

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 830

April 2005

Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data

David W. Berger, Alain P. Chaboud, Sergey V. Chernenko, Edward Howorka, Raj S.
Krishnasami Iyer, David Liu, and Jonathan H. Wright

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/.

Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data

David W. Berger, Alain P. Chaboud, Sergey V. Chernenko, Edward Howorka, Raj S. Krishnasami Iyer, David Liu, and Jonathan H. Wright *

Abstract: We study the association between order flow and exchange rate returns in five years of high-frequency intraday data from the leading interdealer electronic broking system, EBS. While the association between order flow and exchange rate returns has been studied in several previous papers, these have mostly used relatively short spans of daily data from older bilateral dealing systems and, usually, transaction counts instead of actual trading volume. Using a substantially longer span of recent high-frequency data and measuring order flow as actual signed trading volume, we find a strong positive association between order flow and exchange rate returns at frequencies ranging from one minute to one day, and a more modest but still sizeable association at the monthly frequency. We find, however, no evidence that order flow has predictive power for future exchange rate movements beyond, possibly, the next minute. Focusing on the behavior of order flow and exchange rates at the time of scheduled U.S. economic data releases, we find that the surprise components of these announcements are associated with order flow at high frequency immediately after the data releases. This finding seems inconsistent with a simple efficient markets view of how a public news announcement is incorporated into prices.

Keywords: order flow, foreign exchange, high-frequency data, news announcements, private information.

JEL Classification: F31, G14.

* International Finance Division (Berger, Chaboud and Chernenko) and Division of Monetary Affairs (Wright), Board of Governors of the Federal Reserve System, Washington DC 20551, and EBS, 535 Madison Avenue, New York, NY 10022 (Howorka, Iyer and Liu). Contact: alain.p.chaboud@frb.gov. We are grateful to Jon Faust, Mico Loretan, Rich Lyons and Krista Schwarz for helpful discussions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any person associated with the Federal Reserve System.

1. Introduction.

Several recent papers have documented a strong positive association at the daily frequency between exchange rate returns and order flow, measured as the count of buyer-initiated orders net of seller-initiated orders. Chief among these papers is that of Evans and Lyons (2002). They reported that a regression of Deutsche mark/dollar daily returns on daily order flow yielded an R^2 in excess of 60%, an amazingly strong result in the study of price discovery in foreign exchange markets. The authors argued that their finding showed that dispersed information or private information affecting the fundamental value of currencies became embedded in prices via order flow, a measure of net buying pressure. The empirical association between order flow and price changes had, to a less dramatic extent, also been previously documented by Lyons (1995) and Evans (2002).

Evans and Lyons (2002) based their analysis on four months of detailed trading data from Reuters' bilateral direct dealing platform (Reuters Dealing 2000-1) in 1996. In this dealing system, one dealer sends an electronic message to another dealer asking for a two-way price. A deal in which the party initiating the contact ultimately buys/sells is counted as a buyer-initiated/seller-initiated deal, respectively. Evans and Lyons (2002) noted that, in 1996, it was estimated that about 60 percent of interdealer volume was in direct trading, much of which conducted through the Reuters Dealing 2000-1 system, and 40 percent was in brokered trading. Brokered trading was at the time conducted mainly through voice brokers, who matched potential buyers and sellers for a fee, and, to a lesser extent, through one of the newly-introduced electronic broking platforms.

The structure of the interdealer spot foreign exchange market has, however, undergone a radical change since then. For the major currency pairs, the share of interdealer dealing conducted through brokered transactions has risen sharply relative to that conducted through direct dealing. The role of voice brokers in spot trading in the major currency pairs has substantially decreased, and a majority of trading is now conducted through electronic broking platforms. For instance, the Bank of England estimates that, based on its April 2004 survey, 66 percent of the interdealer spot business in the United Kingdom, the largest center of foreign exchange activity, is currently conducted through electronic broking platforms (Bank of England, 2004).¹ Today, two electronic platforms have a leading role in interdealer spot trading, one offered by EBS, and the other offered by Reuters, both effectively electronic limit order books. Over time, trading in each major currency pair has become very highly concentrated on only one of the two systems. Of the most-traded currency pairs, the top two, euro-dollar and dollar-yen are, across the globe, traded primarily on EBS, while the third, sterling-dollar is traded primarily on Reuters Dealing 2000-2.² The process of price discovery for each of these currency pairs now occurs within the computers of the respective electronic brokers -- the modern Walrasian auctioneers. As a result, the reference price at any moment for, say, spot interdealer euro-dollar is the current price on the EBS screen, and all dealers base the prices they quote to their customers on the EBS price. For more details on the EBS system, see Chaboud et al. (2004).

¹ The Bank of England cautions that, as some large dealers did not report this information, that proportion is likely to be underestimated.

² Many more currency pairs are traded actively on each of the systems.

Our goal in this paper is to study the relationship between order flow and exchange rate movements at various frequencies, from one minute to one month, using data from EBS spanning January 1999 through February 2004. Our dataset has several advantages relative to that used by Evans and Lyons (2002), and our paper offers several new contributions. First, as explained above, our data represent, for the first time, a majority of trading in the interdealer spot market for the two most-traded currency pairs under the current market structure. Second, the availability of intradaily data at high frequency over a long period of time allows us to study the impact of order flow at various frequencies, ranging from one minute to one month. We also analyze the forecasting power of order flow at various frequencies. Third, our measure of order flow is signed trading volume, in base currency, rather than the signed deal count which Evans and Lyons used, allowing us to take into account transaction size. Fourth, using minute-by-minute data, we examine the relationship between order flow and exchange rate movements at the times of U.S. macroeconomic data releases, when exchange rates often experience sharp jumps. Fifth, the prices we use to calculate exchange rate returns are based on true executable quotes, and not on indicative quotes which may or may not have represented prices truly available to traders.

As the EBS system is a limit order book, it could potentially have a different price discovery process from the bilateral direct dealing system that Evans and Lyons (2002) considered. In particular, as discussed in several recent papers (e.g., Bloomfield, O'Hara, and Saar (2003)), informed traders in a limit order book setting may optimally choose at times to “make” liquidity, that is to place limit orders in the system, instead of “taking” liquidity, that is to hit existing orders. As order flow is traditionally measured as the net,

per unit of time, of buys and sales by takers, such behavior by informed traders may imply a very different empirical relationship between price movement and order flow in a limit order book relative to a bilateral trading system. However, despite all the potential for differences with previous work, our results at a daily frequency are broadly consistent with those earlier reported by Evans and Lyons. A regression of daily exchange rate returns on contemporaneous daily order flow yields a significant positive coefficient with an R^2 of 45% for euro-dollar and 50% for dollar-yen. Furthermore, at higher frequencies, we again find significantly positive coefficients, with a regression of one-minute returns on one-minute order flow giving an R^2 of about 30% for both currency pairs. We find, however, little, if any, forecasting ability for order flow, with the exception of the first lag of order flow at one-minute frequency. Today's order flow, in particular, does not help predict tomorrow's exchange rate return.

A few other recent papers have also begun to study the relationship between order flow and exchange rates in the EBS and Reuters electronic broking systems that have recently come to be widely used in the interdealer market, generally using shorter samples of mostly daily data. Killeen, Lyons, and Moore (2002) studied daily EBS data on the German mark/French franc exchange rate in 1998, and Hau, Killeen and Moore (2002) studied data from 1998 and 1999 on the German mark and euro exchange rates viz-a-viz the dollar, as did Dunne, Hau and Moore (2004). None of these used intradaily order flow data. Love and Payne (2003) studied order flow and exchange rates using high-frequency intraday data from the Reuters Dealing 2000-2 dealing system from September 1999 to July 2000. Breedon and Vitale (2004) collected and analyzed EBS and Reuters euro-dollar intradaily data from August 2000 to January 2001.

The plan for the remainder of this paper is as follows. In section 2, we briefly describe the data that we use in this paper, referring the reader to the companion paper by Chaboud et al. (2004) for more details. In section 3, we first report our results on the contemporaneous relationship between order flow and exchange rate movements. We then study the predictive power of order flow for future exchange rate returns at various frequencies. In section 4, we consider the link between order flow and the unexpected component of macroeconomic announcements. Section 5 concludes.

2. The Data.

The data used in this paper present an aggregate and purely numerical picture of the EBS system, and do not contain any information on the identity of any market participants. The data consist of three time series at the one-minute frequency: the volume of buyer-initiated trades in each minute, the volume of seller-initiated trades in each minute, and the one-minute exchange rate returns. The data span January 1999 to February 2004 for the dollar-yen and euro-dollar currency pairs, and trading volume is expressed in base currency. The volume of buyer-initiated trades means the total volume transacted where a quote to buy euros for dollars, or dollars for yen, placed in the EBS system by one dealer is dealt on by another dealer, who is then seen as the aggressor, or the initiator of the transaction. The volume of seller-initiated trades is of course defined similarly. Order flow is the difference between the volume of buyer-initiated trades and that of seller-initiated trades. A positive value of order flow therefore indicates a net excess of buyer-initiated trades.

We exclude all data collected from Friday 17:00 New York time through Sunday 17:00 New York time from our sample, as foreign exchange trading activity during these hours is minimal. We also drop certain holidays and days of unusually light volume: December 24-December 26, December 31-January 2, Good Friday, Easter Monday, Memorial Day, Labor Day, Thanksgiving and the following day, and July 4 (or, if this is on a weekend, the day on which the Independence Day holiday is observed). Similar conventions were adopted by Andersen, Bollerslev, Diebold and Vega (2003). To construct a minute-by-minute price series, we use the midpoint of the best bid and the best ask price in the system at the start of each minute and calculate continuously compounded one-minute returns. That is, we compute our returns as 10,000 times the one-minute change in the log exchange rate. Our returns have therefore the interpretation of (approximately) the percentage change in the exchange rate multiplied by 100, and so the units can be thought of as basis points of exchange rate movements.

Table 1 reports some summary statistics of order flow for both currency pairs. The table shows the mean and standard deviation of the one-minute order flow, and of the order flow aggregated to the five-minute, ten-minute, one-hour and one-day frequencies. The proportion of observations for which the order flow is positive is also shown at each frequency. For both currency pairs, the mean order flow is positive. Although it is only very slightly positive at the one-minute frequency, it turns out that if one aggregates to the daily frequency, about 70 percent of days have positive order flow. This is somewhat surprising, and it remains a puzzle at this point. A similar phenomenon was found by

Dunne, Hau and Moore (2004). Evans and Lyons (2002) also found that cumulative dollar-yen order flow in Reuters data was trending upwards over their sample period.³

3. Order Flow and Exchange Rate Returns.

3.1 The contemporaneous relationship

To study the contemporaneous relationship between order flow and exchange rate movements at various frequencies, we ran regressions of the form

$$r_{t,h} = \alpha + \beta o_{t,h} + \varepsilon_{t,h} \quad (1)$$

where $r_{t,h}$ refers to the returns over a horizon h and $o_{t,h}$ refers to the order flow over that same horizon. The horizon h can be one minute, or we can aggregate both the returns and the order flow to five-minute, ten-minute, one-hour, one-day, or one-month frequencies. Note that equation (1) includes an intercept, meaning that our regressions effectively demean the order flow series. The coefficient β is often interpreted as the price impact of a trade and as a possible measure of liquidity (with a higher β reflecting lower liquidity). It is however simply a reduced form coefficient; only under strong identifying assumptions does it actually tell us the effect of an exogenous shock to order flow on exchange rate returns.

The results of this regression for the euro-dollar and dollar-yen currency pairs are shown in Table 2, using heteroskedasticity-robust standard errors. At the one-minute frequency, the coefficient β is significantly positive and the R^2 is 35% for the euro-dollar and 29% for the dollar-yen. An excess of buyer-initiated trades is associated with

³ Brandt and Kavajecz (2004) also find that order flow is positive on average in interdealer trading at most maturities in the U.S. Treasury market, based on GovPX data.

a rising price, with an order imbalance of 1 billion (of the base currency) estimated to lead to about a half-percent appreciation (precisely 55 basis points in euro-dollar and 72 basis points in dollar-yen).⁴ At the daily frequency, the R-squareds are about 50%, and the estimates of β are still significantly positive, about 0.4 percent for each currency pair, a bit less than the Lyons and Evans estimates. At the monthly frequency, the R^2 falls to 19% and 33% for the euro-dollar and dollar-yen currency pairs respectively, and the coefficient β is estimated to be about 0.2 percentage points. The estimate of β is fairly consistent at different frequencies, reflecting the fact that the cross-correlation function of one-minute order flow and one-minute returns is large and positive at lag zero, but is generally very small at other leads and lags.

Figure 1 shows scatterplots of the daily returns against daily order flow in the euro-dollar and dollar-yen currency pairs. A systematic, approximately linear, positive relation can clearly be seen for both currency pairs. The relation is clearly not the result of a small number of outliers, and no nonlinearity is evident in the association.

Table 3 reports the estimated contemporaneous association between order flow and exchange rate returns, for both currency pairs, at one-minute and one-day frequencies over five different subsamples (1999, 2000, 2001, 2002 and 2003), to see how this has evolved over time. The association between order flow and price movements rose slightly from 1999 to 2001, especially at the daily frequency, but fell back a fair bit in

⁴ One could also express these results as the amount of order flow in a given span of time associated with a one basis point movement in the exchange rate. In this case, for instance, an order imbalance of 18.2 million euros in a minute is associated with a one basis point movement in euro-dollar, and an order imbalance of 13.9 million dollars is associated with a one basis point movement in dollar yen. A higher reading of order flow per basis point of movement could be interpreted as evidence of higher market liquidity, although one should keep in mind that we have shown only an empirical association, not formal causation from order flow to exchange rate movement.

2002 and 2003, suggesting that, if anything, average market liquidity has recently increased.

3.2 The (lack of) predictive content of order flow for future returns.

The coefficient β in the regressions above only measures the contemporaneous association between order flow and exchange rate returns, but does not tell us anything about the question of the predictive power of order flow for future exchange rate returns. As a simple test of the predictive content of order flow for future exchange rate returns, we regressed returns at time t on lagged returns and lagged order flow at frequencies ranging from one minute to one day, for both currency pairs. The results of these Granger-causality regressions, using five lags of returns and order flow, are shown in tables 4 and 5.

In regressions at the one-minute frequency in tables 4 and 5, the coefficients on the first lag of order flow are clearly positive and significant, with a magnitude of about one tenth that of the coefficients on contemporaneous order flow found in table 2. Although the F-statistics for these regressions show clear joint statistical significance for lagged order flow, their adjusted R-squareds are, however, essentially zero. With the possible exception of the first lag in dollar-yen at five-minute frequency, other coefficients on lagged order flow at various frequencies in tables 4 and 5 are generally either not statistically significant and/or very small in magnitude. At the daily frequency, in particular, lagged order flow appears to add essentially no explanatory power to the regressions in either euro-dollar or dollar-yen. Figures 2 and 3 show the autocorrelation coefficients of order flow at frequencies ranging from one minute to one day, with 95 percent confidence intervals centered at zero shown as shaded areas. In both euro-dollar

and dollar-yen, the highest levels of (positive) autocorrelation are found at lag 1 at the one-minute frequency, helping to explain the regression results at that frequency.

Overall, our tests therefore fail to detect any predictive power of order flow for future exchange rate returns at the daily frequency, and find only very modest predictability at the one-minute frequency.

4. Order Flow and Exchange Rate Returns After Macroeconomic Announcements.

The association between order flow and exchange rates may well reflect, at least in part, the impounding of private or dispersed information into currency prices. As such, the association may be quite different in the immediate aftermath of a major macroeconomic news announcement, the canonical public news event. Nevertheless, Evans and Lyons (2003) and Love and Payne (2003) argue that much of the impact of macroeconomic news may be transmitted to exchange rates via induced transactions rather than the simple direct adjustment of price that we might expect from the release of public information.

To study the relationship between order flow and exchange rates in the immediate aftermath of macroeconomic announcements, we take eight important regularly scheduled U.S. data releases: initial unemployment claims, CPI, GDP (advance release), housing starts, retail sales, PPI, nonfarm payrolls, durable goods sales and the trade balance, all of which come out at 8:30am. For each of these releases, using the actual released value and the *ex-ante* expectation taken from the Money Market Services survey, we measure the surprise component of the data release as the deviation of the actual release from the survey expectation. Anderson, Bollerslev, Diebold and Vega

(2003) found that the surprise components of many of these data releases had significant systematic price effects.

For each announcement, we regressed the order flow in each of the first ten minutes after the announcement times on the unexpected component of the announcement. The results for the first three minutes are reported in table 6, where the units show the effect of a one standard deviation surprise on the order flow in millions of dollars. All announcements have a significant effect on the order flow in the minute immediately after the announcement, except for CPI and PPI releases. Higher-than-expected claims releases are statistically significantly associated with net orders to buy euros for dollars and to sell dollars for yen. Higher-than-expected payrolls, GDP, trade balance⁵, housing starts and retail sales data are all significantly associated with net (filled) orders to sell euros for dollars and to buy dollars for yen. In other words, news of stronger real economic activity is associated with orders by market “takers” to buy dollars. The effects are quantitatively largest for payrolls and GDP releases, smaller for trade balance and retail sales releases, and smaller again for claims and housing starts, broadly consistent with the relative sensitivity of the exchange rate to the different types of announcements that Andersen, Bollerslev, Diebold and Vega (2003) found. The results are also consistent with the findings of Love and Payne (2003). With our longer sample, however, we are able to analyze the individual effect of each different type of announcement, rather than pooling all the U.S. announcements and scaling them by their standard deviations, as Love and Payne do. We find that the headline data surprise only impacts order flow for a very short interval after each type of announcement. As can be

⁵ Meaning a smaller-than-expected trade deficit.

seen in table 6, most announcement surprises do not have a significant effect on the order flow in the second and third minute after the announcement, and the same is true for subsequent minutes (not shown).

One possible interpretation of the statistical association between order flow in the minute after the announcement and the surprise component of the announcement is that information about the new full-information level of the exchange rate is indeed being transmitted to the market through order flow. In this interpretation, order flow has a causal role. Another possible interpretation is that much of the order flow immediately after an announcement may consist of dealers taking quotes that were placed in the system before the news release. Leaving live quotes near the pre-announcement price in the system at the moment of a news announcement is seemingly not a profit-maximizing strategy, yet Carlson and Lo (2004), studying in great detail the impact in foreign exchange market of one news announcement, argue that some foreign exchange traders choose to do this.⁶ In this interpretation, the observed order flow is simply a byproduct of the price movement arising from the near-instantaneous reaction of the exchange rate to the news announcement, but order flow has no causal role.

Table 7 reports the estimates of equation (1), estimated at the one minute frequency only on announcement days and only in the minute after the release, broken down by announcement type. In light of the fact that data surprises have systematic effects on exchange rate returns (Andersen, Bollerslev, Diebold and Vega (2003)) and of our earlier finding that order flow is correlated with data surprises, we should expect that

⁶ Carlson and Lo (2002) argue that these traders follow at all times a rigidly disciplined strategy, attempting to cover in the interdealer market, with only a small fixed profit margin, positions opened while trading with their customers.

order flow is correlated with exchange rate returns in short windows around news announcements, and table 7 confirms that this is indeed the case. The magnitude of the β coefficients in these regressions is, in many cases, higher than those derived from including all one-minute periods in our sample.

5. Conclusion

We have studied the relationship between order flow and exchange rate returns in a long span of recent data from the EBS electronic trading platform at frequencies ranging from one minute to one month, using order flow measured as signed trading volume and exchange rate returns based on actual executable quotes. The data represent a majority of the worldwide interdealer trading volume in spot foreign exchange in the two leading currency pairs over the period January 1999-February 2004. We confirm the finding of Evans and Lyons (2002), who used a shorter span of daily data on signed deal count on a bilateral trading system, that order flow is significantly associated with contemporaneous exchange rate returns, and we expand their conclusion to a much-wider range of frequencies. One might interpret this as evidence of a mechanism whereby private or dispersed information is embodied in exchange rates, although the causality between prices and order flow is hard to disentangle even with very-high-frequency data. With the possible exception of the first lag at a one-minute frequency, we find no evidence that, in this long span of data, order flow has predictive power for future exchange rate returns

Focusing on the precise time of eight types of scheduled U.S. economic data releases, we also studied the association between the surprise components of scheduled

public data releases and order flow in the EBS system. We find that high-frequency order flow is indeed correlated with the surprise component of these news announcements, with the association generally only statistically significant in the first minute after data releases. This finding seems inconsistent with a simple efficient markets view of how public information becomes embedded into prices, where order flow ought not to be correlated with public news already in the information set.

References

- Anderson, T.G., T. Bollerslev, F.X. Diebold and C. Vega (2003): Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange, *American Economic Review*, 93, pp.38-62.
- Bank of England (2004): The Foreign Exchange and Over-the-Counter Derivatives Markets in the United Kingdom, *Bank of England Quarterly Bulletin*, 44, pp.470-484.
- Bloomfield, Robert J., O'Hara, Maureen and G. Saar (2003): The 'Make or Take' Decision in an Electronic Market: Evidence on the Evolution of Liquidity, NYU Stern School of Business Working Paper No. FIN-02-029.
- Brandt, Michael, and K. A. Kavajecz (2004): Price Discovery in the U.S. Treasury Market: The Impact of Orderflow and Liquidity on the Yield Curve, *Journal of Finance*, 59: pp.2623-2654.
- Breedon, F. and P. Vitale (2004): An Empirical Study of Liquidity and Information Effects of Order flow on Exchange Rates, *European Central Bank Working Paper*, 424.
- Carlson, J.A. and M. Lo (2004): One Minute in the Life of the DM/US\$: Public News in an Electronic Market, forthcoming, *Journal of International Money and Finance*.
- Chaboud, A.P., S.V. Chernenko, E. Howorka, R.S. Krishnasami Iyer, D. Liu and J.H. Wright (2004): The High-Frequency Effects of U.S. Macroeconomic Data Releases on Prices and Trading Activity in the Global Interdealer Foreign Exchange Market, *International Finance Discussion Paper*, 823.
- Dunne P., H. Hau and M. Moore (2004): Macroeconomic Order Flows: Explaining Equity and Exchange Rate Returns, unpublished manuscript.
- Evans, M.D.D. (2002): FX Trading and Exchange Rate Dynamics, *Journal of Finance*, 57, pp.2405-2447.
- Evans, M.D.D. and R.K. Lyons (2002): Order Flow and Exchange Rate Dynamics, *Journal of Political Economy*, 110, pp.170-180.
- Evans, M.D.D. and R.K. Lyons (2003): How is Macro News Transmitted to Exchange Rates? unpublished manuscript.
- Fleming, M.J. and E.M. Remolona (1999): Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information, *Journal of Finance*, 54, pp.1901-1915
- French, K. and R. Roll (1986): Stock Return Variance: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics*, 14, pp.71-100.

- Goodhart, C.A.E., T. Ito and R. Payne (2000): One Day in June 1993: A Study of the Working of the Reuters' 2000-2 Electronic Foreign Exchange Trading System in "The foreign exchange market: Empirical studies with High-Frequency Data", C.A.E. Goodhart and R. Payne (eds.), Macmillan, London.
- Hau, H., W. Killeen and M. Moore (2002): How Has the Euro Changed the Foreign Exchange Market?, *Economic Policy*, 0, pp.149-177..
- Killeen, W., R. Lyons and M. Moore (2001): Fixed versus Flexible: Lessons from EMS Order Flow, NBER Working Paper 8491.
- Love, R. and R. Payne (2003): Macroeconomic News, Order Flow and Exchange Rates, LSE Financial Markets Group Discussion Paper 475.
- Lyons, R.K. (1995): Tests of Microstructural Hypotheses in the Foreign Exchange Market, *Journal of Financial Economics*, 39, pp.321.351.

Table 1: Order Flow Summary Statistics

Horizon, h		Euro	Yen
1 minute	Mean	0.0005	0.0004
	Standard Deviation	0.0194	0.0136
	Proportion Positive	0.52	0.52
5 minutes	Mean	0.0023	0.0022
	Standard Deviation	0.0525	0.0383
	Proportion Positive	0.53	0.53
10 minutes	Mean	0.0045	0.0044
	Standard Deviation	0.0778	0.0590
	Proportion Positive	0.53	0.54
1 hour	Mean	0.0270	0.0263
	Standard Deviation	0.2050	0.1758
	Proportion Positive	0.56	0.57
1 day	Mean	0.6484	0.6301
	Standard Deviation	1.1835	1.2436
	Proportion Positive	0.72	0.70

This table reports some summary statistics for order flow (in billions of base currency) at the one-minute frequency, and aggregated to five-minute, ten-minute, hourly and daily frequencies. Proportion positive means the number of observations at that frequency for which the measured order flow is positive, divided by the total number of observations for which it is nonzero.

Table 2: Estimates of Equation (1) at Different Horizons

Horizon, h		Euro	Yen
1 minute	$\hat{\beta}$	0.55	0.72
	Standard error	0.0006	0.0008
	R-squared	0.35	0.29
5 minute	$\hat{\beta}$	0.54	0.69
	Standard error	0.0010	0.0014
	R-squared	0.46	0.42
10 minutes	$\hat{\beta}$	0.53	0.66
	Standard error	0.0013	0.0017
	R-squared	0.48	0.45
1 hour	$\hat{\beta}$	0.47	0.54
	Standard error	0.0028	0.0032
	R-squared	0.48	0.48
1 day	$\hat{\beta}$	0.39	0.38
	Standard error	0.012	0.010
	R-squared	0.45	0.50
1 month	$\hat{\beta}$	0.18	0.20
	Standard error	0.049	0.036
	R-squared	0.19	0.33

The exchange rate returns are measured as 100 times the log exchange rate return, and so can be interpreted as the return, in percentage points, corresponding to an order flow of +1 billion base currency. Heteroskedasticity-robust standard errors are used.

Table 3: Estimates of Equation (1) at Different Horizons over Different Years

Currency	Horizon, h		Whole Sample	1999	2000	2001	2002	2003
Euro	1 minute	$\hat{\beta}$	0.55	0.55	0.69	0.64	0.46	0.49
		S.E.	0.0006	0.0012	0.0016	0.0015	0.0011	0.0011
		R ²	0.35	0.35	0.35	0.34	0.35	0.37
	1 day	$\hat{\beta}$	0.39	0.36	0.46	0.52	0.34	0.35
		S.E.	0.012	0.025	0.035	0.030	0.018	0.024
		R ²	0.45	0.46	0.41	0.54	0.58	0.47
Yen	1 minute	$\hat{\beta}$	0.72	0.88	0.77	0.77	0.68	0.58
		S.E.	0.0008	0.0022	0.0019	0.0019	0.0017	0.0016
		R ²	0.29	0.30	0.31	0.31	0.33	0.27
	1 day	$\hat{\beta}$	0.38	0.44	0.41	0.47	0.44	0.29
		S.E.	0.010	0.028	0.022	0.026	0.021	0.020
		R ²	0.50	0.49	0.59	0.57	0.65	0.47

The exchange rate returns are measured as 100 times the log exchange rate return, and so can be interpreted as the return, in percentage points, corresponding to an order flow of +1 billion base currency. Heteroskedasticity-robust standard errors are used.

Table 4: Granger Causality Regressions of Euro/Dollar Returns

	1 minute	5 minute	1 hour	1 day
Order Flow:				
L1	5.78 (0.09)	0.60 (0.18)	-0.33 (0.56)	-1.42 (2.26)
L2	0.23 (0.09)	-0.41 (0.18)	0.58 (0.56)	0.55 (2.27)
L3	-0.44 (0.09)	0.07 (0.18)	0.18 (0.56)	2.00 (2.27)
L4	-0.43 (0.09)	-0.33 (0.18)	0.09 (0.56)	2.50 (2.27)
L5	-0.23 (0.09)	-0.42 (0.18)	0.84 (0.56)	-5.01 (2.21)
Return:				
L1	-0.01 (0.00)	-0.02 (0.00)	0.02 (0.01)	-0.06 (0.04)
L2	0.00 (0.00)	-0.02 (0.00)	-0.01 (0.01)	0.02 (0.04)
L3	-0.00 (0.00)	-0.01 (0.00)	-0.01 (0.01)	-0.01 (0.04)
L4	-0.00 (0.00)	-0.00 (0.00)	-0.01 (0.01)	0.01 (0.04)
L5	-0.01 (0.00)	0.01 (0.00)	-0.02 (0.01)	0.09 (0.04)
Constant	-0.00 (0.00)	0.00 (0.01)	0.00 (0.09)	1.77 (3.26)
Num.Obs.	1838670	367265	29345	1276
Adjusted R-sq	0.00	0.00	0.00	0.01
Fstatistic	885.13	4.90	0.81	1.42
p-value	0.00	0.00	0.54	0.21

This table reports the results of regressions of exchange rate returns over four different horizons on five lags of order flow and five lags of returns over the same horizons. Returns are expressed in basis points. Order flow is expressed in billions of base currency. Standard errors are reported in parentheses below the coefficients. The F-statistics and p-values are from the F-tests that lagged order flow coefficients are jointly significantly different from zero.

Table 5: Granger Causality Regressions of Dollar/Yen Returns

	1 minute	5 minute	1 hour	1 day
Order Flow:				
L1	9.04 (0.12)	2.78 (0.24)	0.97 (0.65)	2.22 (2.27)
L2	0.39 (0.12)	0.12 (0.24)	-0.53 (0.65)	0.36 (2.32)
L3	-0.31 (0.12)	-0.20 (0.24)	0.12 (0.65)	-2.66 (2.33)
L4	-0.29 (0.12)	0.14 (0.24)	-0.16 (0.65)	-0.15 (2.32)
L5	-0.25 (0.12)	-0.10 (0.23)	-0.84 (0.63)	3.79 (2.25)
Return :				
L1	-0.04 (0.00)	-0.03 (0.00)	-0.01 (0.01)	-0.07 (0.04)
L2	0.00 (0.00)	-0.03 (0.00)	-0.00 (0.01)	0.02 (0.04)
L3	-0.00 (0.00)	-0.02 (0.00)	0.00 (0.01)	0.02 (0.04)
L4	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.01)	0.02 (0.04)
L5	-0.01 (0.00)	-0.00 (0.00)	0.00 (0.01)	-0.04 (0.04)
Constant	-0.00 (0.00)	-0.01 (0.01)	-0.05 (0.08)	-3.39 (2.95)
Num.Obs.	1838670	367265	29345	1276
Adjusted R	0.00	0.00	-0.00	-0.00
F-statistic	1145.6	27.71	0.89	1.01
p-value	0.00	0.00	0.49	0.41

This table reports the results of regressions of exchange rate returns over four different horizons on five lags of order flow and five lags of returns over the same horizons. Returns are expressed in basis points. Order flow is expressed in billions of base currency. Standard errors are reported in parentheses below the coefficients. The F-statistics and p-values are from the F-tests that lagged order flow coefficients are jointly significantly different from zero.

Table 6: Slope Coefficient in Regression of Order Flow in Different Minutes after News Announcements on the Headline Surprise (t-stats in parentheses)

Announcement	Euro-\$ Order Flow			\$-Yen Order Flow		
	8:30	8:31	8:32	8:30	8:31	8:32
Initial Claims	30.5 (6.10)	9.25 (3.14)	1.99 (0.78)	-5.7 (-3.28)	-1.1 (-0.77)	-0.2 (-0.18)
CPI	3.4 (0.52)	2.00 (0.31)	-2.60 (-0.67)	-0.06 (-0.02)	-5.4 (-1.82)	-1.6 (-0.58)
GDP	-56.1 (-3.26)	-10.4 (-0.94)	-0.11 (-0.01)	39.5 (4.47)	12.7 (2.54)	6.6 (0.90)
Housing Starts	-21.7 (-3.63)	4.30 (0.81)	-4.45 (-0.97)	7.0 (2.62)	0.7 (0.32)	1.5 (0.73)
Nonfarm Payrolls	-53.9 (-2.98)	-5.86 (-0.47)	4.46 (0.48)	43.9 (3.15)	12.2 (1.90)	6.5 (1.61)
PPI	2.69 (0.29)	2.56 (0.69)	1.48 (0.26)	2.6 (0.66)	-2.1 (-1.28)	-2.1 (-0.92)
Retail Sales	-48.7 (-3.68)	-5.03 (-0.72)	-8.38 (-1.49)	20.8 (4.26)	4.0 (1.19)	-0.7 (-0.29)
Trade Balance	-45.4 (-6.17)	-12.34 (-1.59)	-5.34 (-0.96)	23.3 (4.84)	7.9 (1.75)	0.38 (0.14)

The surprises are expressed in standard deviation units and the order flow is expressed in millions of dollars, so the coefficients have the interpretation of the effect of a one-standard deviation surprise on the order flow in millions base currency. The column labeled 8:30 refers to the order flow in the minute from 8:30:00 to 8:31:00, and likewise for the columns labeled 8:31 and 8:32. Results were calculated for minutes up to and including 8:39, but are mostly insignificant, and are not shown, so as to conserve space. The t-statistics are constructed using robust standard errors. Entries that are statistically significant at the 5 percent level are shown in bold.

Table 7: Estimates of Equation (1) at the One-Minute Frequency in the Minutes After Announcements of Different Types (t-stats in parentheses)

Announcement		Euro	Yen
Initial Claims	$\hat{\beta}$	0.58	1.05
	Standard error	0.04	0.06
	R-squared	0.43	0.52
CPI	$\hat{\beta}$	0.59	0.97
	Standard error	0.11	0.17
	R-squared	0.32	0.36
GDP	$\hat{\beta}$	1.04	1.11
	Standard error	0.31	0.22
	R-squared	0.38	0.59
Housing Starts	$\hat{\beta}$	0.54	0.70
	Standard error	0.08	0.10
	R-squared	0.43	0.44
Nonfarm Payrolls	$\hat{\beta}$	0.89	1.01
	Standard error	0.22	0.18
	R-squared	0.22	0.36
PPI	$\hat{\beta}$	1.07	0.74
	Standard error	0.10	0.15
	R-squared	0.69	0.30
Retail Sales	$\hat{\beta}$	0.85	0.66
	Standard error	0.09	0.10
	R-squared	0.62	0.40
Trade Balance	$\hat{\beta}$	0.71	0.59
	Standard error	0.11	0.11
	R-squared	0.42	0.32

The exchange rate returns are measured as 100 times the log exchange rate return, and so can be interpreted as the return, in percentage points, corresponding to an order flow of +1 billion base currency. Regressions are run at the one-minute frequency only in the minutes after announcements of the specified types. Heteroskedasticity-robust standard errors are used.

Figure 1: Returns and Order Flow

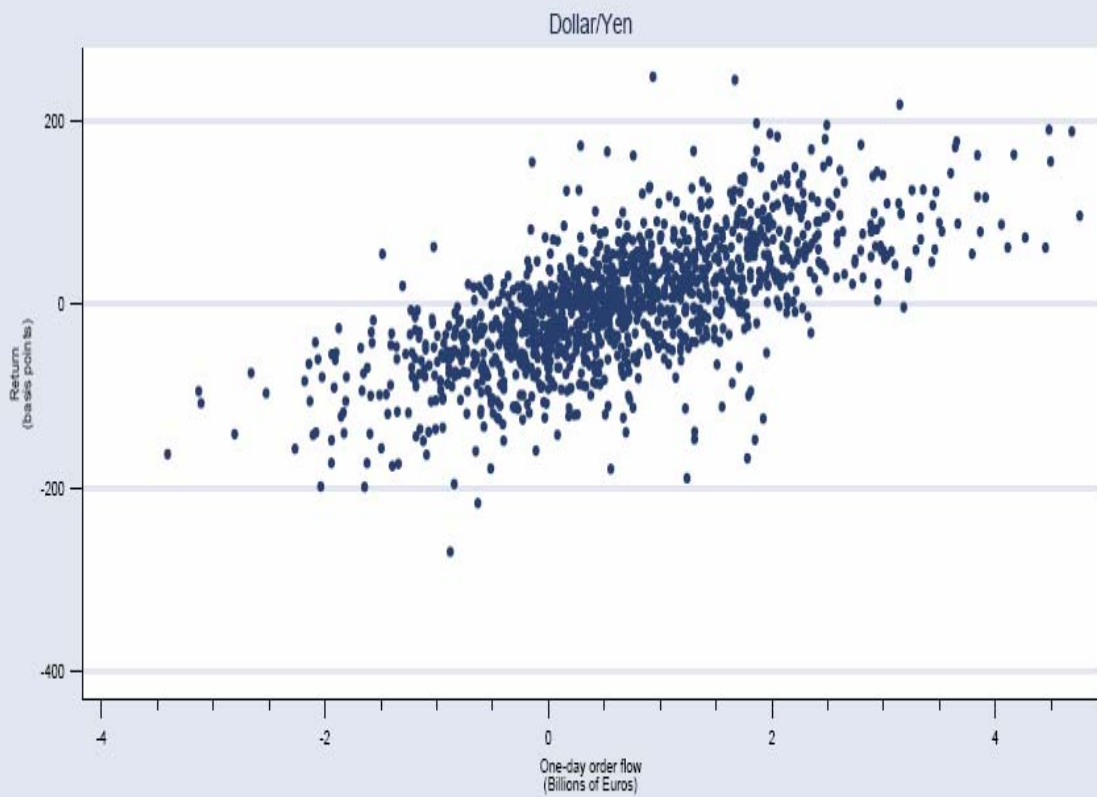
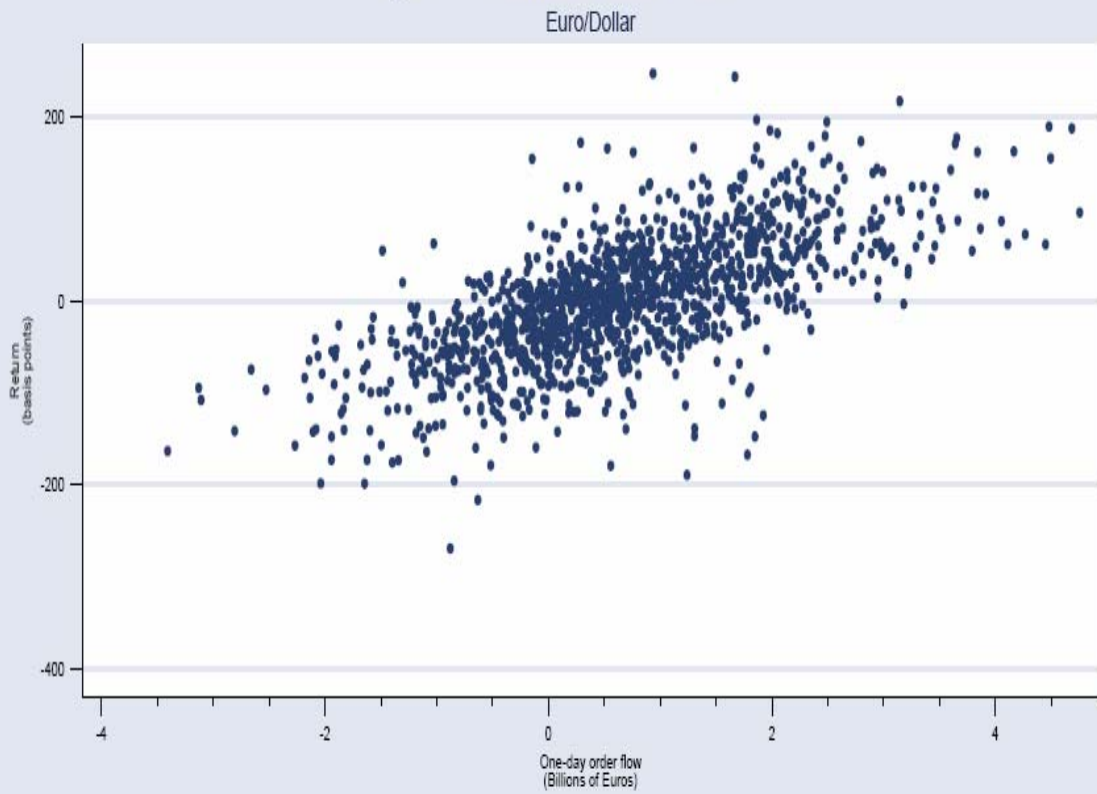
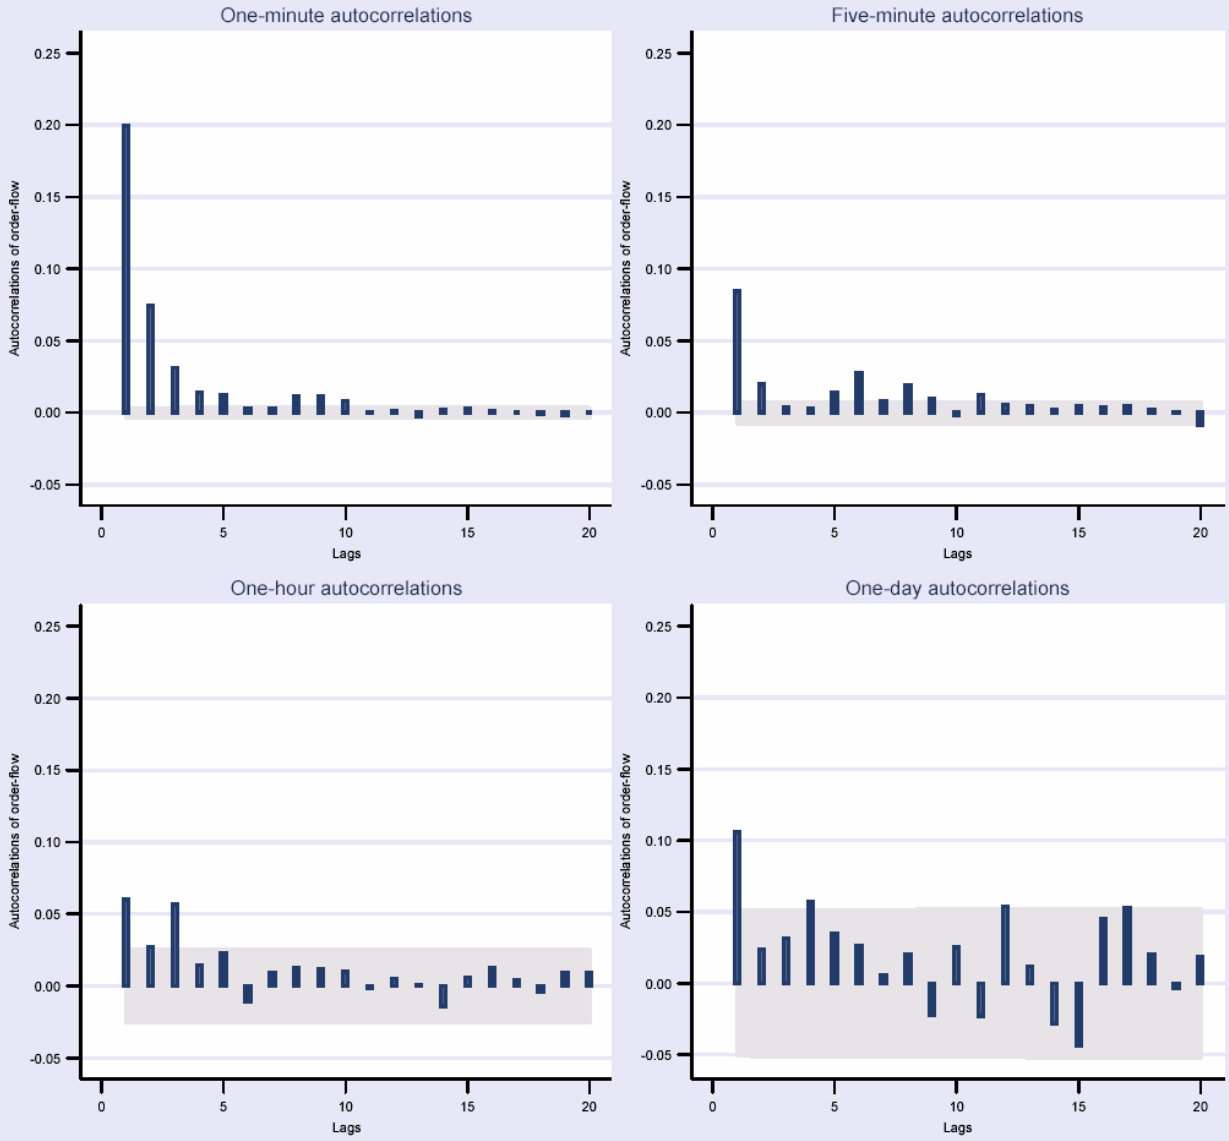
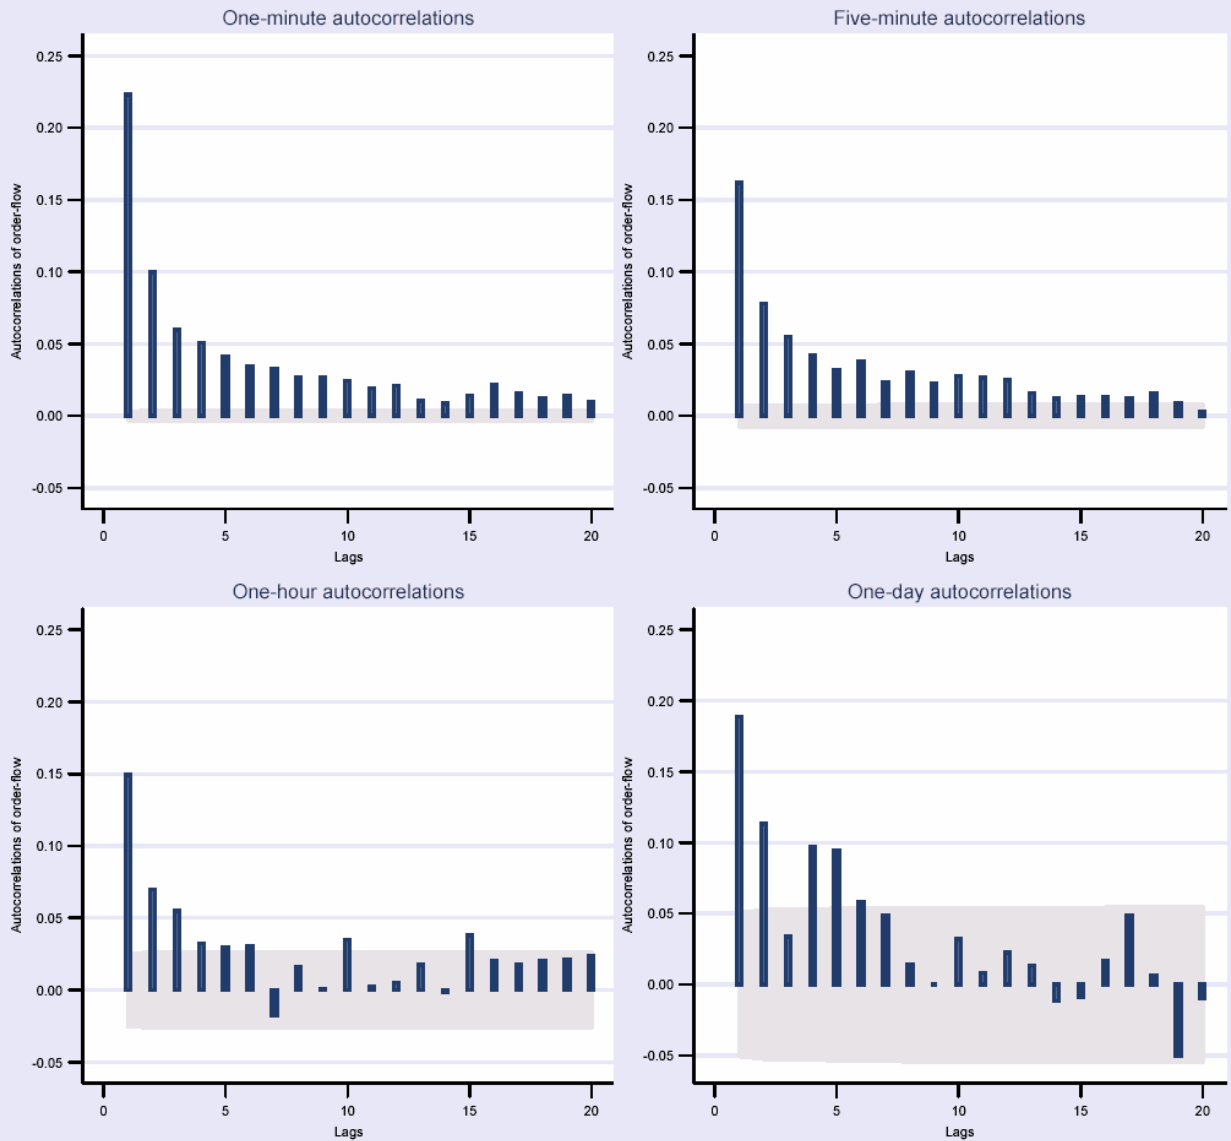


Figure 2: EUR/USD Order-flow Autocorrelations at Different Frequencies



Shaded region denotes 95% confidence interval

Figure 3: USD/JPY Order-flow Autocorrelations at Different Frequencies



Shaded region denotes 95% confidence interval