HEDGE FUNDS AND THE ORIGINS OF PRIVATE INFORMATION IN CURRENCY MARKETS

Carol Osler and Vitaliy Vandrovych

Abstract

This paper provides evidence that hedge funds provide private fundamental information in currency markets and that non-financial corporations and mutual funds do not. On this basis we suggest that the information that customers bring to the foreign exchange market is actively acquired, rather than passively acquired. Our database of currency transactions, the most disaggregated to date, includes ten different categories of market participants of which six correspond to customers. Orders of banks in every size category carry information, consistent with now-standard theory that banks gather information from observing customer trades. Theory does not indicate whether banks should be better informed than their customers. Our results suggest that banks are better informed than their individual customers, possibly because they aggregate information from many customers. [*Key words: microstructure, exchange rates, asymmetric information*] [JEL codes G1, F3.]

February, 2009

[a]Corresponding author: Brandeis International Business School, Brandeis University, Mailstop 32, Waltham, MA 02454. cosler@brandeis.edu

[b]State Street Global Advisors, Boston, MA

All errors are the responsibility of the authors.

HEDGE FUNDS AND THE ORIGINS OF PRIVATE INFORMATION IN CURRENCY MARKETS

This paper addresses two key questions in foreign exchange market microstructure: First, Who brings private information to the market? Second, What is that nature of that information? To address these questions we evaluate the information content of orders executed by the Royal Bank of Scotland, currently ranked fifth among foreign exchange dealing banks worldwide (Euromoney 2007), over a 16-month period in 2001-2002. Following Anand et al. (2005), among others, we measure information content as the signed return immediately after execution ("price impact") over time horizons ranging from five minutes to two weeks. If orders carry exchange-rate relevant information then after buy orders for a given currency that currency should appreciate, and after sell orders it should depreciate.

Our data permit us to disaggregate those placing the orders far more finely than previous studies. We identify ten groups of market participants: six different groups of end users, or equivalently bank customers and four groups of banks. The finest previous disaggregation of end-user order flow considered only three groups of customers: leveraged investors, institutional investors, and corporate customers (Evans and Lyons 2007). The only other study that disaggregates interdealer order flow, Bjonnes et al. (2009), likewise disaggregates banks into four groups.

Standard theory indicates that there is asymmetric information between dealers and their customers, with customers bringing information to the market (Kyle 1985; Glosten and Milgrom 1985). Consistent with this hypothesis, evidence indicates that customer order flow in foreign exchange markets does contain information (Evans and Lyons 2007). It has not yet been ascertained, however, which customers bring the information to the market.

Our six categories of end-users are leveraged investors, institutional investors, broker dealers, governments including central banks, large corporations, and middle-market corporations. Among these groups, only leveraged investors appear to bring information to the market. The estimated return after a $10 million leveraged-investor stop-loss order is four basis points at thirty minutes and roughly double that after six hours. Orders from other end-users have little or no statistically significant price impact at any horizon and thus do not appear to carry exchange-rate relevant information.

Our four groups of banks are RBS itself, large global liquidity providers (e.g., Deutschebank), regional liquidity providers (e.g., Santander), and banks that primarily service customers. Current theory suggests that dealers gather information from observing their customers' trades (Lyons 2001; Evans and Lyons 2002). Dealers might infer that a currency is undervalued, for example, if customers are actively buying it. Consistent with this hypothesis, we find that all the bank groups are informed.

Theory does not indicate whether the banks should be better or worse informed than their customers. On the one hand, dealers get noisy signals of any given customer's information; on the other hand, banks get signals from many customers. The net effect on the strength of their information signal is thus an empirical question. Our evidence suggests that dealers know more than their individual customers or, equivalently, the aggregation effect dominates. Specifically, we find that the price impact of bank trades remains strong for up to two weeks, while the price impact of leveraged-investor trades loses significance after six hours.

Our results also have implications for the nature of information in currency markets. Lyons (2001) hypothesizes that customer order flow carries dispersed, passively-acquired information about fundamentals. The information could concern either real factors, such as aggregate economic activity at home or abroad, or financial factors, such as aggregate risk aversion or wealth. The real-side information that economic activity is strong, for example, could be communicated by strong corporate demand for foreign currency. The financial information that aggregate risk aversion is high, by contrast, could be communicated by low currency demand from financial customers. In either case, the information itself would be acquired passively, meaning the corporation or financial institution would not intentionally seek out the information but would instead reflect the information unknowingly. Individual agents, viewing their own activity in isolation, might not recognize the broader picture. Dealers, who see order flow from many agents, could still recognize the broader pattern.

By distinguishing among end-user types we can test whether the information in customer order flow is passively acquired. Passively-acquired information about the real economy would come from the order flow of large and/or middle-market corporations. Our results indicate, however, that corporate orders do not carry exchange-rate information. Passively-acquired information about financial factors they would learn it from the order flow of all three types of financial customers: institutional investors, broker-dealers, and leveraged investors. Our results indicate, however, that the orders of institutional investors and broker-dealers do not carry exchange-rate information. In short, these results suggest that the information carried by foreign exchange customer order flow is not passively acquired.

The information in customer order flow could also be acquired actively, meaning through the conscious effort of individuals to anticipate the market. The active trading community – meaning hedge funds, currency trading associations (CTAs), and the like – devotes much time and effort to generating private information about exchange rates. Among other things, they focus intensely on upcoming macro statistical releases, for example, Dealers provide their active customers with frequent summaries of upcoming release dates and times, forecasts of key figures, and extensive discussion of related macroeconomic developments. Entire firms are devoted to collecting and disseminating market forecasts. These forecasts are generated by combining existing public information, their currency traders' own observations of the world, and the traders' own interpretive framework. In finding that leveraged investors are the only end users whose trades carry information, we provide direct evidence that information in currency markets is actively acquired.

Indirect evidence for the importance of actively-acquired information comes from two existing studies. Rime *et al.* (2007) provides evidence that information about upcoming statistical releases is embedded in exchange rates. MacDonald and Marsh (1996) show that there is heterogeneity in the exchange-rate forecasts of professional forecasts, despite the forecasters' shared reliance on public information. This heterogeneity "translates into economically meaningful differences in forecast accuracy," and is heavily influenced by "the idiosyncratic interpretation of widely available information" (p. 665). Finally, MacDonald and Marsh show that heterogeneity has a substantial influence on trading volume.

The hypothesis that leveraged investors are better informed than other end-users is consistent with important institutional features of currency markets today. Traders at firms that import and export goods and services, for example, are typically not permitted to speculate in spot and forward markets. As discussed in Carlson, Dahl, and Osler (2008), this policy reflects rogue trader risk and associated high administrative costs associated with such speculation. Since most commercial customers do not trade at high frequencies, the benefits from synthesizing information to create accurate exchange-rate forecasts is unlikely to outweigh those high costs for all but the most active commercial firms.

Even many financial traders are uninterested in basing their trades on exchange-rate relevant information – or at least that is the view of many market participants. Taylor and Farstrup (2006), for example, in their survey of currency management, state:

> [T]here are key participants in foreign exchange markets … that are not always seeking profit derived from their currency positions …. [I]n this category are international equity managers. While some managers factor in currency positions as they go about picking foreign stocks, most are attempting to add value through stock, sector, and region bets rather than currency plays (p. 10, italics in original).

Institutional investors can, of course, justify this focus by citing the familiar academic evidence that exchange rates approximate a random walk and thus returns cannot be predicted. Broker dealers typically trade as a representative for individuals and small institutional investors, so they are no more likely to be informed than such investors. Last but not least, central banks certainly attempt to monitor the market but, since they are not in the business of profiting from their trades, there may be no reason to expect those trades to be informed.

This paper has two sections and a conclusion. Section I describes our data. Section II presents our methodology and discusses our results. Section III concludes.

I. DATA

Our data comprise all price-contingent orders placed at the Royal Bank of Scotland over the period June 1, 2001 through September 20, 2002 in three currency pairs: euro-dollar, dollar-yen, and sterling-dollar. The data include information about each order's placement time, size, trigger rate, direction (buy or sell), the "desk" (group of traders) at which the order was placed, a code identifying the type of agent that placed the order, and the order's status at the end of the sample period: executed, deleted, or open. We estimate the execution time as the time the exchange rate (measured at one-minute intervals) first comes within 2 pips of the order's trigger rate.[1]

Price-contingent orders represent a subset of all order flow. A price-contingent order instructs a dealer to trade a specified amount at the market price once the currency trades at a specified price level ("trigger price"). These orders come in two types: stop-loss orders and take-profit orders. A stop-loss order instructs the dealer to buy (sell) currency if the price rises (falls) to the trigger. A take-profit order instructs the dealer to buy (sell) if the price falls (rises) to the trigger. These orders appear to have a significant influence on exchange rates, amplifying the response to news (Savaser 2006) and creating non-linearities in intraday return dynamics (Osler 2003, 2005).

Our ten groups of order-placing agents include six groups of end-users and four groups from the dealing community. The six end-user groups are: leveraged investors, meaning primarily hedge funds and similar organizations called Commodity Trading Arrangements or CTAs; institutional investors, meaning mutual funds, pension funds, and other low-leverage asset-management firms; large "corporate" (non-financial) customers, examples of which would be General Electric; middle-market corporate customers; broker dealers like Brown Brothers Harriman, who undertake forex transactions associated with the securities transactions of clients who typically range from individual retail investors to modest-sized mutual funds.[2]

The four dealing-community groups are: Royal Bank itself, including its spot/forward dealers at various locations around the world and the exotic options desk, located in London; other major forex dealers like Citibank or Deutschebank, whom we label Global Liquidity

[1] We verified the accuracy of this approach by calculating the frequency of order executions at different time horizons relative to order placement and comparing that frequency with the known frequency in an earlier dataset. The earlier dataset is described in Osler (2003).

[2] A few broker dealers in the dataset were known to be trading exclusively on behalf of one or two leveraged investors. In consequence, their trades are included among "leveraged investors."

Providers; other medium-sized forex dealers like Santander or PNB, whom we label Regional Liquidity Providers; and smaller forex dealers who provide little liquidity to the interbank market but instead trade in the interbank market almost entirely to service customers.[3]

Table 1 provides basic descriptive statistics. During our sample period Royal Bank received 36,806 orders in these currencies worth a total of $193.9 billion. The orders can be roughly evenly divided into three categories: orders placed by end-users, orders placed by other dealers, and internal RBS orders.

Table 2 compares placement patterns between stop-loss and take-profit orders, and between different order sizes. Take-profit orders are more frequent, accounting for 60 percent of all orders. The mean order size, of about $5 million, does not vary much between stop-loss and take-profit orders. Since the size distribution is positively skewed, the median order size of $3 million is quite a bit lower than the average order size. The largest placed order is worth $750 million, though only 2 percent of orders were worth $25 million or more. Very few orders (only ten) were worth less than $100,000. Only 27 percent of orders are actually executed, a figure that is fairly consistent across order types and order sizes. The largest executed order, worth $473 million, was placed by corporate customer in the euro-dollar market. The average order is open about 4 days but the median time open is only about 5 hours; the difference indicates that a minority of orders are open for months.

Table 2 also compares orders placed by customer type. The share of take-profit orders is highest among corporate customers, reaching 75 percent for large corporations and 83 percent

[3] The exotic options dealers place orders in part to hedge their barrier options. These are options that either come into existence, or disappear when the underlying exchange rate touches a certain pre-specified level. When such levels are crossed dealers are required to make large changes in their forward market delta-hedges. If the option disappears, any existing delta-hedge must be unwound; if the option appears, a delta-hedge must be put in place. The option dealers can ensure these transactions take place in a timely manner by placing orders with their spot dealers.

for middle-market corporations. Global liquidity providers – the very biggest banks – and broker-dealers are the only groups that execute more stop-loss orders than take profit orders.

Figure 1, panels A-C, graphically outline intraday patterns in the placement and execution of these orders. Order placement peaks first during Asian trading, then again during the London morning, and finally a third time during the morning in New York/late afternoon in London.[4] This pattern reflects, at least in part, the basic biological rhythms associated sleeping and eating. Order placement falls whenever traders leave work for the day and rises whenever a market opens and traders return to their desks; similarly, it falls when traders leave for lunch and rises when they come back.

The small number of orders placed during Asian trading reflects the bank's London/European base. The high peak in the London afternoon partly reflects the coincidence of active trading in New York and London. It also, however, reflects the formal placement of orders previously received by dealers at RBS and other banks as they leave for the day. Many of the orders would have been received earlier but noted only on a pad of paper; when someone else needs to monitor the order it must be formally entered into the bank's computer system.

Order execution has a distinctly different pattern than order placement. Instead of multiple sharp peaks, it rises to a plateau after London trading begins and only declines again at the end of London trading. This can be explained by analyzing daily patterns in the two factors that jointly dominate order execution: (i) the size of the order book; and (ii) exchange-rate volatility. The order book will be largest from mid-day London through the end of trading for the day, a period that coincides with the highest exchange-rate volatility, as shown in Figure 2.

The statistical analysis below necessarily focuses on executed orders. Table 3 presents the descriptive information for number, average size, and total value of executed orders for each

[4] Placement patterns for stop-loss and take-profit orders, not shown separately, are strongly similar to each other.

customer type. The 9,950 executed orders have average value of roughly $5 millions; in aggregate these orders were worth roughly $47.5 billion.

II. METHODOLOGY, RESULTS, DISCUSSION

This section presents our central tests for the information content of price-contingent orders. We find that leveraged investors and large corporations are the only end-users whose orders have a strong price impact. Members of the dealing community appear to have information similar to that of the leveraged investors. We also find that the price impact of orders is concave.

A. Methodology

We evaluate the price impact of executed orders, disaggregated by customer- and order-type (stop-loss, take-profit). Following the literature (e.g., Anand *et al.* 2005), price impact is defined as the (log) change in price relative to the order's trigger rate immediately following the order's execution. We consider eight time horizons: five minutes, thirty minutes, one, six, and twelve hours, one day, one week, and two weeks.

Our baseline regression aggregates executed orders across all three currencies and assumes that order flow has a linear price impact on returns:

$$(s_{t+k} - s_t)*10000 = \sum_{i=1}^{8} \beta_i * V_{\boldsymbol{i,t}}^{\boldsymbol{SL}} + \sum_{i=1}^{8} \delta_i * \boldsymbol{V}_{\boldsymbol{i,t}}^{\boldsymbol{TP}} + \varepsilon_t \qquad (1)$$

The dependent variable is the signed return measured in basis points, since s_t is the log of the exchange rate expressed as foreign currency per USD. Time horizon is indicated by *k*. $V_{i,t}^{SL}$ ($V_{i,t}^{TP}$) is the signed value of executed stop-loss orders (take-profit orders) of customer type *i*, measured in millions of USD. The raw return is multiplied by plus one if the customer 'buys' US dollars and by minus one otherwise: if customer purchases (sales) bring higher (lower) prices, as indicated in the broader literature, the coefficients should have a positive sign.

The assumption that order flow has a linear impact on price is standard in the empirical literature, but only because it serves as a useful first approximation. There are good reasons to expect the relationship to be concave, instead. It is widely appreciated that splitting large orders into smaller individual transactions and timing the execution of each trade carefully can reduce the impact of large trades. This is demonstrated in the theoretical treatment of Bertsimas and Lo (1998) and it is standard practice among dealers. Indirect evidence of a concave effect is provided in Berger *et al*. (2006), which shows that the proportionate price impact of minute-by-minute interdealer order flow, while positive in all cases, declines with the amount. Hasbrouck (1991) finds that equity order flow declining proportionate impact.

To examine the possibility of a concave relation between order flow and returns we postulate a logarithmic relation, which provides a parsimonious alternative to the linear function form. To conserve the sign of our orders, we take the log of absolute order size and then reassign the buy-sell direction, and denote the resulting variable $log(V_{i,t}^{x})$, $x \in (SL,TP)$:[5]

$$(2) \qquad (s_{t+k} - s_t) * 10000 = \sum_{i=1}^{8} \beta_i * \log(\boldsymbol{V}_{\boldsymbol{i,t}}^{\boldsymbol{SL}}) + \sum_{i=1}^{8} \delta_i * \log(\boldsymbol{V}_{\boldsymbol{i,t}}^{\boldsymbol{TP}}) + \varepsilon_t$$

B. Results

The results from estimating the linear relation between order flow and returns, Equation (1), are presented in Table 4; the results from the non-linear relation, Equation (2), are presented in Table 5. The top of the table shows price impacts for stop-loss orders; the bottom presents

[5] We also estimate a quadratic form of this regression, in which we added the signed squared value of each order size to the right-hand-side of equation (1). The results, suppressed to save space, indicated a concave price impact function but they have the implausible implication that very large orders would have a negative price impact.

price impacts for take-profit orders. For each order type, end-user groups are reported first and dealing-community groups last.

We draw three main conclusions from this analysis. First, leveraged investors are the only informed end-user type, while all bank categories are informed. Second, banks have more information than their informed customers. Third, the relation between order flow and returns is concave.

Who has information? Leveraged investors are the only end users whose orders have a statistically significant price impact. A $10 million leveraged-investor stop-loss buy order appreciates the currency by 4.3 basis points after thirty minutes and by 9.9 basis points after six hours, according to the non-linear estimates. Thereafter the effect is lower and is no longer statistically significant. There is a strong, positive relation between post-trade returns and executed stop-loss orders for all four types of dealers. This evidence is consistent with the hypothesis that leveraged investors are the only end users that bring information to the currency market, and that dealers benefit from that information.

It is not clear, at a theoretical level, whether banks would be more or less informed than their customers. Banks necessarily gain an imperfect signal of any given customer's information, but by aggregating these noisy signals banks might ultimately achieve a signal that is stronger than the ones originally perceived by the customers. Our evidence suggests that the aggregation effect dominates: dealers know more than their individual customers. At any given time horizon, the executed bank orders generally about the same price impact than leveraged-investor orders, with the exception of Regional liquidity providers, whose price impact is statistically larger. For banks, however, the price impact of stop-loss orders generally continues to rise – and to remain statistically significant – even after the twelve-hour horizon at which leveraged-investor trades are no longer informative (Figure 4A). The bank price impacts remain economically large, and are generally statistically significant, even at the two-week horizon.

Take-profit orders do not appear to carry information about upcoming returns, whether they are placed by end-users or by banks. With respect to the end users, for example, only one of the 96 total coefficients (across the linear and concave regressions) is positive and statistically significant, as required for the trades to carry information. A few of these coefficients are negative and significant, of which most concern central banks. This is difficult to interpret, though it certainly suggests that central banks are not exploiting the market for their own profit. The results for dealers are similar. Almost all the price-impact coefficients for take-profit orders are insignificant, and those that reach standard significance levels are generally negative rather than positive.

Why might stop-loss orders predict returns for informed agents while take-profit orders do not? We hazard two guesses based on institutional information. First, the use of take-profit orders may be driven by option-related logic rather than information. Suppose a corporate customer needs currency to import some intermediate product; alternatively, suppose an internationally-invested index fund just received an infusion of new funds and, in order to expand its ownership of the index of some country, it needs to purchase currency. In both cases it is not necessarily wise to trade at the first price of the day since they could, instead, place a take-profit order early in the day which gives them a good chance of trading at a more attractive price later on. (If the order is not executed they would fill the order at the market price at the end of the day.) The option to trade later is free even though intraday volatility makes it valuable: even with no information whatsoever about the exchange rate's likely movements they would be wise to place the take-profit and exploit that option.

Difference in hedging practices across the two order types could also influence measured post-trade returns. Dealers hedge take-profit orders by placing limit orders in the interdealer market, a hedge that is so secure that price risk on take-profit orders is conventionally assigned to the dealers rather than their customers. A limit order for amount Q at price P tells the market that the dealer is willing to trade up to Q so long as he gets price P or better. Thus limit orders act as absorbing barriers: A price rise that triggers a customer take-profit sell order at price P (requiring the dealer to buy the currency in question) would simultaneously trigger the associated limit sell order at price P placed by the dealer in the interdealer market. This limit sell order would normally be executed at the limit price P exactly, causing the price to stop moving but not to reverse course. This could explain why take-profit orders have so little price impact in general, and why any impact they do have is delayed.

Dealers have no useful way to hedge stop-loss orders. Stop-loss orders are not allowed on the interdealer brokers, and placing a stop-loss order with another dealer would not provide good protection since price risk on stop-loss orders is, by market convention, borne by the agent placing the order. Thus when a customer's stop-loss buy order is triggered by a price rise, the dealer (who has to sell to the customer) will quickly eliminate the associated short position by making a purchase in the interdealer market. This generally requires the dealer to choose between placing a market order or a limit order. A market order seems most likely, since stop-loss orders can trigger price cascades (Osler 2005) and in general dealers consider these orders risky ("every stop-loss order is a potential relationship breaker," says one dealer). So a price rise that triggers a customer's stop-loss buy order would thus trigger market buy orders in the interdealer market, pushing prices up further.[6]

[6] We note in passing that the two market conventions discussed here are confirmed in separate orders data from the Royal Bank of Scotland, described in Osler (2003, 2005), that include execution prices.

Linear or Concave? A comparison of Tables 4 and 5 shows that our qualitative conclusions about information are consistent regardless of whether we assume a linear or concave relation between order flow and returns. Stop-loss orders placed by leveraged investors and by anyone in the dealing community have a significant impact that begins almost immediately and broadly rises with time horizon, while take-profit orders do not carry information.

Though the *F* statistics are highly statistically significant, the regressions' explanatory power, reported at the bottom of each table, is small. This is to be expected, given the high level of noise commonly associated with financial markets. Explanatory power is almost invariably low when financial returns are examined at such high frequencies – the most notable exception, the daily regressions using aggregate interdealer order flow (e.g., Evans and Lyons 2002, Berger et al. 2006) prove the rule. Their order flow aggregates millions of trades while our regressions examine single transactions.

The explanatory power varies with time horizon, rising between the 5-minute and 1-hour horizons, leveling off, and then generally falling. A similar pattern is observed in Berger et al.'s (2006) analysis of aggregate interdealer order flow.

Notably, the explanatory power is consistently higher when the relation is concave than when the relation is linear, with the difference averaging roughly ten percent of the linear value. Figure 4B juxtaposes the estimated price impact of leveraged-investor stop-loss orders from Equations (2) with corresponding estimates from Equation (1). The price impact values are roughly the same for orders below $20 million but diverge sharply for larger orders. While the linear relation predicts that a $100 million buy leveraged-investor stop-loss order brings an 85 basis point appreciation, the non-linear relation predicts an appreciation of only 21 basis points.

III. CONCLUSIONS

This paper provides evidence that leveraged investors are the only type of currency end user that brings private information to the markets. Our data comprise price-contingent orders placed at the Royal Bank of Scotland over 16 months during 2001-2002. Our data allow the agents placing orders to be disaggregated into ten categories: six categories of end-users – leveraged investors, institutional investors, broker-dealers, large corporations, middle-market corporations, and governments and central banks – and four categories of banks – Royal Bank itself, other major banks, regional banks, and smaller banks that primarily service local customers. We evaluate the price impact of orders placed by each group over time horizons ranging from five minutes to two weeks. We address three questions: Who brings private information to the market? Is the information acquired actively or passively? Do dealers know more or less than their informed customers?

We find that leveraged investors are the only end users whose trades consistently carry information. Their executed stop-loss orders have statistically and economically significant price impact, which rises from approximately four basis points per $10 million order at thirty minutes to ten or more basis points at the one-day horizon. Though it is standard in the literature to estimate price impact as a linear function of order flow, our results indicate that the price impact of individual trades is in fact a concave function of order size.

Theory suggests that members of the dealing community learn exchange-rate relevant information from observing customers order flow (Evans and Lyons 2002). Our results are consistent with this hypothesis, since stop-loss orders from all four bank groups have a statistically significant price impact that is statistically indistinguishable from the impact of leveraged investors' stop-loss orders at horizons below 12 hours. At longer horizons, however, dealers trades appear to carry information while the leveraged-investor trades do not. This suggests that dealers, by observing the trades of many customers, are ultimately better informed than their customers taken individually.

Our finding that corporations and institutional investors do not appear to bring information to the market suggests that private information is not a passive reflection of existing economic conditions. Our finding that leveraged investors do bring information to the market suggests that private information is instead actively acquired through the informed interpretation of publicly available information.

REFERENCES

Anand Amber, Sugato Chakravarty, Terence 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: 289-309.

Anand, Amber, and Avanidhar Subrahmanyam, 2006. "Information and the Intermediary: Are Market Intermediaries Informed Traders in Electronic Markets?" Typescript, Syracuse University.

Berger, David, Alain Chaboud, Sergei Chernenko, Edward Howorka, and Jonathan Wright (2006). "Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data," Board of Governors of the Federal Reserve System, International Finance Discussion Papers No. 830.

Bertsimas, Dimitris, and Andrew W. Lo, 1998. Optimal control of execution costs. Journal of Financial Markets 1: 1-50.

Carlson, John A., Christian Dahl, and Carol Osler. 2008. Short-run exchange-rate dynamics: Theory and Evidence. Working paper.

Dominguez Kathryn and Freyan Panthaki. 2006. What Defined 'News' in Foreign Exchange Markets, *Journal of International Money and Finance* 25: 168-198.

Evans and Lyons. 2002. Order Flow and Exchange Rate Dynamics. *Journal of Political Economy* 110: 170-180.

Evans and Lyons. 2007. Exchange Rate Fundamentals and Order Flow. *NBER Working Paper 13151*. National Bureau of Economic Research.

Foucault, Thierry, 1999. Order Flow Composition and Trading Costs in a Dynamic Limit Order Market, *Journal of Financial Markets* 2: 99-134.

Froot Keneth A. and Tarun Ramadorai. 2005. Currency Returns, Intrinsic Value, and Institutional-Investor Flow. *Journal of Finance* 60(3): 1535-1566.

Glosten, Lawrence R. and Paul R. Milgrom. 1985. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents. Journal of Financial Economics 14(1): 71-100.

Harris, Lawrence, 1998. Optimal dynamic order submission strategies in some stylized trading problems, *Financial Markets, Institutions, and Instruments* 7: 1-75.

Hasbrouck Joel. 1991. Measuring the information content of stock trades. *The Journal of Finance* 46: 179-207.

Ho T. and H.Stoll. 1983. The Dynamics of Dealer Markets Under Competition. Journal of Finance 38: 1053-1074.

Kyle A. 1985. Continuous Auctions and Insider Trading. Econometrica 53: 1315-1335.

Lyons Richard K. 2001.The Microstructure Approach to Exchange Rates. MIT Press, Cambridge, MA. 333pp.

MacDonald, Ronald and Ian Marsh. 1996. Currency Forecasters are Heterogeneous: Confirmation and Consequences. Journal of International Money and Finance 15(5): 665-685.

Osler Carol. 2003. Currency Orders and Exchange-Rate Dynamics: Explaining the Success of Technical Analysis. Journal of Finance. 58: 1791-1819.

_______________ 2005. "Stop-Loss Orders and Price Cascades in Currency Markets,' Journal of International Money and Finance 24: 219-41.

Osler, Carol, 2007. "Foreign Exchange Microstructure: A Survey," forthcoming in Bruce Mizrach, Ed., Encyclopedia of Complexity and System Science (New York: Springer Science & Business Media)

Osler, Carol, Alexander Mende, and Lukas Menkhoff, 2007. "Price Discovery in Currency Markets," Typescript, Brandeis University

Rime, Dagfinn, Lucio Sarno, and Elvira Sojli, 2007. "Exchange-rate Forecasting, Order Flow, and Macro Information," Norges Bank Working Paper 2007-2.

Savaser Tanseli. 2006. "Exchange-Rate Response to Macro News: Through the Lens of Microstructure," Typescript, Williams College.

Shleifer, Andrei, 1986. "Do Demand Curves Slope Down?" The Journal of Finance 41: 579-90.

Taylor, Andrew, and Adam Fartstrup (2006), Active Currency Management: Arguments, Considerations, and Performance for Institutional Investors (CRA RogersCasy International Equity Research: Darien, Connecticut).

Table 1. Placed Orders By Client Type

Data underlying this table comprise all stop-loss and take-profit orders placed at the Royal Bank of Scotland over the period from June 1st 2001 till September 20th 2002 in three currency pairs: euro-dollar, sterling-dollar, and dollar-yen. Totals may not add due to rounding.

	Number of orders	Percent of orders	Dollar Value ($ Billions)	Percent of Order Value
All Orders	36,781	100.0	193.7	100.0
End Users	12,061	32.9	62.6	32.3
Leveraged Investors	2,356	6.4	12.6	6.5
Institutional Investors	1,352	3.7	7.8	4.0
Brokers/Dealers	2,384	6.5	13.7	7.1
Central Banks and Government Agencies	544	1.5	3.9	2.0
Large Corporations	3,623	9.9	20.1	10.4
Middle Market Corp.	1,802	4.9	4.5	2.3
Banks	24,720	67.2	104.6	67.7
Royal Bank Internal	13,105	35.6	77.3	39.9
Global Liquidity Providers	4,162	11.3	19.8	10.2
Regional Liquidity Providers	1,352	3.7	7.5	3.9
Customer-Service Banks	6,101	16.6	26,5	13.7

Table 2. Descriptive Statistics, All Placed Orders

Data underlying this table comprise all stop-loss and take-profit orders placed at the Royal Bank of Scotland over the period from June 1st 2001 till September 20th 2002 in three currency pairs: euro-dollar, sterling-dollar, and dollar-yen. Totals may not add due to rounding.

		All Orders	Stop-Loss	Take-Profit	Very Small Orders (<=$100,000)	Large Orders (> $25 Mill.)
Number of Orders		36,781	16,064	20,717	10	831
Share of Orders (%)		100	43.7	56.3	0.03	2.3
Size ($ Mill.):	Mean	5.3	5.5	5.1	0.05	47.5
	Median	3.0	3.1	2.6	0.05	32.8
Distance to Mkt. (in pips):	Mean	77.4	78.4	76.7	76.1	143.5
	Median	43.9	40.0	47.0	71.4	59.2
Share Within 1/2 std of 1-Day Return (%)		28.7	30.4	27.4	20.00	18.8
Share Within 1 std of 1-Day Return (%)		57.5	60.6	55.0	30.00	44.9
Share Within 2 std of 1-Day Return (%)		82.9	83.5	82.4	70.00	70.2
Days Open:	Mean	4.2	3.4	4.7	0.3	11.2
	Median	0.2	0.1	0.2	0.3	0.4
Share Executed (%)		27.1	24.6	29.0	20.0	21.7
Share by End-User Category (%):						
Leveraged Investors		6.4	35.9	64.1	0.00	2.0
Institutional Investors		3.7	39.8	60.2	0.00	4.0
Brokers/Dealers		6.5	67.5	33.5	0.12	3.8
Central Banks and Government Agencies		1.5	27.4	72.6	0.00	7.3
Large Corporations		9.9	28.6	71.4	0.00	2.3
Middle-Market Corporations		4.9	18.6	81.4	0.03	0.2
Share by Bank Category (%):						
Royal Bank Internal Orders		35.6	44.7	55.3	0.04	3.0
Global Liquidity Providers		11.3	64.4	35.6	0.00	1.6
Regional Liquidity Providers		3.7	45.2	54.8	0.00	0.5
Customer-Service Banks		16.6	37.8	62.2	0.04	1.3

Table 3. Descriptive Statistics, Executed Orders

Data underlying this table comprise all stop-loss and take-profit orders executed by the Royal Bank of Scotland over the period June 1st 2001 till September 20th 2002 in three currency pairs: euro-dollar, sterling-dollar, and dollar-yen. Totals may not add due to rounding.

			Stop-Loss			Take-Profit		
	Total Number	**Stop-Loss Share**	**Number**	**Average ($ mill.)**	**Total ($ mill.)**	**Number**	**Average ($ mill.)**	**Total ($ mill.)**
End-Users								
Leveraged Investors	750	28.1	211	5.5	1,168	539	5.2	2,816
Institutional Investors	413	41.4	171	4.6	786	242	3.9	1,141
Broker/Dealers	485	59.4	288	4.6	1,333	197	4.7	833
Central Banks and Government Agencies	164	31.1	51	3.9	199	113	7.9	890
Large Corporations	1,274	24.7	315	6.6	2,089	959	4.4	4,213
Middle-Market Corporations	754	17.5	132	2.9	291	622	1.9	1,173
Banks								
Royal Bank Internal	3,527	38.8	1,368	6.2	8,501	2,159	5.2	11,269
Global Liquidity Providers	942	68.9	649	4.0	2,604	293	4.0	1,172
Regional Liquidity Providers	315	43.5	137	5.8	795	178	5.3	938
Customer-Service Providers	1,325	47.4	628	4.6	2,888	697	3.4	2,349
Total	**9.859**	**40.1**	**3,950**	**5.3**	**20,747**	**5,909**	**4.5**	**26,794**

Table 4. Price Impact of Price-Contingent Orders: Linear Estimates

This table summarizes the results of the regressions testing for a difference in the price impact of the orders from different client types:

$$(s_{t+k} - s_t)*10000 = \sum_{i=1}^{8} \beta_i * V_{i,t}^{SL} + \sum_{i=1}^{8} \delta_i * \boldsymbol{V}_{i,t}^{TP} + \varepsilon_t ,$$

where s_t is the log exchange rate expressed as foreign currency units per USD. k is the time horizon over which we measure price impact. $V_{i,t}^{SL}$ is the signed dollar value of stop-loss order of customer of the type i; $V_{i,t}^{TP}$ is the signed dollar value of take-profit order of customer of the type i. Size is positive (negative) if the client instructs the dealer to buy the commodity (denominator) currency. Coefficients can be interpreted as the impact of $1 million in basis points. * 10 percent significance; ** 5 percent significance; *** 1 percent significance.

	5 min	30 min	1 hour	6 hours	12 hours	1 day	1 week	2 weeks
Stop-Loss Orders								
End-Users								
Leveraged Investors	0.058	0.173**	0.163*	0.486***	0.269	0.229	0.166	0.292
Institutional Investors	-0.055	-0.236	-0.392	-0.249	-0.533**	-0.275	-0.219	-0.739
Broker-Dealers	-0.043	0.096	0.032	0.249	0.316	0.030	0.544	1.576
Central Banks, Gov’ts	-0.037	0.152	0.368	0.615	1.187	1.262	0.264	-0.392
Large Corporations	-0.024	-0.011	-0.003	0.049	0.033	0.078	0.076	0.186**
Middle-Market Corp.	-0.023	0.138	0.151	-0.254	-0.075	0.137	1.424**	3.059**
Banks								
Royal Bank Internal	0.018	0.073*	0.159***	0.335***	0.348***	0.500***	0.323*	0.745**
Global	0.108***	0.309***	0.440***	0.935***	0.842***	1.095**	0.795	1.109
Regional	0.207***	0.489***	0.832***	2.079***	2.063***	1.585**	2.316***	2.298*
Customer-Service	0.026	0.117*	0.287**	0.605***	0.734***	0.720**	0.492	1.332**
Take-Profit Orders								
End-Users								
Leveraged Investors	0.014	0.004	-0.017	-0.075	-0.098	-0.686**	-0.341	-0.933
Institutional Investors	-0.030	-0.031	0.076	0.219	0.165	-0.329	-0.444	-0.451
Broker-Dealers	0.072	-0.058	-0.031	-0.431	-0.689	-0.432	0.038	-0.395
Central Banks, Gov’ts	0.036*	-0.201***	-0.231***	-0.085	-0.372**	-0.010	-0.235	0.013
Large Corp.	0.008	0.004	-0.029	-0.018	0.189	0.333**	0.287	0.301
Middle-Market Corp.	-0.134**	-0.217*	-0.211	-0.182	-0.167	-0.304	-0.318	-1.404
Banks								
Royal Bank Internal	0.009	-0.011	-0.036	-0.072	-0.089	0.069	0.336	-0.129
Global	0.053	0.179	0.221	0.614*	0.744*	0.588	-0.039	-1.044
Regional	-0.014	-0.324**	-0.121	0.085	-0.120	-0.019	0.397	2.155**
Customer-Service	-0.050	-0.044	-0.070	-0.092	-0.046	-0.037	0.608	-0.487
Prob > F	0.0019	0.0001	0	0	0	0.0007	0.0368	0.0074
R-squared	0.0032	0.0064	0.0085	0.0088	0.0062	0.0046	0.0026	0.0028

Table 5. Price Impact of Price-Contingent Orders: All Currencies, Log-Linear Estimates

This table summarizes the results of the regressions testing for a difference in the price impact of the orders from different client types:

$$(s_{t+k} - s_t) * 10000 = \sum_{i=1}^{8} \beta_i * \log(\boldsymbol{V}_{i,t}^{SL}) + \sum_{i=1}^{8} \delta_i * \log(\boldsymbol{V}_{i,t}^{TP}) + \varepsilon_t \,,$$

where s_t is the log exchange rate expressed as foreign currency units per USD. k is the time horizon over which we measure price impact. We take the log of the order size and then reassign the buy-sell direction. The resulting variable is $log(V_{i,t}^{x})$, where $x \in (SL,TP)$. It is positive (negative) if the client instructs the dealer to buy the commodity (denominator) currency. * 10 percent significance; ** 5 percent significance; *** 1 percent significance.

	5 min	30 min	1 hour	6 hours	12 hours	1 day	1 week	2 weeks
Stop-Loss Orders								
End-Users								
Leveraged Investors	0.399	1.849***	1.909***	4.306***	3.184*	3.295	2.638	2.249
Institutional Investors	-0.096	-0.520	-1.281	-1.138	-2.492	-0.485	0.631	2.312
Broker-Dealers	-0.041	0.562	0.637	1.402	1.390	0.417	2.747	7.810*
Central Banks, Gov'ts	0.090	1.152	2.201	3.928	6.069	5.301	0.440	-1.468
Large Corporations	0.023	0.365	0.444	-0.149	-1.672	-1.371	-2.742	-1.474
Middle-Market Corp.	-0.422	0.582	0.967	-0.963	0.682	-0.131	5.051	5.873
Banks								
Royal Bank Internal	0.156	0.644***	1.229***	2.428***	3.009***	3.313***	2.237*	6.094***
Global	0.613***	1.606***	2.268***	4.099***	3.725***	4.679**	3.003	5.680
Regional	0.209	0.924**	2.021***	3.896***	4.317***	4.599***	4.214**	8.514**
Customer-Service	0.894***	2.185***	3.925***	8.686***	8.834***	5.554*	9.109**	8.284
Take-Profit Orders								
End-Users								
Leveraged Investors	0.111	-0.050	-0.347	-1.060	-1.257	-4.335**	-3.657*	-6.855*
Institutional Investors	-0.312	-0.439	0.051	1.170	0.367	-2.918	-3.282	-1.331
Broker-Dealers	0.309	-0.451	-0.204	-0.703	-1.466	-0.481	1.236	2.032
Central Banks, Gov'ts	0.125	-1.402**	-1.995**	-0.824	-3.812*	-1.400	-3.824	2.243
Large Corp.	0.155	0.025	-0.028	0.685	1.322	1.071	1.123	3.717
Middle-Market Corp.	-0.103	0.069	0.478	1.400	1.538	1.836	1.448	0.319
Banks								
Royal Bank Internal	0.000	-0.004	-0.162	-0.494	-0.451	0.739	1.874	-1.503
Global	0.311	0.643	0.638	2.379	3.375	3.156	-0.365	-3.964
Regional	-0.265	-0.157	-0.331	-0.051	0.367	0.635	3.168	-2.523
Customer-Service	0.027	-0.659	0.017	0.789	-0.804	0.049	0.283	8.886
Prob > F	0.0218	0	0	0	0	0.0007	0.065	0.0491
R-squared	0.0031	0.0071	0.0105	0.0093	0.0072	0.0047	0.0029	0.0031

Figure 1 A. Distribution of placed and executed orders over the day, USD/EURO

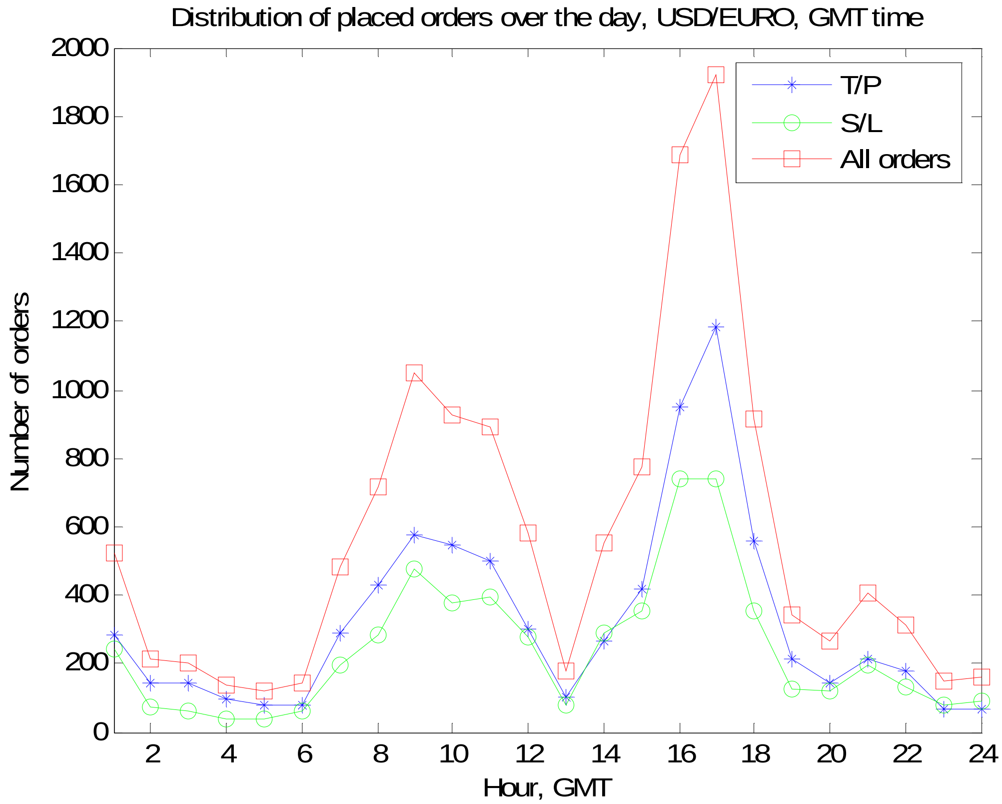

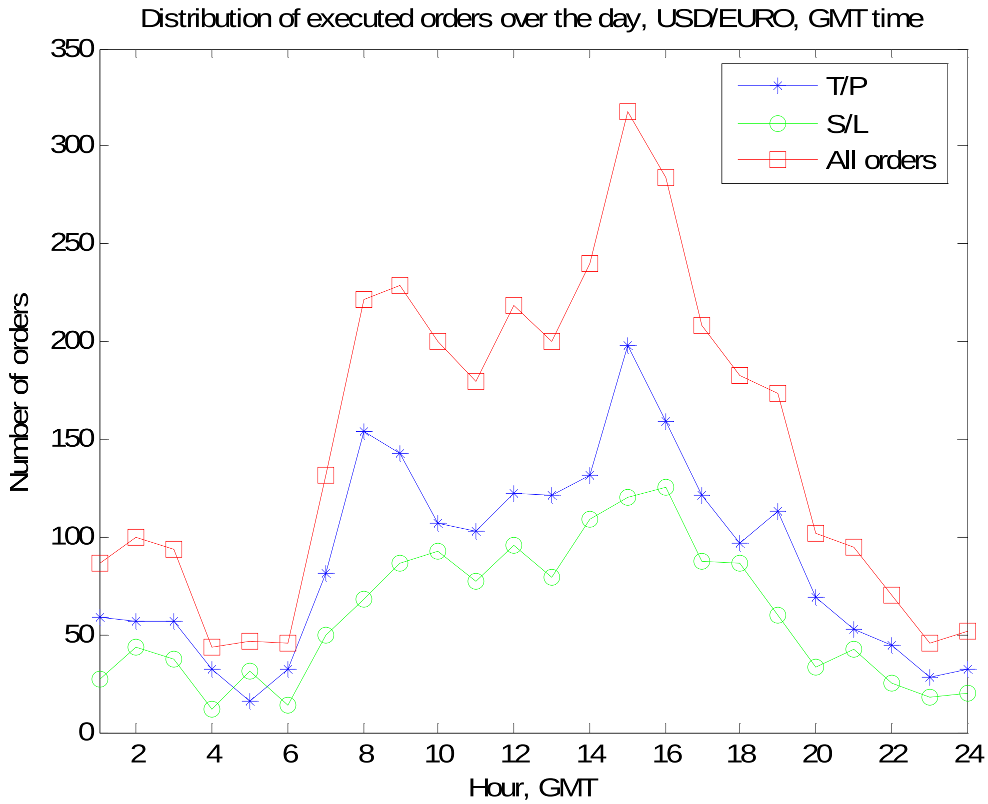

Figure 1 B. Distribution of placed and executed orders over the day, USD/GBP

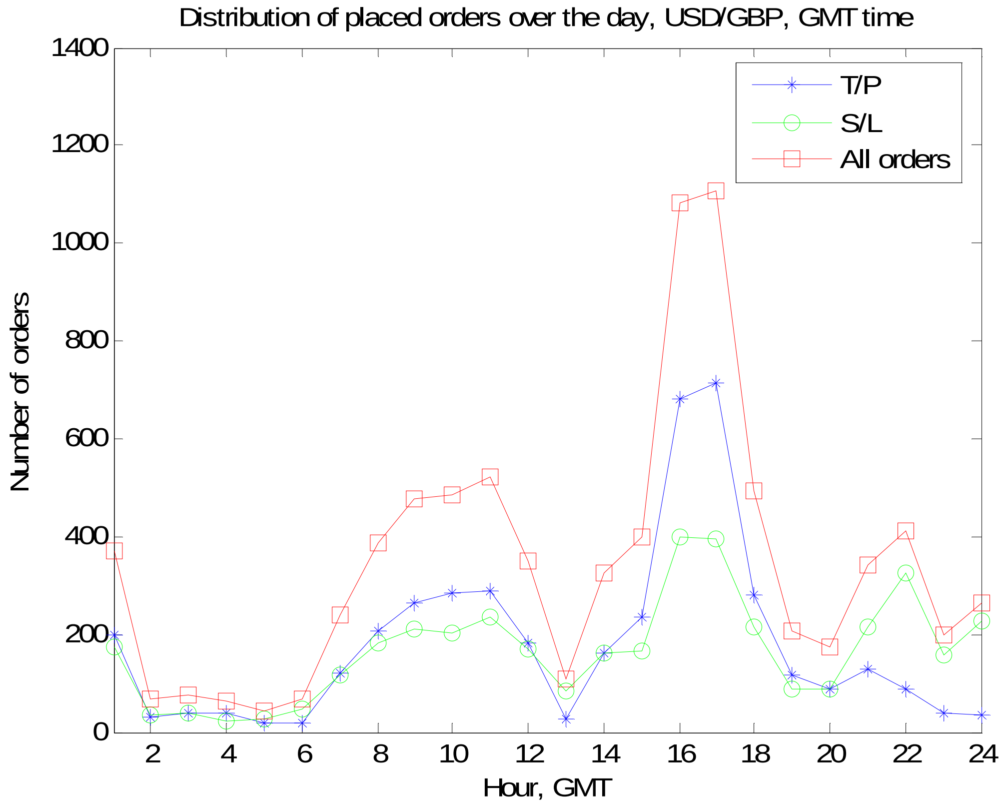

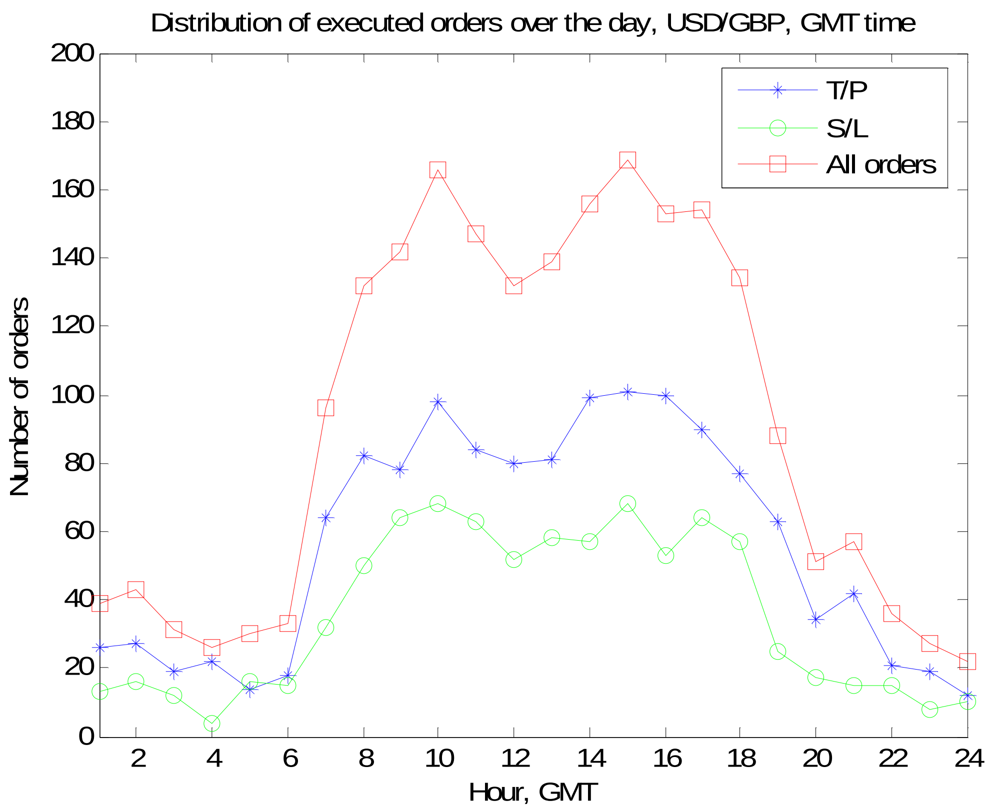

Figure 1 C. Distribution of placed and executed orders over the day, JPY/USD

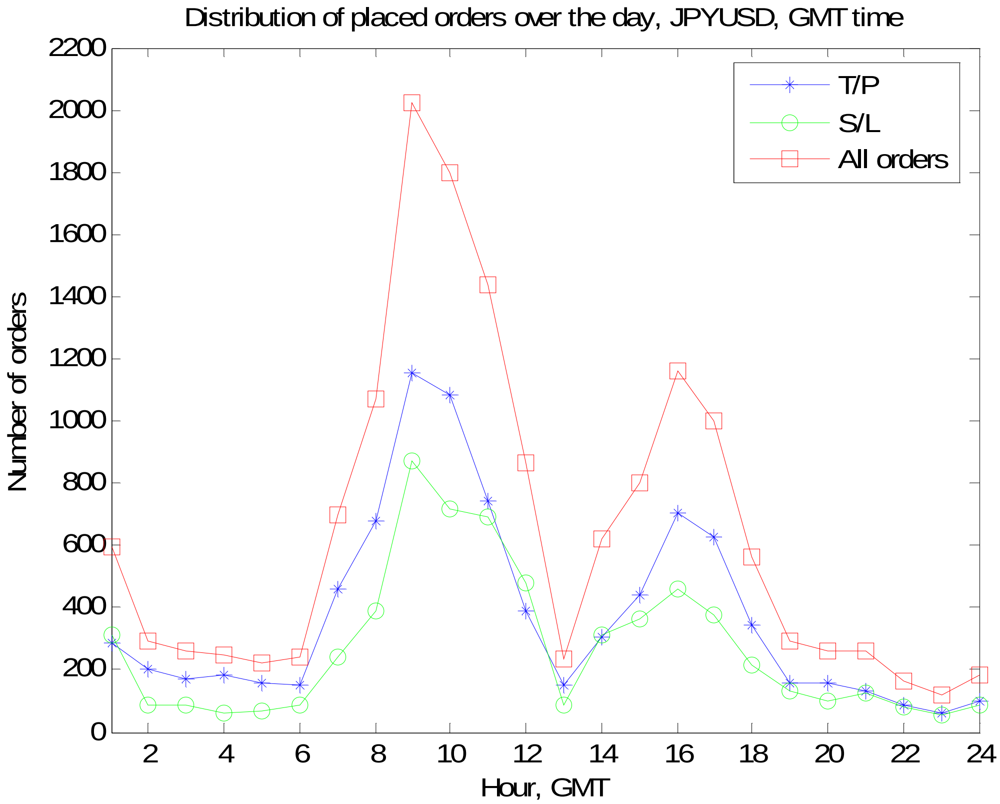

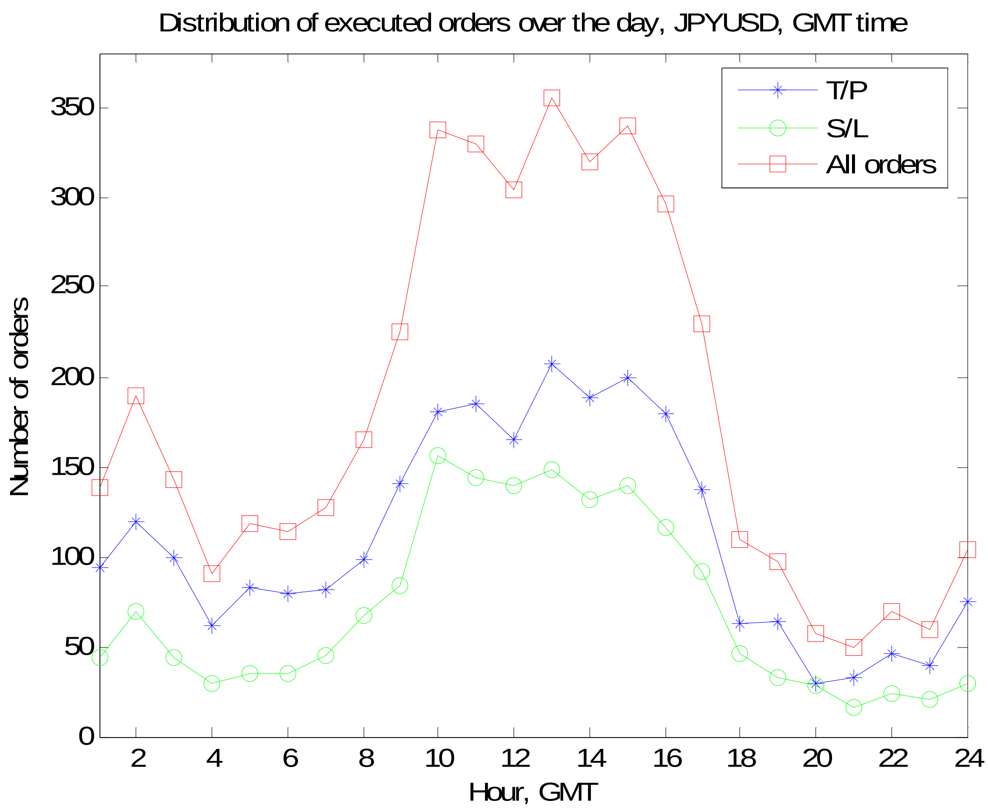

Figure 2. Intraday pattern of exchange rate volatility

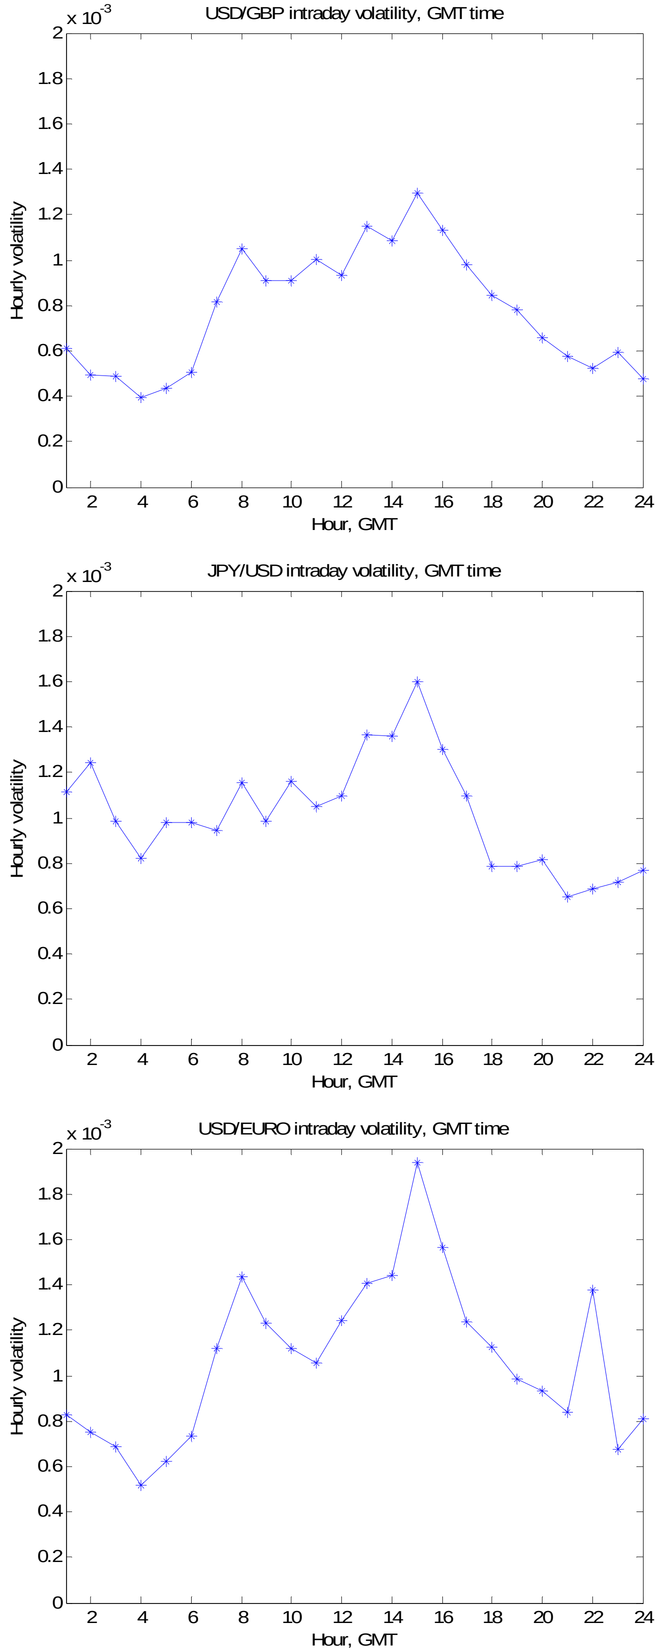

Figure 3: Price impact of informed orders

Figure shows estimated coefficients from regressions of the following equation:

$(s_{t+k} - s_t)*10000 = \sum_{i=1}^{8} \beta_i * V_{i,t}^{SL} + \sum_{i=1}^{8} \delta_i * \boldsymbol{V_{i,t}^{TP}} + \varepsilon_t$, where s_t is the log exchange rate expressed as foreign currency units per USD. k is the time horizon over which we measure price impact. $V_{i,t}^{SL}$ is the signed dollar value of stop-loss order of customer of the type i; $V_{i,t}^{TP}$ is the signed dollar value of take-profit order of customer of the type i.

4A: Price Impact of "Informed" Stop-Loss Orders, Linear Estimates

Based on point estimates only.

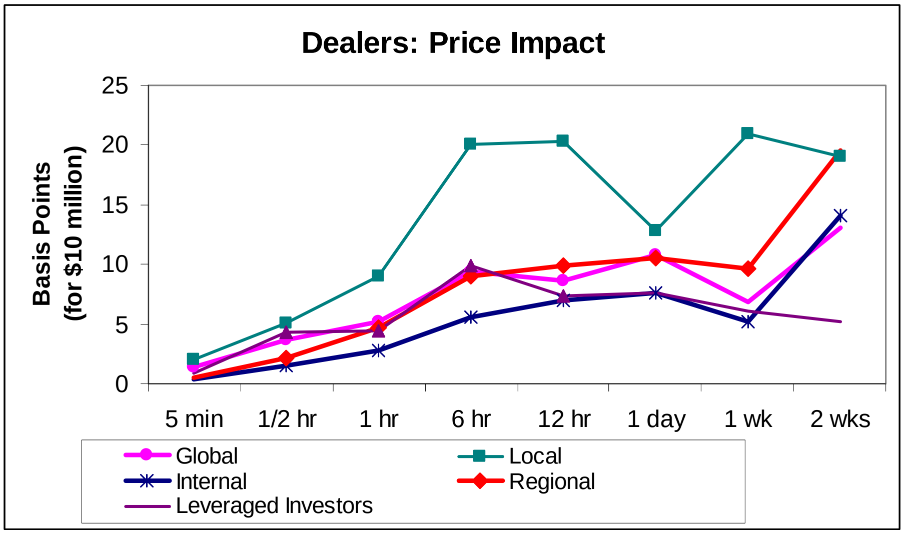

4B: Price Impact of Leveraged Investor Stop-Loss Orders, Linear vs. Concave Estimates