their own electronic trading platforms. Understanding the trading behavior of these investors is therefore of major importance for these companies when creating trading protocols and their own hedging algorithms.

Through in-sample and out-of-sample forecasting analyses we find that investors' future order flow is significantly affected by past price movements and that these predictive patterns last for several hours. This observation clearly shows that investors try to learn from previous price movements and exploit this knowledge for future high-frequency trading decisions, especially in the absence of private information, and news.² Our study does not permit us to analyse the extent to which this might be explained by the use of technical analysis or simply eyeball assessment of

¹Lyons (2002) has already pointed out, that there has been a shift in the interdealer market from direct trading towards electronic brokerage trading. This shift is partially explained by more transparency on electronic brokerage systems. In the customer market, a similar argument applies to explain the shift from customer-to-dealer-bank trading towards electronic internet trading platforms. These platforms are also more transparent and try to offer small (interbank) spreads to all of their customers independently of their transaction volume and thus order handling costs.

²The relationship between learning, feedback effects, information cascades, technical analysis and price bubbles is discussed for instance in De Long, Shleifer, Summers & Waldmann (1990), Bikhchandani, Hirshleifer & Welch (1992) Avery & Zemsky (1998), Lee (1998), and Shiller (2002).

potential price trends. We note, however, that Taylor & Allen (1992) report that about 90% of professional currency traders consider technical analysis as a valuable tool when designing their trading strategies. Our observations shed some light on the mechanisms of how market efficiency (within the retail trader segment) is actually obtained on a high frequency level while past prices seem to influence traders' decisions significantly for up to several hours. It turns out that market orders are easier to predict than limit orders.

We provide very clear evidence that prior price changes forecast order flows related to position closures better than order flows related to position initiations. We also show through in-sample investigations that the explanatory power of past price changes is about 3-5 times higher for position closures than position initiations. We interpret this finding as evidence for the existence of a monitoring effect. Such a monitoring effect can be related to the literature on referenced based decision making and endowment effects (c.f. Thaler (1980), Kahneman, Knetsch & Thaler (1990)). The endowment effect argues that an individual attributes a higher than its true value to a good once it becomes part of the individual's own endowment. In a situation where costs are related to the possession of a good for which the investor requires compensation, an endowment effect might be diminished and explained in a more rational way. Monitoring costs are exactly such costs especially in the context of security trading. Our observations can also be related to an asymmetric learning effect in the sense that the learning process is more intense, when investors already have a certain risk exposure through an open position.

The existence of a monitoring effect has implications for a large body of the market microstructure literature. A relatively recent branch of this literature focuses on optimal order placement strategy and on the question of the optimal mix at the equilibrium between market and limit orders. Parlour (1998) considers a dynamic model with strategic traders, explaining various patterns observed in order placement strategies and transaction prices. Hollifield, Miller & Sandas (2004) provide a model of optimal order submission, in which traders' optimal order placements depend on the valuation of the asset and the trade-offs between prices, execution probabilities and the risk of being picked off. Foucault (1999) derives the equilibrium in a trading game where traders arrive sequentially and choose to submit either a market or a limit order with a one-period life. He provides a complete closed form characterisation of traders order placement strategies, which allows him to analyse

the order flow composition (mix between market and limit orders) and trading costs. In his model he explicitly incorporates the risk of price misspecification into traders strategies. He also finds that the proportion of limit orders in order flow is positively related to asset return volatility. A cross-sectional analysis of order flow composition and trading costs in limit order markets allows him to test whether the proportion of limit orders in order flow is positively related to volatility or to the average size of the spread.

Foucault, Kadan & Kandel (2005) develop dynamic models of a limit order market populated by strategic liquidity traders. They analyse the question of how a trader's impatience affects order placement strategies, and show that at the equilibrium patient traders tend to provide liquidity to the impatient. They also address the question of waiting costs implied by limit orders, which is particularly important for traders who choose between limit and market orders (Demsetz (1968)), and show that liquidity suppliers tend to submit more aggressive limit orders to reduce their waiting time, and therefore their waiting costs. The idea of waiting costs is not new, Demsetz (1968) addressed this question: "*Waiting costs are relatively important for trading in organized markets, and would seem to dominate the determination of spreads*". Rosu (2009) observes that the waiting cost of a patient trader is smaller than the waiting cost of an impatient trader. He shows that the price impact of a market order is about four times higher than the price impact of a limit order.

While many of the studies mentioned above ignore the presence of information asymmetry, another branch of the theoretical literature considers such models in which informed traders and uninformed investors consider the trade-off between the use of market and limit orders in a dynamic setting. Among others Anand, Chakravarty & Martell (2005), Bloomfield, O'Hara & Saar (2005) and Kaniel & Liu (2006) show that limit orders play an important role in the order submission strategies of informed traders and therefore might convey information about future price movement and volatility. Chakravarty & Holden (1995), and Harris (1998) provide theoretical models in which informed investors are allowed to place both market and limit orders. A combined strategy of market and limit orders seems to be more profitable than only submitting market orders.

Neither this part of the literature nor that which ignores asymmetry of information addresses the question of whether order submission strategies are affected by

reference mechanisms such as holding an open position, or the current inventory of the traders, or by managing of active risk. Do traders exhibit different strategies when they hold an open position? The existence of a monitoring effect suggests that this question should be answered in the affirmative and that such effects should be addressed in theoretical market microstructure models.

Traditional studies on order flow (e.g., Evans & Lyons (2002a,b, 2005, 2006), Rime (2003) and Daniélsson, Payne & Luo (2002), Bjønnes & Rime (2005), Payne (2003), Berger et al. (2008)) 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 (2005) and Marsh & O'Rourke (2005) use a dataset of 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 for a given sampling frequency. Using standard instrumental variable techniques to estimate a vector autoregressive model that allows for contemporaneous feedback trading Daniélsson & Love (2006) show that when data are sampled at the one and five minute frequencies, the price impact of order flow is underestimated when feedback trading is not incorporated into the model, implying that trades carry more information than previous estimates suggest. All these studies concentrate on the question of how order flow can be used to predict future prices changes, but **not** in the way price changes affect order flow.

The paper is organized as follows: Section 2 explains the trading mechanism and the different order types on the OANDA FXTrade platform. Section 3 presents our economic hypotheses. Section 4 presents the empirical results and evidence towards the verification of the hypotheses, while Section 5 concludes.

2 OANDA FXTrade

OANDA FXTrade is an electronic trading platform for currencies, without limits on trade size, operating 24 hours, 7 days per week. This platform is a market mak-

ing system that executes orders using the exchange rate prevalent in the market determined either by their own inventory book and/or by predicted prices relying on a proprietary forecasting algorithm based on an external data-feed. OANDA FXTrade offers immediate settlement of trades and tight spreads as low as 2 to 3 pips for all transaction sizes. Given various boundary conditions, such as sufficient margin requirements, orders are always executed. The legal counterparty of a trade, however, is always OANDA FXTrade. The OANDA FXTrade platform is based on the concept of margin trading. This means that a trader can enter into positions larger than his available 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 a leverage 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, an automatic margin call order is used to close all open positions using the prevalent market rates at the current 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. A server manages the inventory book, the current exchange rates, and the current market orders to match existing limit orders. A 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. 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 as a result of further trading activity. Take-profit orders are typically set to close an existing position after a certain profit has been realized. Stop-loss orders, in contrast, specify that the position should be closed after the realization of a certain loss to avoid further losses. Table 1 provides an overview of the transactions and further activities of traders on OANDA FX-Trade, which are recorded in an activity record file. We get detailed information on whether an order is submitted to open (close) an existing position and thus reflects an increase (or decrease) in risk exposure.

The group of traders on OANDA is quite heterogeneous, varying from small retail investors, to small professional traders and smaller institutions and businesses. The attractiveness of OANDA FXTrade is that it provides instant access to currency trading without the involvement of banks, brokers or other intermediaries and thus without additional trading costs.

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 current market rates. This happens if the trader does not have 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 does not have 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 one buys the base currency (*EUR*) and sells the quote currency (*USD*), whereas selling *EUR/USD* means that one sells the base currency (*EUR*) and buys the quote currency (*USD*). Recorded units always refer to the base currency.

3 Motivation

We construct order flow using the standard definition given by Lyons (1995). He defines aggregated net order flow as the difference between buyer and seller initiated

trades (within a given period) or, stated differently, as the cumulative sum of signed trades where buyer initiated and seller initiated trades receive positive and negative signs, respectively. Focusing on the initiating party of a trade, this definition aims to capture 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 seller is treated as being more important than the buyer, who might not have the latest information. This standard measure of order flow has been used by Daniélsson et al. (2002) to predict future prices in the interbank market. Chordia & Subrahmanyam (2004) show that daily US equity order flow is helpful in predicting one period ahead price returns.

Transaction Record	Standard Order Flow Signs
Buy market (open)	+
Sell market (open)	-
Buy market (close)	+
Sell market (close)	-
Limit order: Buy	not used
Limit order: Sell	not used
Buy limit order executed (open)	-
Sell limit order executed (open)	+
Buy limit order executed (close)	-
Sell limit order executed (close)	+
Buy take-profit (close)	-
Sell take-profit (close)	+
Buy stop-loss (close)	-
Sell stop-loss (close)	+
Buy margin call (close)	not used
Sell margin call (close)	not used
Change order	not used
Change stop-loss or take-profit	not used
Cancel order by hand	not used
Cancel order: insufficient funds	not used
Cancel order: bound violation	not used
Order expired	not used

Table 2: Standard Order Flow Signs. Col. 1 states the record entries, and col. 2 contains the signs for the construction of the standard net order flow measure.

In this standard order flow measure buy and sell orders are treated symmetrically (opposite signs), and following the same logic, we compute, in addition, order flow imbalances for every order category. For example, we can 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 a specific sampling pe-

riod. Altogether, we obtain eight category specific order flow measures 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: Description of the category order flow. Col. 1 states the number of the category and col. 2 gives the category description.

The category specific order flow measures allow insight into several aspects of a trader’s preference structure. In particular, we are able to exploit the information of whether trades are executed to open or close a position, which allows us to analyze the existence of a monitoring effect. We claim that there is a monitoring effect in the sense that traders react faster when they fear losing something (i.e., when they already hold a position) than when they are only planning to take a position. Within a fully rational setup such fear should not exist or at least, there should be no difference between entering the market at the wrong price (thereby implying subsequent losses) and leaving the market at the wrong price (thereby realizing losses). Differences in both types of fear are related to the existence of endowment effects (Thaler (1980)), which ultimately can be seen as a manifestation of loss aversion (Kahneman & Tversky (1979)) and reference based decision making. On the contrary continuous monitoring of the market is costly, for which the investor expects to be compensated by realizing a higher profit through picking a potentially better execution opportunity. The existence of a costly monitoring effect might partially explain the existence of endowment effects within a rational trading strategy.

Provided that such a monitoring effect exists, we should observe that the order flow in the (**close**) categories is easier to predict (based on the information contained in the price process) than order flow in the (**open**) categories. On the one hand, one could argue that this effect should be even more pronounced for market orders than for limit orders since market orders are usually submitted by active and impatient investors who trade for liquidity reasons and watch the market more closely, whereas

limit orders are thought to be submitted by passive traders who might not monitor the market continuously (c.f. Glosten (1994) and Seppi (1997)), even if they have an open position. On the other hand, one could argue that limit order traders might also be very keen to monitor the price process, when their limit orders are outstanding. After placing their limit orders, they may monitor the market and so hark back to stop-loss or take-profit orders, either by placing such orders or by changing their outstanding stop-loss or take-profit orders.

With the category specific order flow measures, we can furthermore, investigate whether we observe self-reinforcing price movements as reported by Osler (2005) on OANDA FXTrade, in the sense that executed stop-loss orders contribute to self-reinforcing price movements whereas executed take-profit orders impede them. Provided that there exist a self-reinforcing price movements, then

- i) based on their own histories, order flow in the stop-loss order category should lend itself more readily to prediction than order flow in the take-profit order category,
- ii) if stop-loss orders induce self-reinforcing price movements and take-profit orders do not, then (in addition to their own histories) information on the price process itself should be more valuable, for predicting order flow of take-profit orders than for predicting order flow of stop-loss orders.

The idea behind this is that there are local downward or upward trends in the price process. Those trends are accelerated by the execution of stop-loss orders, which generate positive feedback trading (De Long et al. (1990)), and are decelerated by the execution of take-profit orders, which generate negative feedback trading or resistance barriers. For 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 neither execution of further stop-loss nor take-profit orders, which yields a decelerated downward movement or even an upward moving price process (bounce back).

4 Empirical Findings

4.1 Description of the Dataset

The dataset used in our analysis is constructed from the trading activity record of OANDA FXTrade from 1st of October 2003 to 14th of May 2004 (227 days). This record contains the complete trading activities for 30 currency pairs on a second by second basis and allows us to distinguish the transaction types listed in Table 1. In addition, depending on the order type, we receive information on transaction prices (market orders, limit orders executed, stop-loss, take profit, margin call), bid and ask quotes (limit orders pending), additional transaction units, and the limits of stop-loss and take-profit orders.

In our analysis we focus on the most actively traded currency pair *EUR/USD*, which accounts for nearly 39 % of all records with an average inter record duration of 8.5 seconds. Table 4 contains the descriptive statistics for the dataset and the transaction volumes for each order category. All figures are daily averages computed over the whole dataset containing 227 days. The average number of different traders per day amounts to 744 for the *EUR/USD* currency pair.

One observation, which is striking when considering the descriptive statistics is that the traders on OANDA FXTrade submit more market orders than limit orders. This characteristic might well be explained by the fact that most traders on this platform are small retail investors who are more impatient and more willing to submit market orders which are executed immediately than limit orders which can be pending for up to one month.³ Another interesting observation in Table 4 is that on average more than 22% of all actions are changes of stop-loss or take-profit limit orders. This figure provides evidence from a descriptive point of view that not only market order traders but also limit order traders do monitor the market closely, when they have an outstanding limit order. This is also in support of the existence of the monitoring effect.

From a series of quotes from the interbank market, we construct the corresponding mid-quotes series for 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). The quotes from the interbank market are provided by Olsen Financial Technologies and represent tradeable quotes stemming from different electronic brokerage systems including Reuters Deal-

³After one month the order expires on OANDA FXTrade.

Transaction Record	%	Obs	Trading Volume in <i>EUR</i> per Day								
			Total	Mean	Min	5% Qtl	25% Qtl	50% Qtl	75% Qtl	95% Qtl	Max
Buy market (open)	13.10	1322	37930860	25854	82	113	515	2065	9240	85854	2220414
Sell market (open)	10.61	1072	30816226	27218	44	89	592	2138	9861	96214	1759412
Buy market (close)	8.27	835	25074760	27468	163	201	672	2326	9553	95940	1630034
Sell market (close)	10.27	1037	31839764	29534	29	66	564	2164	10063	97248	1930846
Limit order: Buy	5.41	546	14041270	28876	24	63	549	2053	9469	95436	1934417
Limit order: Sell	4.76	482	11080825	34283	237	267	515	1662	7509	117914	1511133
Buy limit order executed (open)	3.22	325	5416146	17484	41	79	422	1410	6267	67127	735479
Sell limit order executed (open)	2.92	295	3231307	10554	58	84	242	824	3652	34607	584303
Buy limit order executed (close)	0.46	46	1382690	32718	4800	4824	5313	7020	17994	80426	506182
Sell limit order executed (close)	0.46	46	1470630	32287	407	436	927	3440	16816	93447	452512
Buy take-profit (close)	3.14	317	2918779	9779	144	170	310	704	2724	30314	583296
Sell take-profit (close)	3.49	352	4404025	12857	61	75	256	796	3960	43028	820876
Buy stop-loss (close)	2.18	220	4488496	16433	126	175	667	2535	9837	70968	513989
Sell stop-loss (close)	2.55	258	5309807	16667	23	59	503	2255	9424	66743	650061
Buy margin call (close)	0.12	12	166375	7263	1006	1010	1185	1817	3718	14211	71133
Sell margin call (close)	0.17	17	275282	6381	1369	1372	1440	2351	4409	17266	77231
Change order	3.01	305	13898910	49771	105	203	1295	4888	18181	162927	1622712
Change stop-loss or take-profit	22.36	2260	60965013	26748	10	79	867	3694	14163	95983	1703030
Cancel by hand	2.41	243	10043949	42295	211	272	1031	4186	16003	148571	1662224
Cancel: insufficient funds	0.28	28	2439586	67905	4938	4953	5431	7354	66280	186280	622650
Cancel bound violation	0.20	20	195118	14803	571	571	627	2650	6860	29909	98308
Order expired	0.65	66	1063061	19942	44	54	443	1682	7204	68648	355982

Table 4: Descriptive statistics of the OANDA FXTrade trading activity dataset for the *EUR/USD* currency pair. All numbers are daily averages over 227 days and all transaction volume statistics are denominated in *EUR*.

ing 3000 and EBS. These 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 which also includes tradeable quotes from Reuters Dealing 3000 and EBS. In addition to the price series, we construct the order flow measures (standard and category specific) defined in the previous section for the corresponding 12 frequencies.

Figure 1 depicts the empirical bivariate autocorrelation functions for lags up to 20 periods between price changes and standard order flow for a 1 minute frequency. The analysis of the bivariate autocorrelation functions sheds light on the dynamic interaction of order flow and price change.

- In the upper left panel, we observe the autocorrelation function of the order flow measure itself. We see a very clear slowly declining pattern of the autocorrelation function, showing that order flow itself is persistent.
- The lower left panel depicts the cross-correlation function of lagged order flow with price changes. We observe that only the first order cross-correlation coefficient is significantly positive, which shows that future price changes are driven by current order flow. This supports the literature concentrated on predicting/explaining price changes with order flow (e.g., Evans & Lyons (2002a,b, 2005, 2006, Daniélsson et al. (2002))).
- The upper right panel depicts the cross-correlation function of lagged price changes with order flow.⁴ We observe significant cross-correlation coefficients, which show that future order flow is driven by current price changes. This observation supports heuristically the idea that investors update their beliefs and place their orders based on the past development of the price process, and hence is another indication for the existence of a learning effect (c.f. Bikhchandani et al. (1992) and Avery & Zemsky (1998)). From this observation, this effect seems to be a short run effect, since the cross-correlation coefficients between future order flow and current price changes are significant for only up to 5 lags.
- In the lower right panel, we observe the autocorrelation function of the price changes themselves. The price changes are positively first order auto-correlated.

⁴For both cross-correlation functions, we plot lag 0 through 19. The value at lag 0 is the same in both cross-panels and represents the contemporaneous correlation between order flow and price changes.

Thus, we observe a kind of short term positive feedback trading pattern for the price process itself, and we do not observe a traditional bid-ask bounce effect, since we consider mid-quotes on a 1 minute frequency.

Standard Order Flow vs. Price Changes




Figure 1: Empirical bivariate autocorrelation functions of price changes and standard order flow for an aggregation level of 1 minute. The upper left panel depicts the autocorrelation function (lag: 1–20) of the order flow measure and 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.

4.2 Estimation Framework

Although the analysis above provides some insight into the dynamic relationship between order flow and price change, giving for a first impression of the validity of the hypotheses stated at the outset, we now investigate them in detail with the help of in-sample and out-of-sample forecasting analyses.

Based on the results of the descriptive analysis⁵, we now consider AR(p) benchmark models in which only the history of the order flow measures themselves serve to explain and to predict future order flows. These predictions are then compared, using the modified Diebold-Mariano (mDM) test of Harvey, Leybourne & Newbold (1997), to the predictions based on including the history of price changes. This enables us to identify whether the inclusion of **additional** information contained in past prices improves order flow forecasts significantly.

The AR(p) benchmark specification is given by

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

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

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

where Δy_t denotes the interbank price change process, and $B_q^y(L)$ the lag-polynomial of order q .

The forecasting study is executed in the following way: Altogether we consider a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. We divide these 32 weeks into 8 periods of 4 weeks each, where the first 3 weeks are considered as the in-sample estimation period and the last weeks is considered as the out-of-sample forecasting period. Table 5 summarizes the setup of the in- and out-of-sample periods. We choose this forecasting setup with alternating in-sample and out-of-sample periods in order to guarantee robust forecasting results as compared to studies with only one estimation and one forecasting period. Our forecasting setup is particularly conservative in the sense that we estimate the model parameters only once for every in-sample period and do not use a rolling window regression technique with continuous updating of model parameter estimates. Holidays and week-ends are excluded from the sample.

⁵We observed that the order flow process is persistent.

Period	In-Sample	Out-of-Sample
1	Mo. 6. Oct. 2003 – Fr. 24. Oct. 2003	Mo. 27. Oct. 2003 – Fr. 31. Oct. 2003
2	Mo. 3. Nov. 2003 – Fr. 21. Nov. 2003	Mo. 24. Nov. 2003 – Fr. 28. Nov. 2003
3	Mo. 1. Dec. 2003 – Fr. 19. Dec. 2003	Mo. 22. Dec. 2003 – Fr. 26. Dec. 2003
4	Mo. 29. Dec. 2003 – Fr. 26. Jan. 2004	Mo. 19. Jan. 2004 – Fr. 23. Jan. 2004
5	Mo. 26. Jan. 2004 – Fr. 13. Feb. 2004	Mo. 16. Feb. 2004 – Fr. 20. Feb. 2004
6	Mo. 23. Feb. 2004 – Fr. 12. Mar. 2004	Mo. 15. Mar. 2004 – Fr. 19. Mar. 2004
7	Mo. 22. Mar. 2004 – Fr. 9. Apr. 2004	Mo. 12. Apr. 2004 – Fr. 19. Apr. 2004
8	Mo. 19. Apr. 2004 – Fr. 7. May 2004	Mo. 10. May 2004 – Fr. 14. May 2004

Table 5: In-sample and out-of-sample periods of the forecasting study.

The results we are going to present are very robust to the choice of the in-sample estimation and out-of-sample prediction periods. Moreover, the same conclusions can be drawn if the hypotheses are evaluated separately with respect to the eight individual forecasting setups presented in Table 5. The results are furthermore robust to the inclusion or exclusion of the overnight periods. In our analysis we have chosen to include the overnight periods.

4.3 In-Sample Estimation Results

Let us first consider the in-sample estimation results. In order to be able to compare the results, we choose to present them for a lag-polynomial of order one, implying that equation (IP- k) simplifies to

$$x_t^k = c + \rho x_{t-1}^k + \phi \Delta y_{t-1} + \varepsilon_t, \quad \text{for } k = 1, \dots, 8,$$

and we are going to interpret the magnitude of the estimated coefficients.

Tables 6 to 8 present the estimated coefficients of the in-sample estimations of the category specific order flow measures for each period on the 1, 2, and 5 minutes sampling frequencies, respectively. We observe that the estimated coefficients on lagged price changes and lagged order flows have the expected signs, for each order flow category and each period. Let us consider, for example, the stop-loss order flow category. We see that the estimated coefficients for the price change are positive. Assume for the moment that the price change is positive, and remember that we defined the dependent variable in our regression as the difference between the number of buy and sell stop-loss orders. A positive price change leads in this case to an increase of the dependent variable, meaning that we have an increase in the execution of buy stop-loss orders. On the opposite, if the price change is negative,

the dependent variable decreases inducing a larger execution of sell stop-loss orders. The reverse is true if we consider the take-profit order flow category.

Moreover, we see that the impact of price changes on limit order flows (executed open and close) is always lower than the impact of price changes on the market order flows (open and close). Specifically, the estimated coefficients of the price changes on **market (open)** order flow is roughly three to four times larger than the estimated coefficients for the **limit executed (open)** order flow category. In addition, if we only have a look at the market order flow categories, we note that the price impact is always about 3-5 times larger for the **market (close)** order flow category than for the **market (open)** category. This is another observation in support of the existence of a monitoring effect, in the sense that traders monitor the market more closely when they already hold an open position. We also observe that the impact of the price process is larger for take-profit order flow than for stop-loss order flow for all three sampling frequencies. This shows that prices are more valuable for the explanation of take-profit orders than stop-loss orders, which supports our hypothesis that take-profit orders seem to impede price movements and stop-loss orders seem to self-reinforce them.

Altogether, these effects are relatively stable over all periods and over all order flow categories, and confirm within an in-sample analysis our hypotheses. Our analysis, furthermore, sheds light on the discussion whether prices already reflect all relevant information and can be considered as sufficient statistics. Our results show that they are not. If prices were a sufficient statistic, one should observe a significant coefficient only for the lagged price changes.⁶ Below we are going to discuss the results of our out-of-sample forecasting study which provides stronger evidence in support of our hypotheses.

4.4 Forecasting Study

To validate our claims, we compare the Root-Mean-Squared-Predictions-Errors (RM-SPE) of the AR(p) benchmark models with those of the corresponding forecasting models for the in-sample and out-of-sample forecasting studies.

⁶See Caplin & Leahy (1996).

Periods/ Coefs	Standard Order Flow	Limit Order	Limit Order executed open	Limit Order executed close	Market Order Open	Market Order Close	Stop-Loss Order	Take-Profit Order	Margin Call Order	
1	c	-0.0669	0.0443	0.0183	0.0094	0.0887	-0.0128	0.0054	0.0018	
	ρ	0.0799	0.2431	0.1432	0.1014	0.1553	0.1615	0.2471	0.0396	
	ϕ	-0.5212	-0.2208	-0.2839	-0.0785	-0.0186	-1.4091	0.6460	-0.8058	0.0135
	R^2	0.0069	0.0625	0.0307	0.0011	0.0240	0.0559	0.0478	0.0894	0.0021
2	c	0.0960	0.0241	0.0141	0.0027	0.0789	-0.0602	0.0169	-0.0853	0.0038
	ρ	0.1035	0.1721	0.1002	0.0004	0.1968	0.1188	0.1383	0.2367	0.0761
	ϕ	-1.8104	-0.1135	-0.3830	-0.0573	-0.6455	-2.3909	0.8272	-1.0846	0.0952
	R^2	0.0178	0.0305	0.0245	0.0005	0.0430	0.0455	0.0409	0.0742	0.0103
3	c	0.1166	0.0270	0.0030	-0.0074	0.2324	-0.1876	-0.0060	-0.0245	0.0264
	ρ	0.0901	0.3776	0.1327	0.0693	0.1886	0.1727	0.1672	0.1846	0.0312
	ϕ	-1.4268	-0.0894	-0.5164	-0.1553	-0.8601	-2.3096	0.6273	-1.2807	0.3071
	R^2	0.0122	0.1428	0.0326	0.0094	0.0432	0.0708	0.0453	0.0646	0.0042
4	c	0.2078	0.0833	0.0678	0.0036	0.3556	-0.2200	-0.0807	-0.0344	-0.0020
	ρ	0.1217	0.4167	0.0781	-0.0018	0.2188	0.2375	0.2028	0.2827	0.1071
	ϕ	-1.4371	-0.2959	-0.3396	-0.1561	-0.8523	-2.0128	0.7321	-1.2813	0.3159
	R^2	0.0190	0.2681	0.0120	0.0015	0.0534	0.1066	0.0590	0.1252	0.0199
5	c	0.1828	0.0043	-0.0049	0.0032	0.2657	-0.1936	-0.0493	-0.0306	0.0020
	ρ	0.1323	0.3899	0.1694	0.0189	0.2763	0.2598	0.1665	0.2667	0.1132
	ϕ	-1.3272	-0.1742	-0.0318	-0.0801	-0.1310	-2.7899	0.7590	-1.2043	0.2026
	R^2	0.0189	0.1527	0.0287	0.0016	0.0552	0.1249	0.0471	0.0977	0.0397
6	c	0.1358	0.0929	0.0939	-0.0020	0.2663	-0.2048	-0.0736	-0.0917	-0.0255
	ρ	0.1619	0.4201	0.1833	0.0570	0.2821	0.1630	0.1169	0.2260	0.5365
	ϕ	-1.6794	-0.1937	0.0913	-0.0999	-0.7343	-3.8426	1.5667	-2.5881	0.2685
	R^2	0.0275	0.1775	0.0339	0.0046	0.0739	0.1016	0.0520	0.1140	0.2946
7	c	0.2039	0.0276	0.0425	-0.0096	0.1955	-0.1185	-0.0356	-0.1154	-0.0115
	ρ	0.1564	0.5497	0.2329	0.0172	0.2288	0.1764	0.1230	0.2293	0.1397
	ϕ	-1.2569	-0.2079	0.2041	-0.1799	-0.9590	-3.0670	1.6381	-3.1767	0.3999
	R^2	0.0013	0.3032	0.0553	0.0005	0.0552	0.0883	0.0489	0.1093	0.0511
8	c	0.0105	-0.0447	-0.1118	0.0117	0.0367	-0.0459	-0.0023	0.0864	-0.0012
	ρ	0.1029	0.5178	0.2210	0.0776	0.2505	0.1865	0.1248	0.2419	0.1789
	ϕ	-0.9889	-0.1914	2.3429	-0.3399	-0.7982	-3.5646	0.8586	-4.4995	0.2493
	R^2	0.0107	0.2681	0.0820	0.0126	0.0650	0.1066	0.0304	0.1162	0.0505

Table 6: Estimated coefficients of the **in-sample** estimations of the standard order flow and the category specific order flow measures on the **1 minute** sampling frequency for each forecasting period. c denotes the estimated constant, ρ the estimated coefficient of the lagged order flow, and ϕ the estimates coefficient of the price change. The estimated coefficients in bold are significant at the 5% significance level.

Periods/ Coefs	Standard Order Flow	Limit Order	Limit Order executed open	Limit Order executed close	Market Order Open	Market Order Close	Stop-Loss Order	Take-Profit Order	Margin Call Order
1	c	-0.1337	0.0880	0.0378	0.0185	0.1720	-0.0272	0.0123	0.0034
	ρ	0.0646	0.2501	0.1211	0.1251	0.1854	0.2059	0.1104	0.1943
	ϕ	-1.6515	-0.3798	-0.3681	-0.1185	-0.5498	-2.5280	0.7220	-1.1802
	R^2	0.0140	0.0713	0.0296	0.0180	0.0348	0.1138	0.0339	0.0787
2	c	0.1850	0.0470	0.0277	0.0054	0.1465	-0.1127	0.0365	-0.1753
	ρ	0.1576	0.1973	0.1278	-0.0003	0.2660	0.1531	0.0577	0.2170
	ϕ	-3.0206	-0.1795	-0.4880	-0.0547	-1.2523	-3.2364	1.0380	-1.0574
	R^2	0.0447	0.0411	0.0413	0.0005	0.0862	0.0820	0.0248	0.0644
3	c	0.2450	0.0094	0.0064	-0.0152	0.4628	-0.3610	-0.0132	-0.0486
	ρ	0.0556	0.3388	0.1246	0.0133	0.1951	0.1970	0.1147	0.1805
	ϕ	-2.4092	-0.2329	-0.6470	-0.3357	-1.4253	-3.2542	0.7633	-1.4940
	R^2	0.0162	0.1572	0.0367	0.0159	0.0584	0.1158	0.0321	0.0722
4	c	0.3913	0.1679	0.1259	0.0070	0.6386	-0.3864	-0.1861	-0.0766
	ρ	0.1724	0.4121	0.1449	0.0083	0.2982	0.3314	0.0809	0.2075
	ϕ	-1.6346	-0.2597	-0.2126	-0.4436	-1.1155	-2.6758	1.1846	-2.1920
	R^2	0.0353	0.1725	0.0254	0.0145	0.0985	0.1598	0.0283	0.1243
5	c	0.2715	0.0094	-0.0104	0.0061	0.4859	-0.3752	-0.1029	-0.0669
	ρ	0.1498	0.3388	0.1617	0.0635	0.3396	0.2808	0.1288	0.1971
	ϕ	-2.9081	-0.2329	0.1108	-0.0390	-0.7715	-3.5961	0.5452	-1.3821
	R^2	0.0265	0.1157	0.0262	0.0045	0.1072	0.1723	0.0283	0.0650
6	c	0.2715	0.1665	0.1953	-0.0044	0.5170	-0.3938	-0.1514	-0.1994
	ρ	0.1498	0.4796	0.1513	0.0265	0.3010	0.1978	0.1049	0.1453
	ϕ	-2.9081	-0.2666	0.0914	-0.1608	-1.1733	-3.9745	0.7991	-1.7344
	R^2	0.0323	0.2320	0.0232	0.0047	0.0810	0.1392	0.0272	0.0509
7	c	0.4128	0.0455	0.0878	-0.0196	0.3650	-0.2386	-0.0747	-0.2528
	ρ	0.1477	0.6263	0.2088	0.0002	0.2806	0.1711	0.0821	0.1592
	ϕ	-0.8253	-0.4406	0.3284	-0.1714	-0.7977	-3.1338	1.6239	-4.0912
	R^2	0.0353	0.3954	0.0457	0.0040	0.0796	0.0978	0.0378	0.0965
8	c	0.0191	-0.0878	-0.2259	0.0227	0.0729	-0.0909	-0.0053	0.1893
	ρ	0.1217	0.5283	0.2156	0.1021	0.2522	0.1993	-0.0022	0.1740
	ϕ	-1.6207	-0.5594	1.8946	-0.3719	-0.9890	-3.6624	1.0755	-4.2999
	R^2	0.0162	0.2799	0.0744	0.0211	0.0681	0.1289	0.0098	0.0802

Table 7: Estimated coefficients of the **in-sample** estimations of the standard order flow and the category specific order flow measures on the **2 minute** sampling frequency for each forecasting period. c denotes the estimated constant, ρ the estimated coefficient of the lagged order flow, and ϕ the estimates coefficient of the price change. The estimated coefficients in bold are significant at the 5% significance level.

Periods/ Coefs	Standard Order Flow	Limit Order	Limit Order executed open	Limit Order executed close	Market Order Open	Market Order Close	Stop-Loss Order	Take-Profit Order	Margin Call Order	
1	c	-0.3258	0.2263	0.0872	0.0483	0.3848	-0.5615	-0.0749	0.0339	0.0076
	ρ	0.0783	0.2284	0.1802	0.0921	0.2729	0.2638	0.0399	0.1650	0.1879
	ϕ	-2.0463	-0.2974	-0.2012	-0.1774	-0.5175	-3.1197	1.0414	-1.6232	0.0935
	R^2	0.0216	0.0584	0.0408	0.0128	0.0732	0.1782	0.0262	0.0808	0.0560
2	c	0.4630	0.1252	0.0673	0.0135	0.3217	-0.2910	0.0895	-0.4655	0.0166
	ρ	0.1509	0.1473	0.1399	-0.0037	0.3555	0.1216	0.1031	0.1603	0.1766
	ϕ	-2.6838	-0.2317	-0.3450	-0.0347	-1.1251	-3.5067	0.6781	-1.6963	0.1552
	R^2	0.0381	0.0246	0.0359	0.0002	0.1393	0.0771	0.0261	0.0513	0.0443
3	c	0.6305	0.1391	0.0126	-0.0392	1.1215	-0.8945	0.0895	-0.1312	0.1307
	ρ	0.0435	0.3761	0.2135	0.0027	0.2207	0.2020	0.1031	0.1349	0.0464
	ϕ	-3.6552	-0.4762	-0.3444	-0.2567	-1.4583	-3.5703	0.6781	-1.2263	0.2030
	R^2	0.0322	0.1436	0.0583	0.0084	0.0675	0.1205	0.0184	0.0404	0.0037
4	c	0.9889	0.4863	0.3176	0.0173	1.4171	-1.0935	-0.4569	-0.1875	-0.0101
	ρ	0.1637	0.3197	0.1373	0.0412	0.3776	0.2436	0.0978	0.2213	0.0850
	ϕ	-2.0823	-0.2437	-0.1944	-0.2745	-0.6348	-3.3677	0.9881	-1.7302	0.4784
	R^2	0.0363	0.1042	0.0225	0.0091	0.1454	0.1184	0.0269	0.1108	0.0263
5	c	0.9039	0.0217	-0.0236	0.0155	1.0724	-0.9821	-0.2799	-0.1722	0.0090
	ρ	0.1455	0.3946	0.2268	0.0686	0.4170	0.2463	0.0585	0.1603	0.2223
	ϕ	-1.7444	-0.1504	-0.0539	-0.1535	-0.6012	-3.8549	1.0151	-2.2359	0.1913
	R^2	0.1560	0.0713	0.0514	0.0048	0.1659	0.1451	0.0197	0.0622	0.0797
6	c	0.6953	0.4200	0.4439	-0.0103	1.1955	-0.9952	-0.3777	-0.4956	-0.2269
	ρ	0.1501	0.4782	0.2281	0.0531	0.3593	0.1930	0.1089	0.1562	0.1529
	ϕ	-0.5469	0.1770	0.3988	-0.0666	0.6281	-4.6819	0.6560	-2.3292	1.4470
	R^2	0.0223	0.2281	0.0556	0.0039	0.1389	0.1631	0.0247	0.0664	0.0492
7	c	1.0280	0.1204	0.2157	-0.0484	0.8050	-0.5818	-0.1876	-0.6327	-0.0584
	ρ	0.1491	0.5907	0.2184	0.0178	0.3661	0.1993	0.0794	0.1583	0.1287
	ϕ	-1.6289	-0.4199	0.4697	-0.1730	-0.5313	-4.2164	1.9232	-4.6241	0.3723
	R^2	0.0256	0.3509	0.0514	0.0048	0.1344	0.1448	0.0438	0.0991	0.0457
8	c	0.0496	-0.2703	-0.6011	0.0616	0.2001	-0.2357	-0.0143	0.5409	-0.0073
	ρ	0.0426	0.4253	0.1727	0.0398	0.1754	0.1676	-0.0980	0.0619	0.1234
	ϕ	-2.1192	-1.6292	0.3279	-0.1811	-1.2644	-3.2275	1.2247	-3.7611	0.0964
	R^2	0.0074	0.1892	0.0325	0.0042	0.0387	0.0946	0.0074	0.0254	0.0199

Table 8: Estimated coefficients of the **in-sample** estimations of the standard order flow and the category specific order flow measures on the **5 minute** sampling frequency for each forecasting period. c denotes the estimated constant, ρ the estimated coefficient of the lagged order flow, and ϕ the estimates coefficient of the price change. The estimated coefficients in bold are significant at the 5% significance level.

The results over all 8 periods of the in-sample and out-of-sample studies using the standard order flow measure are presented in Table 9. The in-sample results based on the category based order flow are given in Tables 10 and 11, and the corresponding out-of-sample results can be found in Tables 12 and 13. The first cell entry in Table 9 as well as in all other following Tables is the RMSPE of the associated forecasting model. The second cell entries in parenthesis are the p-values 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 $AR(p)$ benchmark model. P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the benchmark model.

We clearly observe that the information contained in the history of the price process in addition to the information contained in the order flow measures themselves is helpful in predicting the aggregated standard order flow as well as the eight category specific order flow measures. This statement is based on the observations that

- i) for the in-sample prediction of the standard order flow measure (Table 9 col. 2 and 3) we observe that on all frequencies the forecasting models containing the additional information on the historical interbank price changes are able to beat the benchmark specification (AR) in terms of smaller RMSPEs (bold cell entries). Considering the p-values of the mDM-tests, we see that these models outperform the AR benchmark model, with 5 exceptions, always at the 10% significance level.
- ii) for the out-of-sample prediction of the standard order flow measure (Table 9 col. 4 and 5) the RMSPEs are for the 3 forecasting horizons up to 5 minutes smaller than those of the $AR(p)$ benchmark models, when additional information on the interbank price change process is incorporated in the forecasting models. We see that the RMSPEs for 1 and 2 minutes forecasting horizons are significantly smaller using a 1% significance level in the mDM test.
- iii) for the in-sample prediction of the category specific order flow measures (Tables 10 and 11) we find that the RMSPEs of the $AR(p)$ benchmark models are, except for 10 cases, always higher than those of the forecasting models with the historical prices, and 67 are significant at the 10% significance level.
- iv) for the out-of-sample prediction of the category specific order flow measures (Tables 12 and 13) we observe that over all eight categories 59 RMSPEs are

smaller than those of the $AR(p)$ benchmark models; 19 of them are significantly smaller at the 1% level, 25 of them are significantly smaller at the 5% level and even 31 at the 10% level.

These results enable an interesting interpretation: The aggregated standard order flow as well as the category specific order flow measures can generally be easier predicted with the help of historical price information being, in particular in the absence of macroeconomic news and private (customer order flow) information, the only source of information available to the traders. This might indicate that traders update their beliefs and place orders based on their interpretation of recent historical price movements, and therefore might rely on either technical analysis as pointed out by Taylor & Allen (1992),⁷ or simply attempt to intuitively extrapolate recent price movements information. Especially on OANDA FXTrade, such information is more valuable than on the interbank market, since most of the traders do not have own customer order flow, which can serve as private information. Our observation also shows that it might take up to several hours to process new information and gain market efficiency.

Moreover, regardless of whether traders hold an open position or not, we find that the price process contributes more to the prediction of market order flow than to the prediction of limit order flow, in both in-sample and out-of-sample periods. Following Foucault (1999), who shows that in equilibrium impatient traders tend to submit market orders and patient traders limit orders, and taking the descriptive statistics into account, a likely interpretation is that traders on OANDA-FXTrade are very impatient. The evidence supports the view that small, not sophisticated traders, who place their orders based on their interpretation of recent price movements, and thus react quickly to recent price changes, prefer to submit market orders which are executed immediately, instead of limit orders. This is reflected at the end of the day by the fact that market orders are more predictable than limit orders. This finding is also consistent with the results of Rosu (2009).

Our postulate of the existence of a monitoring effect is clearly supported. Indeed, we see, in Tables 12 and 13, that the price process contributes more to the prediction

⁷The survey study of Taylor & Allen (1992) shows that at least 90% of the London based dealers rely, in addition to private and fundamental information, on information from technical analyses to design their trading strategies.

of market (**close**) order flow than to the prediction of market (**open**) order flow.⁸ Thus, we find evidence that traders employ different strategies when they hold an open position. In detail, we observe that only the RMSPEs for the 1 minute and the 2 minute frequencies of the price changes for the market order (open) forecasting models are significantly smaller (1% level) than those of the corresponding AR(p)-benchmark models, but that the first 5 RMSPEs corresponding to frequencies up to 15 minutes are significantly smaller than those of the benchmark models for the market order (close) order flow category at the 1% level. From the remaining (lower) frequencies only one is significant on the 10% level of the mDM-test for price changes for the market order (close) forecasting models. A similar observation, but not as pronounced, can be made for limit order (close) order flow and limit order executed (open) order flow as well. The reason why this effect is not as clear as for market orders is that limit orders are posted to the system before market orders, and their execution is later than simply implied by the price process. Market orders, however, reflect changes in price preferences directly since they are executed immediately. Our results complete the work of Feng & Seasholes (2005), who investigate the question of whether investors sophistication and trading experience attenuate behavioral biases in financial markets. They analyse if investors actively monitor stocks they have sold, and their results provide indirect evidence that they do. Our analysis is more focused on a small set of securities and not a large universe of stocks, but we also provide evidence that the investors monitor the price process closer when they already hold an open position.

Let us now consider the hypothesis that executed stop-loss orders contribute whereas executed take-profit orders impede, self-reinforcing price movements. Table 13 shows that at all forecasting frequencies, the RMSPEs of the benchmark models for the stop-loss order flow category are smaller than those of benchmark models for the take-profit order flow category. This observation supports the existence of price cascades since stop-loss order flow, based on its own historical order flow, is better predictable (in terms of smaller RMSPEs) than take-profit order flow. This is foreseen if stop-loss orders contribute to self-reinforcing price movements causing a sequence of further stop-loss order executions.

Comparing the RMSPE pattern of the stop-loss and take-profit order flow categories forecasting models containing additional information on the history of the price change process, we observe that there is essentially no difference in the value

⁸The same conclusion can be drawn from the in-sample forecasting studies.

of the history of the price process in predicting take-profit or stop-loss order flow. This observation thus provides little additional evidence for the validity of our statement, since we expected that the information on the direction of the price change process should already be included in the historical stop-loss order flow. It therefore should be of less importance in predicting future stop-loss order flow in contrast to the case when take profit order flow is considered. Taken also the results of the in-sample analysis into account, altogether, we find weak evidence that stop-loss orders self-reinforce price movements, but we cannot draw a clear cut conclusion here.

Freq	In-Sample		Out-of-Sample	
	BM	hist. prices	BM	hist. prices
1 min (AR)	. 4.3485	4.3275 (0.0000)	. 4.2978	4.2726 (0.0000)
2 min (AR)	. 6.5269	6.4809 (0.0000)	. 6.5110	6.4540 (0.0000)
5 min (AR)	. 11.1317	11.0924 (0.0005)	. 11.2270	11.2066 (0.1451)
10 min (AR)	. 16.6981	16.6804 (0.1550)	. 16.8088	16.9430 (0.9996)
15 min (AR)	. 20.9063	20.8962 (0.3847)	. 21.5097	21.6271 (0.9888)
20 min (AR)	. 25.0626	25.0171 (0.0958)	. 25.8728	26.0218 (0.9914)
25 min (AR)	. 28.7239	28.7046 (0.3424)	. 29.8577	29.9160 (0.7327)
30 min (AR)	. 32.0495	31.8593 (0.0141)	. 33.4052	33.5240 (0.8807)
45 min (AR)	. 41.4252	41.2203 (0.1617)	. 44.6980	44.9230 (0.7714)
1 hr (AR)	. 48.3953	48.0999 (0.1627)	. 52.1271	52.8020 (0.9685)
2 hr (AR)	. 72.2455	71.5105 (0.0333)	. 85.9340	86.3036 (0.6968)
4 hr (AR)	. 116.3498	115.7275 (0.0704)	. 143.6758	144.5649 (0.7909)

Table 9: Results for the **standard** order flow measure **in-sample** (col. 2 and 3) **out-of-sample** (col. 4 and 5) predictions on different sampling frequencies (Freq). The forecasting study is conducted over a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. These 32 weeks are divided into 8 periods of 4 weeks each, where the first 3 weeks are always considered as the in-sample estimation periods and the last weeks are always considered as the out-of-sample forecasting periods. 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 and third cell entries in parenthesis are the p-value from the modified Diebold-Mariano (mDM) test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the AR(p) benchmark model (BM). P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the AR(p) benchmark model.

Freq	Limit Orders		Limit Orders Executed Open		Limit Orders Executed Close		Market Orders Open	
	BM	hist. prices	BM	hist. prices	BM	hist. prices	BM	hist. prices
1 min (AR)	. 1.4454	1.4445 (0.0000)	. 1.6270	1.6175 (0.0976)	. 0.6330	0.6315 (0.0010)	. 1.9857	1.9770 (0.0000)
2 min (AR)	. 2.4518	2.4484 (0.0002)	. 2.5267	2.5158 (0.0726)	. 0.9106	0.9073 (0.0019)	. 3.0876	3.0700 (0.0000)
5 min (AR)	. 5.1312	5.1219 (0.0599)	. 4.5255	4.5236 (0.2658)	. 1.4973	1.4948 (0.0011)	. 5.5784	5.5518 (0.0000)
10 min (AR)	. 9.0484	9.0444 (0.3588)	. 7.1240	7.1464 (0.9509)	. 2.1762	2.1722 (0.0010)	. 9.1244	9.0985 (0.0019)
15 min (AR)	. 12.2393	12.2766 (0.8578)	. 9.1686	9.1475 (0.1799)	. 2.7164	2.7143 (0.3744)	. 12.1184	12.1397 (0.7651)
20 min (AR)	. 15.1778	15.1480 (0.1501)	. 10.7343	10.7400 (0.6278)	. 3.1611	3.1526 (0.0032)	. 14.5709	14.4889 (0.0335)
25 min (AR)	. 17.9797	17.9766 (0.4771)	. 12.1639	12.2353 (0.8521)	. 3.5789	3.5638 (0.0105)	. 17.4819	17.3621 (0.0088)
30 min (AR)	. 19.8530	19.8267 (0.2631)	. 14.5193	14.5874 (0.8520)	. 3.9716	3.9339 (0.0237)	. 20.1194	19.9743 (0.0061)
45 min (AR)	. 26.7754	26.8938 (0.7577)	. 18.8809	18.8847 (0.5364)	. 4.9625	4.9513 (0.0902)	. 26.7991	26.5359 (0.0112)
1 hr (AR)	. 32.6755	32.6807 (0.5225)	. 22.3827	22.3710 (0.3974)	. 5.8760	5.8693 (0.2387)	. 33.6136	33.3265 (0.0318)
2 hr (AR)	. 52.3839	52.2478 (0.1935)	. 36.2961	36.3494 (0.6072)	. 8.3189	8.2904 (0.1125)	. 57.5232	56.8558 (0.0050)
4 hr (AR)	. 83.9541	83.7328 (0.2998)	. 66.0293	65.9525 (0.3041)	. 11.9968	11.9914 (0.4525)	. 102.0218	101.0879 (0.0485)

Table 10: Results of the **in-sample** predictions of the category specific order flow measures on different sampling frequencies (Freq). The forecasting study is conducted over a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. These 32 weeks are divided into 8 periods of 4 weeks each, where the first 3 weeks are always considered as the in-sample estimation periods and the last weeks are always considered as the out-of-sample forecasting periods. 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 and third cell entries in parenthesis are the p-value from the modified Diebold-Mariano (mDM) test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the AR(p) benchmark model. P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the AR(p) benchmark model.

Freq	Market Orders Close		Stop-Loss Orders Close		Take-Profit Orders Close		Margin Call Orders Close	
	BM	hist. prices	BM	hist. prices	BM	hist. prices	BM	hist. prices
1 min (AR)	. 2.7454	2.6613 (0.0000)	. 1.9344	1.9238 (0.0000)	. 3.2039	3.1693 (0.0000)	. 0.7265	0.7239 (0.0002)
2 min (AR)	. 4.2203	4.0575 (0.0000)	. 3.0333	3.0210 (0.0000)	. 5.2630	5.2056 (0.0000)	. 1.2003	1.1954 (0.0022)
5 min (AR)	. 7.8304	7.5485 (0.0000)	. 5.2662	5.2454 (0.0000)	. 9.8193	9.7153 (0.0000)	. 2.6045	2.5830 (0.0779)
10 min (AR)	. 12.5372	12.1819 (0.0000)	. 7.9106	7.8903 (0.0002)	. 15.3615	15.3124 (0.0667)	. 3.8643	3.8418 (0.0347)
15 min (AR)	. 16.5733	16.2111 (0.0000)	. 9.9829	9.9745 (0.0798)	. 19.8503	19.8464 (0.4749)	. 4.9016	4.8678 (0.0371)
20 min (AR)	. 20.1138	19.6528 (0.0000)	. 11.7201	11.6999 (0.0208)	. 23.0231	22.8464 (0.0001)	. 6.0462	5.9750 (0.0642)
25 min (AR)	. 23.3486	22.9671 (0.0000)	. 13.2341	13.2237 (0.2025)	. 27.0023	26.9249 (0.0567)	. 6.4416	6.3820 (0.0544)
30 min (AR)	. 27.4590	27.1266 (0.0000)	. 15.1372	15.1554 (0.6431)	. 32.6134	32.4378 (0.0015)	. 7.0998	7.0261 (0.0401)
45 min (AR)	. 35.8041	35.3963 (0.0014)	. 19.0306	19.0036 (0.0917)	. 41.8647	41.7748 (0.0203)	. 9.5666	9.4286 (0.0681)
1 hr (AR)	. 44.7921	44.2294 (0.0007)	. 22.0938	21.9956 (0.0693)	. 50.5215	50.3161 (0.0214)	. 10.9656	10.8131 (0.0540)
2 hr (AR)	. 73.2669	72.8933 (0.0425)	. 32.0076	31.6867 (0.0437)	. 83.8739	83.6420 (0.1951)	. 16.1938	16.0723 (0.1540)
4 hr (AR)	. 115.5674	114.7439 (0.0385)	. 47.2952	46.6408 (0.0262)	. 137.1682	136.1951 (0.1650)	. 22.9814	22.8589 (0.1321)

Table 11: Results of the **in-sample** predictions of the category specific order flow measures on different sampling frequencies (Freq). The forecasting study is conducted over a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. These 32 weeks are divided into 8 periods of 4 weeks each, where the first 3 weeks are always considered as the in-sample estimation periods and the last weeks are always considered as the out-of-sample forecasting periods. 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 and third cell entries in parenthesis are the p-value from the modified Diebold-Mariano (mDM) test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the AR(p) benchmark model. P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the AR(p) benchmark model.

Freq	Limit Orders		Limit Orders Executed Open		Limit Orders Executed Close		Market Orders Open	
	BM	hist. prices	BM	hist. prices	BM	hist. prices	BM	hist. prices
1 min (AR)	. 1.5059	1.5050 (0.0031)	. 1.8412	1.8370 (0.1212)	. 0.7415	0.7409 (0.1797)	. 1.8765	1.8708 (0.0000)
2 min (AR)	. 2.5477	2.5435 (0.0002)	. 2.8676	2.8699 (0.6597)	. 1.1982	1.1969 (0.2627)	. 2.8995	2.8844 (0.0000)
5 min (AR)	. 5.5135	5.4916 (0.0118)	. 5.1697	5.1626 (0.0662)	. 1.7648	1.7621 (0.0305)	. 5.2898	5.2926 (0.6315)
10 min (AR)	. 9.8622	9.8526 (0.3319)	. 8.3812	8.4009 (0.9185)	. 2.8620	2.8594 (0.1706)	. 8.7090	8.7002 (0.2641)
15 min (AR)	. 13.5125	13.4816 (0.3837)	. 11.7627	11.7542 (0.3281)	. 3.5404	3.5243 (0.1543)	. 11.6636	11.7545 (0.9845)
20 min (AR)	. 16.5926	16.4909 (0.0745)	. 12.8240	12.8136 (0.4148)	. 4.1443	4.1438 (0.4413)	. 14.1298	14.1836 (0.7615)
25 min (AR)	. 19.9732	19.8155 (0.1639)	. 16.1600	16.2733 (0.9573)	. 4.6501	4.6562 (0.8752)	. 16.7700	16.9292 (0.9079)
30 min (AR)	. 22.2530	22.1888 (0.2247)	. 17.6244	17.8720 (0.9036)	. 5.3137	5.2601 (0.1672)	. 19.5854	19.7348 (0.9072)
45 min (AR)	. 30.3289	30.2350 (0.2603)	. 23.3965	23.5171 (0.9981)	. 6.5633	6.5685 (0.7099)	. 25.6376	25.9390 (0.9258)
1 hr (AR)	. 36.3297	36.1042 (0.1467)	. 27.4002	27.5341 (0.9314)	. 7.4748	7.4618 (0.2858)	. 31.2881	31.9349 (0.9688)
2 hr (AR)	. 58.4704	58.5918 (0.6544)	. 49.5621	49.8873 (0.9539)	. 10.9815	10.8932 (0.0080)	. 53.8287	55.6411 (0.9771)
4 hr (AR)	. 107.3412	105.0585 (0.0482)	. 81.5804	80.6359 (0.0775)	. 17.0937	17.0386 (0.1817)	. 101.5470	103.8547 (0.8717)

Table 12: Results of the **out-of-sample** predictions of the category specific order flow measures on different sampling frequencies (Freq). The forecasting study is conducted over a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. These 32 weeks are divided into 8 periods of 4 weeks each, where the first 3 weeks are always considered as the in-sample estimation periods and the last weeks are always considered as the out-of-sample forecasting periods. 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 and third cell entries in parenthesis are the p-value from the modified Diebold-Mariano (mDM) test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the AR(p) benchmark model (RW, AR). P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the AR(p) benchmark model.

Freq	Market Orders Close		Stop-Loss Orders Close		Take-Profit Orders Close		Margin Call Orders Close	
	BM	hist. prices	BM	hist. prices	BM	hist. prices	BM	hist. prices
1 min (AR)	. 2.6810	2.6100 (0.0000)	. 1.8428	1.8324 (0.0000)	. 3.5039	3.4894 (0.0091)	. 0.7707	0.7687 (0.0693)
2 min (AR)	. 4.1646	4.0264 (0.0000)	. 2.8420	2.8272 (0.0000)	. 5.8825	5.8311 (0.0004)	. 1.1404	1.1373 (0.1203)
5 min (AR)	. 7.9576	7.6980 (0.0000)	. 4.9441	4.9223 (0.0000)	. 10.8902	10.7474 (0.0001)	. 2.1087	2.0974 (0.2095)
10 min (AR)	. 12.9142	12.6485 (0.0000)	. 7.4257	7.4011 (0.0002)	. 17.3042	17.2111 (0.0217)	. 3.1301	3.1080 (0.1991)
15 min (AR)	. 17.3679	17.1217 (0.0000)	. 9.4822	9.4698 (0.1016)	. 23.7763	23.8357 (0.6523)	. 3.9560	3.9300 (0.1594)
20 min (AR)	. 21.1034	21.0014 (0.1051)	. 11.0635	11.0712 (0.7544)	. 26.5161	26.6077 (0.7757)	. 4.5620	4.6193 (0.9044)
25 min (AR)	. 25.6057	25.4827 (0.1614)	. 12.7121	12.6941 (0.2426)	. 32.6157	32.7224 (0.7022)	. 5.6977	5.6005 (0.0596)
30 min (AR)	. 28.3720	28.2159 (0.0538)	. 14.0931	14.0252 (0.0062)	. 36.6937	36.4566 (0.0071)	. 5.7202	5.7616 (0.7729)
45 min (AR)	. 37.7322	37.9234 (0.8651)	. 17.6807	17.6993 (0.8152)	. 47.9815	47.8084 (0.0373)	. 7.0983	7.2341 (0.9954)
1 hr (AR)	. 44.0571	44.2877 (0.8383)	. 20.3353	20.4799 (0.9346)	. 55.7560	55.8053 (0.5955)	. 9.0440	9.1169 (0.8867)
2 hr (AR)	. 77.4225	77.4581 (0.5472)	. 30.9472	30.8688 (0.2532)	. 98.1701	98.2703 (0.5990)	. 13.4674	13.8434 (0.8713)
4 hr (AR)	. 23.6795	125.0778 (0.7443)	. 43.0342	42.0985 (0.0150)	. 145.2135	144.3137 (0.3253)	. 20.6128	21.4649 (0.8248)

Table 13: Results of the **out-of-sample** predictions of the category specific order flow measures on different sampling frequencies (Freq). The forecasting study is conducted over a period of 32 weeks starting on Monday the 6th of October 2003 and ending on Friday the 14th of May 2004. These 32 weeks are divided into 8 periods of 4 weeks each, where the first 3 weeks are always considered as the in-sample estimation periods and the last weeks are always considered as the out-of-sample forecasting periods. 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 and third cell entries in parenthesis are the p-value from the modified Diebold-Mariano (mDM) test with the null hypothesis that the RMSPE of the associated forecasting model is not smaller than the RMSPE of the AR(p) benchmark model. P-values in bold correspond to those cases where the RMSPE of the associated forecasting model is smaller than the RMSPE of the AR(p) benchmark model.

5 Conclusion

We have investigated the predictive power of past price changes for the aggregated order flow measure (Lyons (1995)), and eight transaction category specific order flows, based on a unique dataset from the currency trading platform, OANDA FXTrade. Our data contains detailed information on generally small retail investors' trading characteristics, currency positions and detailed information on order flows of several different order types. The main focus of this paper lies in investigating whether investors behave differently when they already hold a position compared to when they do not. The key question asked is: Does the current inventory matter? Answer: Yes!

We conduct forecasting studies on 12 intraday frequencies and find that those forecasting models incorporating information on order flow and price changes provide significantly better forecasts than benchmark models using only information on past order flow through AR specifications. Our in-sample and out-of-sample forecasting analyses show that investors' future order flow is affected by past price movements, and that market and limit order flows are much better predictable if those orders are submitted to close an existing position than if they are used to open one. This suggests evidence for the existence of a monitoring effect stating that investors value price information with respect to their current inventory. Monitoring effects are generally ignored in theoretical market microstructure models in which decisions regarding the submission of market and limit orders are modelled regardless of the inventory of the investor. Our study shows that monitoring effects play an important role, and that theoretical market microstructure models can be improved by the incorporation of monitoring or inventory effects. Furthermore, we find some evidence that stop-loss orders contribute to and take-profit impede self-reinforcing price movements, results which support the hypothesis of Osler (2005).

References

- ANAND, A., S. CHAKRAVARTY, & T. MARTELL (2005): “Empirical Evidence on the Evolution of Liquidity: Choice of Market versus Limit Orders by Informed and Uninformed Traders,” *Journal of Financial Markets*, 8, 288–308.
- AVERY, C. & P. ZEMSKY (1998): “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 88 (4), 724–48.
- BERGER, D., A. CHABOUD, S. CHERNENKO, E. HOWORKA, & J. WRIGHT (2008): “Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data,” *Journal of International Economics*, 75, 93–109.
- BIKHCHANDANI, S., D. HIRSHLEIFER, & I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades,” *Journal of Political Economy*, 100 (5), 992–1026.
- BJØNNES, G. H. & D. RIME (2005): “Dealer Behavior and Trading Systems in Foreign Exchange Markets,” *Journal of Financial Economics*, 75, 571–605.
- BLOOMFIELD, R., M. O’HARA, & G. SAAR (2005): “The “Make or Take” Decision in an Electronic Market: Evidence on the Evolution of Liquidity,” *Journal of Financial Economics*, 75 (4), 165–199.
- CAPLIN, A. & J. LEAHY (1996): “Trading Costs, Price, and Volume in Asset Markets,” *The American Economic Review*, 86, 192–196.
- CHAKRAVARTY, S. & C. HOLDEN (1995): “An Integrated Model of Market and Limit Orders,” *Journal of Financial Intermediation*, 4, 213–241.
- CHORDIA, T. & A. SUBRAHMANYAM (2004): “Order Imbalance and Individual Stock Returns,” *Journal of Financial Economics*, 72, 485–951.
- DANÍELSSON, J. & R. LOVE (2006): “Feedback Trading,” *International Journal of Finance and Economics*, 11, 35–53.
- 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.

- DE LONG, J. B., A. SHLEIFER, L. H. SUMMERS, & R. J. WALDMANN (1990): "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45 (2), 379–95.
- DEMSETZ, H. (1968): "The Cost of Transacting," *Quarterly Journal of Economics*, 82, 33–53.
- 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.
- (2005): "Meese-Rogoff Redux: Micro-Based Exchange-Rate Forecasting," *American Economic Review*, 95, 405–414.
- (2006): "Understanding Order Flow," *International Journal of Finance and Economics*, 11, 2–23.
- FENG, L. & M. SEASHOLES (2005): "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?" *Review of Finance*, 9, 305–351.
- FOUCAULT, T. (1999): "Order Flow Composition and Trading Costs in a Dynamic Limit Order Market," *Journal of Financial Markets*, 2, 99–134.
- FOUCAULT, T., O. KADAN, & E. KANDEL (2005): "Limit Order Book as a Market for Liquidity," *Review of Financial Studies*, 18, 1171–1217.
- GLOSTEN, L. R. (1994): "Is the Electronic Open Limit Order Book Inevitable?" *The Journal of Finance*, 49 (4), 1127–1161.
- HARRIS, L. E. (1998): "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems," *Financial Markets, Institutions & Instruments*, 7, 1–76.
- HARVEY, D., S. LEYBOURNE, & P. NEWBOLD (1997): "Testing the Equality of Prediction Mean Squared Errors," *International Journal of Forecasting*, 13, 281–291.
- HOLLIFIELD, B., R. MILLER, & P. SANDAS (2004): "Empirical Analysis of Limit Order Markets," *Review of Economic Studies*, 71, 1027–1063.

- KAHNEMAN, D. & A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–292.
- KAHNEMAN, D., J. L. KNETSCH, & R. H. THALER (1990): “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, 98 (6), 1325–48.
- KANIEL, R. & H. LIU (2006): “So What Orders Do Informed Traders Use?” *Journal of Business*, 79, 1867–1913.
- LEE, I. H. (1998): “Market Crashes and Informational Avalanches,” *Review of Economic Studies*, 65 (4), 741–59.
- LYONS, R. (1995): “Tests of Microstructural Hypothesis in the Foreign Exchange Market,” *Journal of Financial Economics*, 39, 321–351.
- LYONS, R. K. (2002): “The Future of the Foreign Exchange Market,” *in*: LITAN, R. & R. HERRING, eds., “Brookings-Wharton Papers on Financial Services,” Brookings Institution Press: Washington, DC, pp. 253–280.
- MARSH, I. W. & C. O’ROURKE (2005): “Customer Order Flow and Exchange Rate Movements: Is There Really Information Content?” Working Paper, Cass Business School, London.
- OSLER, C. L. (2005): “Stop-Loss Orders and Price Cascades in Currency Markets,” *Journal of International Money and Finance*, 24, 219–241.
- PARLOUR, C. (1998): “Price Dynamics in Limit Order Markets,” *Review of Financial Studies*, 11, 789–816.
- PAYNE, R. (2003): “Informed Trade in Spot Foreign Exchange Markets: An Empirical Investigation,” *Journal of International Economics*, 61, 307–329.
- RIME, D. (2003): “New Electronic Trading Systems in the Foreign Exchange Markets,” *New Economy Handbook*, 469–504.
- ROSU, I. (2009): “Liquidity and Information in Order Driven Markets,” Working paper, University of Chicago.
- SEPPI, D. J. (1997): “Liquidity Provision with Limit Orders and a Strategic Specialist,” *Review of Financial Studies*, 10, 103–150.

SHILLER, R. J. (2002): “Bubbles, Human Judgment, and Expert Opinion,” *Financial Analysts Journal*, 58 (3), 18–26.

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.

THALER, R. (1980): “Toward a Positive Theory of Consumer Choice,” *Journal of Economic Behavior & Organization*, 1 (1), 39–60.