PREDICTABLE ORDER FLOW AND EXCHANGE RATE DYNAMICS

C. L. Osler

Abstract

Recent empirical research shows that order flow is a critical determinant of high-frequency exchange rate movements. The two commonly-cited potential sources for this connection are information and inventory effects. This paper looks for the presence of inventory effects by estimating how exchange rates respond to predictable order flow.

Stop-loss and take-profit orders, which are common in currency markets, cluster in predictable ways at round numbers (Osler (2001)), so the news conveyed by such clusters should be zero on average. If information effects are the only ones that matter, exchange rate behavior at round numbers should be no different from behavior elsewhere, on average. If inventory effects matter, the clustering of orders suggests three testable hypotheses: (1) trends should be reversed when rates reach round numbers; (2) trends should be intensified once rates cross round numbers; (3) these special behaviors should be temporary.

These hypotheses are tested with the bootstrap technique on minute-by-minute data for dollar-mark, dollar-yen, and dollar U.K. pound over January 1996 through April 1998. Tests consistently support the presence of inventory effects. The paper supports its attribution of these special behaviors to order clustering by noting that the two alternative explanations for such behavior suggested in the literature are not supported empirically. It also notes evidence that stock prices are affected by related patterns of order clustering. (Key words: Order Flow, Information, Exchange Rates, High-frequency, Path Dependence.) (JEL numbers: F1, G3.)

May, 2001

The views expressed in the paper are those of the author and are do not necessarily reflect views at the Federal Reserve Bank of New York or the Federal Reserve System. The author thanks Priya Gandhi for excellent research assistance. Any errors or omissions are the responsibility of the author.

PREDICTABLE ORDER FLOW AND EXCHANGE RATES

Recent empirical research has shown that order flow is a critical determinant of high-frequency exchange rate movements: a currency's value tends to rise when buy orders dominate order flow, and tends to fall when sell orders dominate (Evans and Lyons (2001); Rime (2000); Lyons (2001); Evans (2001)). When examining the connection between order flow and financial prices in theoretical detail, researchers' attention has often focused on the way order flow can signal private information (Kyle (1985); Lyons (1995); Evans and Lyons (2001); Rime (2000); Lyons (2001); Evans (2001)). Inventory effects have also received extensive theoretical and empirical attention (Garman (1976), Amihud and Mendelson (1980), Ho and Stoll (1981), Lyons (1995), Yao (1998), Bjonnes and Rime (2000), Evans and Lyons (2001), Evans (2001)).

This paper provides evidence that the connection between order flow and exchange rates is not exclusively based on information effects, by showing that exchange rates seem to react to predictable order flow. The paper also provides evidence consistent with a role for inventories, by showing that the apparent effects of this predictable order flow are consistent, in direction and duration, with the predictions of inventory models.

The predictable, non-informative order flow examined here comprises clusters of stop-loss and take-profit orders. A take-profit order instructs the dealer to sell (buy) if the exchange rate rises (falls) to a particular level. A stop-loss order instructs the dealer to buy (sell) if the exchange rate rises (falls) to a particular level.[1] Osler (2001) shows that these orders are not uniformly distributed across the spectrum of possible exchange rates. Instead, their requested execution prices are strongly clustered at round numbers.

Further, these distributions differ across order types. Take-profit orders tend to be strongly clustered exactly at round numbers, such as \$1.6600/£ or ¥123.00/\$; stop-loss buy orders are strongly clustered just above round numbers, at rates such as \$1.6605/£ or \$1.6610/£; stop-loss sell orders are strongly clustered just below round numbers.

The main hypothesis of this paper is based on the observation, made previously by Hasbrouck (1988), that orders should only provides private, market-relevant information when they differ from the expected amount conditional on existing information. These clustering patterns are well-known to market participant, since they create the patterns themselves. If the rate reaches a round number such as \$1.6600/£, market participants should expect many conditional orders; if it reaches a non-round number such as \$1.6641/£ they should expect just a few. Under the joint hypothesis that expectations are unbiased, so that expectational errors average close to zero, and that information is the only source of influence from order flow to exchange rates, the average effect of order clusters should be zero.

Alternatively, inventory theoretic models imply three predictions about exchange rate behavior at round numbers. First, exchange rates should reverse course relatively frequently upon reaching such numbers, since the clusters of take-profit orders at the round numbers should act as partially reflecting barriers. Second, exchange rate trends should generally intensify once such round numbers are crossed, since the clusters of stop-loss orders just beyond round numbers should intensify trends. Third, the trend-reflecting and trend-intensifying tendencies at round numbers should be temporary (Kraus and Stoll (1972); Hasbrouck (1988)). The results presented here indicate that order clusters do affect exchange rates, consistent with these three predictions. This suggests that information effects are not the sole source of influence from order flow to exchange rates, and that inventory effects also matter.

The results that predictable order flow has only temporary effects on exchange rates differs from existing evidence about effects of aggregate order flow, which have been documented to last multiple days (Evans and Lyons (1999); Rime (2000); Evans (2001)). This would be expected, since aggregate order flow includes informative orders and the effects of non-orders information should be permanent.

The analysis here cannot provide a complete proof that order clustering is the source of the special behaviors of exchange rates around round numbers. To evaluate the possibility that some other factor might be responsible, the paper closely examines both of the other possible reasons for such behavior suggested by the literature: central bank intervention and chaotic underlying exchange rate processes. Strong reasons are provided for rejecting both potential explanations.

The idea that order clustering explains the special behaviors of exchange rates near round numbers also gains credibility from the fact that related results have been found for prices of firms listed on the New York Stock Exchange. Prices for equity limit orders have also tended to cluster, in a pattern that closely parallels the clustering examined here (Osborne (1962); Neiderhoffer (1965), Neiderhoffer and Osborne (1966); Harris (1991); Cooney et al. (2000)). Limit orders should cause rates to pause at their specific rate or turn back, a tendency which has been documented (Neiderhoffer and Osborne (1966)).[2]

The paper suggests that the definition of market-relevant information can be dichotomized into order flow information per se, on the one hand, and information about other factors relevant to exchange rates, such as economic fundamentals or non-fundamentals-based trading behavior, that can be partially communicated through order flow.

The statistical analysis relies on minute-by-minute exchange rate quotes for dollar-mark, dollar-yen, and dollar-U.K. pound covering January, 1996 through April, 1998. The statistical methodology is a variant of the bootstrap (Efron 1979, 1982), in which I compare the exchange rates' behavior at round numbers with their behavior at large sets of arbitrary numbers, under the null hypothesis that there is no difference.

The paper is organized as follows. Section I puts this research in the context of the existing literature, and discusses some of the concepts of "information" involved. Section II discusses the data and methodology. Section III presents the results. Section IV discusses alternative possible sources for the exchange-rate behavior documented here. Section V offers concluding remarks.

II. EXISTING LITERATURE

This section first provides a useful dichotomy between two types of information processed by dealers. It then reviews existing research on the influence of order flow on exchange rate changes, focusing on the information and inventory effects. It also reviews existing evidence on the tendency for take-profit and stop-loss orders to cluster at round numbers, an observation that is critical to the claim that such clusters are predictable.

Before delving into these matters, it is helpful to divide information into two categories. The first includes market-relevant news that can be communicated *through* order flow, but which does not pertain to order flow itself. This category, labeled I_N for brevity, includes news about fundamental factors such as economic growth, expected inflation, trade and investment flows, expected asset returns, and risk tolerance. News about technical trading signals would also be included in this category, since technical trading is so widespread.

The second category of information is news about order flow per se—the news that $5 million was or will be sold against yen at ¥1.2323/$, for example. Within this category it is important to distinguish order flow information that conveys I_N, I_{ON}, and other order flow information. The second subcategory will be denoted I_{O*}.

If order flow only affects exchange rates because it is a conduit for I_N, then I_O should have no independent effect on exchange rates. Conversely, if I_{O*} does affect exchange rates, then order flow must have an effect on exchange rates independent of information.

This dichotomy can be related to the dichotomy between fundamental and semi-fundamental private information:

<table>
<tr><td>Fundamental</td><td colspan="3">Semi-Fundamental</td></tr>
<tr><td></td><td>Trading Signals</td><td>Order Flow News That Conveys I_N</td><td>Other Order Flow News</td></tr>
<tr><td colspan="2"></td><td>I_{ON}</td><td>I_{O*}</td></tr>
<tr><td colspan="2">I_N</td><td colspan="2">Order Flow Information</td></tr>
</table>

Evans and Lyons (2001) and Evans (2001) provide models in which pure order flow information, I_{O*}, affects prices because inventory effects matter.

I.A. Kyle-Type Information Effects

Most of the information relevant to exchange rate determination is dispersed among market participants. For example, each goods trading firm knows its own imports or exports, but the aggregate picture of current international trade is not known by anyone. Given the shortage of public information about current exchange rate fundamentals and noise trading, it is rational for currency dealers to try to infer such private information from order flow. In Kyle (1985), and many subsequent microstructure papers, "informed traders" have a private signal about the true value of a given stock. A central market-maker must distinguish between informed and uninformed traders, while simultaneously attempting to learn the noise traders' private signal. The difficulty of this learning task is enhanced in currency markets by the multiplicity of factors relevant to exchange rate behavior, by the large number of currency dealers, and by the physical dispersion of deals across time.

Empirical evidence linking currency order flow directly to the communication of private information is scarce. Lyons (1995) shows that dealers' shade their prices more as order size rises, which he interprets as an information effect. Bjonnes and Rime (2000) show that bid/ask spreads widen with order size, and interpret this as an information effect. However, these observations are also predicted by inventory-theoretic models such as Ho and Stoll (1981).

Lyons (2001) shows that orders from financial institutions have a greater impact on exchange rates than orders from other sources. Likewise, Bjonnes and Rime (2000) find that orders from Norwegian customers have a greater effect on Krone exchange rates than orders from non-Norweigian sources, and that orders from the Norwegian central bank have an even greater effect on those exchange rates. It is possible to interpret these results in terms of the information content of orders from the sources in question. Since financial institutions have a closer focus on holding-period returns than agents such as importing and exporting firms, orders from financial firms should carry more concentrated information about likely future exchange rates than orders from other firms.[3] Likewise, residents of Norway could reasonably be assumed to know more about the factors affecting the Krone than foreigners, and the Norwegian central bank is likely to be the best informed of all.

I.B. Inventories and Demand Curves

Another possible connection from orders to exchange rates is based on inventory effects. In early models of inventory effects, dealers choose an optimum inventory level (Garman (1976), Amihud and Mendelson (1980), Ho and Stoll (1981)), balancing the risk involved in holding inventories of financial assets whose prices vary at high frequencies against the costs of eliminating inventories.[4] Dealers in these models face a downward-sloping demand curve for stocks, derived from carefully aggregating agents with stochastic, ad hoc underlying fundamental demands. If a dealer receives an unanticipated customer sale (purchase) and ends up with a long (short) position, she will quote lower (higher) bid as well as offer rates. This connection from inventories to prices holds if dealers are risk averse (Ho and Stoll (1981)), if they are profit-maximizers facing the gambler's ruin problem (Garman (1976)), or if they are profit maximizers facing position limits (Amihud and Mendelson (1980)).

In reality, most currency dealers are free to choose a non-zero end-of-day inventory position within specified limits, but they typically prefer a zero inventory position and eliminate their excess inventories quickly (Lyons (1995), Bjonnes and Rime (2000)). Evans and Lyons (2001) presents a model of the foreign exchange market in which dealers inventories are exogenously required to reach zero at the end of the trading period, but (some) customer demand is derived from utility maximization. Inventory constraints are also exogenous in Evans (2001). The price effects of inventories in these models are consistent with those sketched above.[5]

Empirical support for the prediction that dealers shade their prices in response to unwanted inventory accumulation comes from Lyons (1995), who examines the dealing behavior of one currency dealer during a week in 1992. He finds clear evidence that the dealer shaded his outgoing prices in response to his inventory levels.[6] Bjonnes and Rime (2000) suggest that such price shading may have become less frequent with the widespread use of electronic broking systems, partly because shading prices provides private inventory information to other dealers. In their examination of four dealers over one week in 1998, the authors find no statistically significant effect of inventory on price quotes. Instead, Bjonnes and Rime find that dealers tend to dispose of unwanted inventory via outgoing contacts on electronic brokers.

If dealers have shifted away from shading prices in response to unwanted inventory accumulation, and towards the use of electronic brokers, the mechanics of the connection from inventories to prices may also have changed. Goodhart et al. (1996) find that one factor driving quote revisions on the Reuters D2000-2 electronic broking system is the direction and/or magnitude of previous trades.[7] This behavior could be rational, in part because trades sometimes come in waves of purchases and sales (Goodhart et al. (1996); see also Hasbrouck (1991)), so that a dealer's conditional expectation of the direction of the next trade should depend positively on the direction of the most recent trade.[8] The existence of such behavior, together with the results of Bjonnes and Rime (2000), suggests the following connection from inventory to exchange rates: A dealer ends up with unwanted inventory due to some incoming customer trade. To eliminate the unwanted inventory, the dealer makes an outgoing trade, taking the bid. Other dealers observe the trade at the bid and lower their prices.

II. DATA

This section first reviews the orders and exchange rate data used for the statistical tests. It then describes the general approach to testing the null hypothesis that predictable order flow, represented by conditional order clusters, has no effect on exchange rates. Finally, it presents the statistical methodology in detail.

II.A. Orders Data

The orders data underlying the statistical tests used here represent all take-profit and stop-loss orders ("conditional" orders, for brevity) processed by a major currency dealing bank during September, 1999 through April 11, 2000. The data-set includes approximately 9,700 orders in three currency pairs: dollar-yen, dollar-U.K. pound, and euro-dollar. The bank in question informally estimates that these orders account for roughly 5 percent of its total (absolute) order flow. However, many conditional orders never appear on any order book, so the true fraction of conditional orders in all order flow is likely to be higher.[9]

A take-profit order instructs a dealer to buy (sell) a certain amount at the market rate once the exchange rate falls (rises) to a certain level, which will be called the "requested execution rate." A stop-loss order instructs a dealer to buy (sell) a certain amount at the market rate once the exchange rate rises (falls) to a certain level. Though the use of these orders could be inefficient in a perfect market with fully rational players (Dybvig (1988)), the foreign exchange market does not fit this description. The foreign exchange market has non-frictionless trading, asymmetric information, an excess of information relative to the amount that can be absorbed by any single individual (because it operates around the clock), and an excess of information from minute to minute relative to the amount that can be efficiently absorbed by a non-dealer. Furthermore, many agents in the foreign exchange market regard themselves as imperfectly rational, and use stop-loss and take-profit orders as a method for managing self-perceived tendencies to time inconsistency.[10]

Osler (2001) shows that the requested execution rates for stop-loss and take-profit orders have a pronounced tendency to concentrate at round numbers. Over 10 percent of all orders are placed at rates ending in 00 (such as ¥123.00/$ or $1.4300/£); about 3 percent of orders are placed at each of the other rates ending in 0 (such as ¥123.20/$ or $1.4370/£); about 2 percent of orders are placed at rates ending in 5.

Figure 1 disaggregates orders according to type (take-profit or stop-loss) and direction (buy or sell), and Table 1 summarizes the information relevant to our analysis. Two critical asymmetries are apparent. First, take-profit orders have a more pronounced tendency to cluster at rates ending in 00 than stop-loss orders. About 9.3 percent of take-profit orders (weighted by value) are placed exactly at such rates; the corresponding figure for stop-loss orders is 4.4 percent. Second, stop-loss buy orders tend to cluster at rates just above those ending in 00 or 50, and stop-loss sell orders tend to cluster at rates just below such round numbers. For example, over 14 percent of stop-loss buy orders are placed at rates ending in the range [01,10], while only 7.4 percent of those orders are placed at rates ending in the range [90,99]. Osler (2001) shows that these asymmetries are statistically significant.

As noted by Hasbrouck (1988), only the unexpected component of the clusters should convey news. Active market participants like dealers (who execute the orders), currency salespersons (who advise customers on placing their orders), and customers (who place the orders), are all familiar with the clustering tendencies of conditional orders, so their expectation errors should average close to zero. The average effect on exchange rates of conditional order clusters should also be zero, if order flow only affects prices when it communicates news. Alternatively, if conditional order clusters affect exchange rates, on average, then there must be something about the orders besides news that makes dealers change their prices. In short, Predictable order flow affects exchange rates if and only if information is not the only source of influence from order flow to exchange rates. This insight forms the basis of the tests reported below.0

II.B. Exchange Rate Data

I test whether exchange rates behave unusually at round numbers using minute-by-minute exchange rate quotes over January, 1996 through April, 1998 covering the New York trading hours of 9 a.m. to 4 p.m. for three currency pairs: dollar-mark, dollar-yen, and dollar-U.K. pound. The quote for a given minute was taken to be the quote posted at or most recently before the exact beginning of the minute. Round numbers are defined as exchange rates with final two digits equal to 00 or 50. These include exchange rates such as DM1.4300/$, ¥111.00/$, and $1.6500/£.

Quote data are not ideal, since they are technically only indicative (though many banks only post quotes that can be interpreted as firm), and they do not behave identically to actual transactions prices. However, the available transactions price series are at most only four months long, which is time insufficient to permit reliable hypothesis testing in this context. Further, the divergences of quote levels from transactions prices are generally small: indeed, the first attempt to characterize the difference concluded that "the time path of the indicative quotes can ... be taken as a very good an close proxy for that in the underlying firm series" (Goodhart et al. (1996) p. 110). More recent investigations conclude that the quote "returns are consistently more volatile than their [transaction] counterparts and display strong moving average effects which are not present in the [transaction] returns" (Danielsson and Payne (1999) p. 1). To insulate the results from problems associated with such discrepancies, I compare quote behavior at round numbers with quote behavior at arbitrary numbers.

III. STATISTICAL METHODOLOGY

Two sets of empirical tests are run below. In the first, I examine the direction and magnitude of the exchange rate's movements during the fifteen minutes after reaching or crossing a round number. In the second, I examine the longevity of the effects.

III.A. Exchange-Rate Behavior Immediately After Reaching Round Numbers

The null hypothesis for the first set of tests is that order flow only matters to the extent that it communicates news, in which case clusters of conditional orders should have no effect on exchange rates, on average. I test this by investigating whether exchange rate behavior in the first fifteen minutes after reaching or crossing a round number differs from behavior at arbitrary levels. I examine two alternative hypotheses, both derived from inventory models.

Alternative Hypothesis 1: Exchange rates trends are reversed at round numbers more frequently than at other numbers. This alternative hypothesis supposes that clusters of take-profit orders at round numbers reverse exchange rate trends, as predicted by inventory models. To evaluate the null against this alternative, I find every instance that the exchange rate arrives at a round number, defining "arrival" as coming within 0.01 percent. I then define the exchange rate as having reversed course after rising (falling) to a round number if it is below (above) the number 15 minutes later. The proportion of times the rate reverses course is called the "reversal frequency."

Finally, I calculate the corresponding reversal frequency for a large number of exchange rates chosen arbitrarily. If this is "usually" higher for the round numbers than for the arbitrary numbers, in a rigorous statistical sense described below, I conclude that round numbers tend to reverse trends, on average.

Alternative Hypothesis 2: Exchange rates typically trend more rapidly after crossing round numbers than after crossing other numbers. This alternative hypothesis supposes that clusters of stop-loss orders just beyond round numbers tend to intensify trends, as would be predicted by inventory models. To evaluate the null against this alternative, I calculate the average exchange rate move during the 15 minutes following its arrival at a round number, conditional on the rate failing to reverse course. Movements are signed so that a larger positive number means a faster movement in the direction consistent with the alternative hypothesis. If the exchange rate reaches a particular number from above (below), the subsequent move is measured as $s_t - s_{t+15}$ ($s_{t+15} - s_t$), where s_t represents the log of the exchange rate. If the average signed exchange rate movement after reaching a round number usually exceeds the corresponding average for arbitrary numbers, I conclude that trends tend to intensify are round numbers have been crossed.

II.B Longevity of Special Behaviors at Round Numbers

To anticipate, the results of the first set of tests indicate that rates tend to bounce when they arrive at round numbers, and that if they cross the round number they tend to trend rapidly. These results are consistent with the hypothesis that information is not the only source of influence from order flow to exchange rates. They are also consistent with the hypothesis that inventory effects matter.

To examine a possible role for inventories more closely, I examine whether the special behaviors at round numbers are temporary or permanent, because all sources from Garman (1976) through Evans (2001) agree that inventory effects should be temporary. I repeat the tests summarized above at four more time horizons: 30 minutes, 1 hour, 2 hours, and 24 hours.

III.C. The Bootstrap

The use of arbitrary numbers as a benchmark to examine behavior at round numbers is an application of the bootstrap technique (Efron (1979), (1982)). Since the distribution of exchange rate returns is not known to fit any parametric distribution, the bootstrap has the advantage of being agnostic about the correct statistical distribution for hypothesis testing.[11]

In this application of the bootstrap technique, I create an approximation to the correct statistical distribution for hypothesis testing based on the actual distribution of exchange rates. My approximation to the behavior of exchange rates at round numbers, under the null hypothesis that round numbers are not special, is the behavior of exchange rates at arbitrary numbers. To generate this approximating distribution, I first choose a very large number of random exchange rates: 10,000 sets of 30 arbitrary rates. I then evaluate the behavior of exchange rates around these arbitrary levels. If behavior at round numbers is sufficiently different from behavior at arbitrary numbers, by common statistical standards, I conclude that behavior at round numbers is special.[12]

The simplest possible bootstrap test would compare the exchange rate's average behavior at round numbers with the distribution of behaviors at the 10,000 sets of arbitrary levels, each evaluated over the entire sample period. However, this test would be statistically unreliable because the statistics to be examined are all in the form of averages. Under the central limit theorem, the variance of each associated distribution is sensitive to the number of times the exchange rate hits the specified levels. It proved infeasible to control the number of hits for arbitrary numbers sufficiently closely to be sure that the variance of the approximating distribution was appropriate. For this reason I apply an alternative approach that uses only first moments.

I divide the sample period into 58 intervals of 10 consecutive trading days. For each interval I compare the average reversal frequency and post-crossing move for round numbers and arbitrary numbers. Under the null hypothesis, the figures for round numbers have a 50 percent chance of exceeding the corresponding figures for arbitrary numbers in each time interval. Thus, under the null hypothesis each 10-day interval can be viewed as a Bernoulli trial with probability ½, and results for the combined set of trials should conform to the binomial distribution with parameters (½,N), where $N \leq 58$ is the number of ten-day intervals in which round numbers were reached at least once.

IV. RESULTS

Results of the first set of tests indicate that exchange rates tend to reverse course when they reach round numbers, and that trends tend to accelerate after round numbers are crossed. This is inconsistent with the null hypothesis that information effects are the only source of influence from order flow to exchange rates, and consistent with the hypothesis that inventory effects matter. Results of the second set of tests show that the special behaviors at round numbers are no longer evident one day later, which is also consistent with the hypothesis that inventory effects matter.

IV.A. Exchange-Rate Behavior Immediately After Reaching Round Numbers

Alternative Hypothesis 1: Exchange rate trends tend to be reflected at round numbers. During the sample period, exchange rates reached round numbers an average of 4,000 times per currency.[13] Fifteen minutes after these events, exchange rates had stopped trending or reversed course 59.0 percent of the time, on average (Table 2). The corresponding figure for arbitrary numbers is 54.8 percent.

For dollar-mark, the reversal frequency for round numbers exceeds the corresponding average frequency for arbitrary numbers in 42 of the 58 relevant ten-day intervals (I denote this pair as: [42,58]); this outcome is statistically significant, since under the null hypothesis it would be observed less than 0.1 percent of the time. For the yen and the pound, the corresponding pairs are [38,57] and [37,58], respectively, both of which are also statistically significant.

Alternative Hypothesis 2: After crossing round numbers, exchange rates tend to trend rapidly. The average exchange-rate move during the 15 minutes after hitting a round number, conditional on a failure to reverse course during that interval, is 0.00058 percent (on average across the three currencies). The corresponding move after hitting any arbitrary number and failing to bounce is 0.00044 percent.

For dollar-mark, the average post-crossing move for round numbers exceed the corresponding figure for arbitrary numbers in 46 of the 58 relevant intervals, which is statistically significant at the 0.001 percent level. Corresponding figures for dollar-yen and dollar-pound are [46,57] and [46,58], which are also highly statistically significant.

III.B. Longevity of Special Behaviors at Round Numbers

The special behaviors associated with round numbers seem to be temporary. Table 3 shows that the exchange rate's behavior at round numbers is indistinguishable from its behavior at arbitrary numbers within a day after reaching the round number, and often within an hour. This is true for both the trend reversal and trend-intensifying behaviors identified above, and for all three currencies, with one minor exception: For the U.K. pound, the trend reversal behavior initially disappears within half an hour, but then becomes statistically significant once again at one day. It would be hard to find an economic explanation for this observation. Overall these results are consistent with the possibility that inventory effects are an important source of the connection between order flow and exchange rates.

IV. DISCUSSION

So far this paper has shown that exchange rates exhibit special behaviors near round numbers, and has interpreted these behaviors as indicating (1) that the effects of order flow on exchange rates is not solely related to information, and (2) that inventory effects may be important. Neither of these implications would be warranted if the special exchange rate behaviors at round numbers are unrelated to the presence of conditional order clusters. To provide further support for the hypothesis that conditional order clusters are the source of these exchange rate behaviors, this section examines two other hypothetical sources of such exchange behavior: central bank intervention and chaotic underlying exchange-rate processes. To my knowledge, these are the only other relevant hypotheses to be gleaned from the literature.

IV.A. Central Bank Intervention

Standard exchange rate models provide one alternative explanation for the special behavior of exchange rates at round numbers: central bank intervention. That these models provide any explanation at all may seem counterintuitive, since they generally assume that only the *log* exchange rate matters for economic outcomes. In effect, this assumes away the quotation conventions and other institutional features required for the concept of a "round number" to be well-defined. Round numbers could only have special relevance within this literature if a central bank tended to intervene frequently at round numbers.

To generate the behavioral tendencies documented in Section III, the intervention at round numbers would have to conform to a special pattern. When the rate arrives at a round number the central bank must initially trade "against the wind." If the rate crosses a round number, the central bank must then trade "with the wind." Such behavior seems unlikely on purely conceptual grounds, since it would fulfill none of the standard motivations for central bank intervention, such as calming disorderly markets by providing short-term liquidity, smoothing the path of the exchange rate, or targeting particular exchange rate levels or bands.

Despite the shaky conceptual foundation for this hypothesis, it is worth looking closely at the empirical evidence on this matter. For dollar-mark and dollar-pound there was no reported intervention at all in these markets during the sample period, according to reports in *The Wall Street Journal* and the *Financial Times*. The Bank of Japan did intervene on a number of occasions during the sample period, and possibility of intervention was occasionally discussed publicly by officials from the Bank of Japan and the Ministry of Finance. I ran the two tests reported in Section III over the subset of 20 months in which there was neither reported intervention nor the public discussion of intervention, according to the Wall Street Journal.[14] This had no effect on the qualitative conclusions presented earlier (Table 4). Indeed, the exchange rates' special behaviors at round numbers are more pronounced in the reduced sample period, rather than less, though the differences are not large. This evidence suggests very strongly that central bank intervention was not the primary source of the exchange-rate behavior identified in Section III.

IV.B. Chaos

Clyde and Osler (1997) show that technical trading signals based on visual patterns, like head-and-shoulders patterns, might derive their forecasting ability from an underlying chaotic structure in the financial price series. Some chaotic processes will form clear patterns, called "attractors," when charted in phase space. If such series are graphed over time certain visible patterns might have predictive power.

There are two strong reasons why this hypothesis is unlikely to explain the behaviors identified in Section III. First, the hypothesis could not apply to signals that depend on rates being at round numbers, since chaotic processes, by their nature, are not constrained by such factors. Second, empirical tests have generally provided little support for the hypothesis that exchange rates are chaotic (Hsieh (1989); Cecen and Erkal (1996)).

V. CONCLUSIONS

Important empirical work of the past few years indicates that order flow is a critical determinant of high-frequency exchange rate movements (Goodhart et al. 1996; Evans and Lyons (1999); Rime (2000); Lyons (2001); Evans (2001)). This paper asks why order flow matters for exchange rates, focusing on the two major hypotheses advanced in the literature: information effects (Kyle 1985) and inventory effects (Garman 1976). Evidence provided here suggests that information effects are not the sole source of influence from order flow to exchange rates, and that inventory effects may also be present.

To help distinguish between information and inventory effects the paper tries to identify the effects of predictable clusters of take-profit and stop-loss orders. As shown in Osler (2001), take-profit orders tend to be clustered exactly at round numbers, while stop-loss orders tend to be clustered just beyond round numbers. If order flow only affects exchange rates because it communicates news, and if market participants have reasonably accurate expectations about order clusters, then the average effect of these order clusters should be zero. More specifically, exchange rates should behave the same near round numbers as elsewhere. By contrast, if inventory effects matter, then exchange rate trends should tend to be reflected at round numbers, and to be intensified once rates have crossed round numbers. Further, these special behaviors at round numbers should be temporary. The predictions of inventory models are supported throughout.

The results of this paper have additional implications for our understanding of exchange rates and for the definition of market-relevant information. First, the results provide further confirmation that order flow is a strong influence, possibly the dominant influence, on high-frequency exchange rate movements. Second, the paper highlights the significance of an institutional feature of the market that is typically ignored in exchange rate research, the quotation convention. The concept of a round number that might matter for exchange rates would not be well-defined in the absence of a consistent quotation convention in the dominant, wholesale market. Such quotation conventions certainly exist—for example, the exchange rate between the dollar and the pound is universally quoted as dollars per pound among dealers—and can easily be rationalized in terms of the fast pace of the market. Nonetheless, the sensitivity of exchange-rate behavior to quotation conventions is inconsistent with theoretical exchange-rate models, where exchange rate behavior is universally assumed to be invariant to institutional features of the market.

Third, the results indicate that order flow need not be monotonically related to the flow of private information into the market. Since order clusters can be at least partially predicted by market participants, based on past order clustering patterns around round numbers, surges of order flow can be expected at high frequencies that are unconnected with corresponding surges in private information. This contrasts with the view of the stock market developed formally in Diamond and Verrecchia (1987) and Easley and

O'Hara (1992), and the view of currency markets presented informally in Rime (2000), in which higher trading intensity is monotonically related to the arrival of news.

Finally, the paper's results suggest that high-frequency exchange rate movements are path dependent, since the conditional distribution of future exchange rates depends in part on the rate's current level. This has potential implications for both theoretical and empirical work. For example, path dependence is inconsistent with the random walk assumption (Evans 2001) and the related assumption of a simple diffusion process (Andersen et al. 2001). The path dependence of exchange rates may also help explain why exchange rates are so difficult to forecast at high frequencies, and why technical analysis, which bases forecasts on recent price movements and levels, has a track record of forecasting success while standard exchange rate models have not. These are areas for future research.

REFERENCES

Amihud, Yakov, and Haim Mendelson, "Dealership Market: Market-Making With Inventory," *Journal of Financial Economics* 8 (1980): 31-53.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, "Modeling and Forecasting Realized Volatility," *NBER Working Paper* 8160 (March 2001).

Boothe, P., and D. Glassman, "The Statistical Distribution of Exchange Rates," *Journal of International Economics,* 22 (1987): 297-319.

Brock, W., J. Lakonishok, and B. LeBaron, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance* Vol. 48, No. 5 (December 1992): 1731-64.

Cecen, A.A., and Erkal, C., "Distinguishing Between Stochastic and Deterministic Behavior in High Frequency Foreign Exchange Rate Returns: Can Non-linear Dynamics Help Forecasting?" *International Journal of Forecasting* 12 (1996): 465-73.

Clyde, W.C., and C.L. Osler, "Charting: Chaos Theory in Disguise?" *Journal of Futures Markets* 17 (1997): 489-514. (NIL)

Danielsson, Jon, and Richard Payne, "Real Trading Patterns and Prices in Spot Foreign Exchange Markets," Mimeo, London School of Economics (March 1999).

De Grauwe, P., and D. Decupere, "Psychological Barriers in the Foreign Exchange Markets," *Journal of International and Comparative Economics* 1 (1992): 87-101.

Diamond, Douglas W., and robert E. Verrecchia, "Constraints on Short-Selling and Asset Price Adjustment to Private Information," *Journal of Financial Economics* 18 (1987): 277-311.

Dooley, Michael P., and Jeffrey Shafer, "Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981," in *Floating Exchange Rates and the State of World Trade and Payments*, Bigman and Taya, eds. (Ballinger Publishing Company, Cambridge, MA: 1984): pp. 43-70.

Dybvig, Philip H., "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market," *The Review of Financial Studies* 1 (1988): 67-88.

Easley, David, and Maureen O'Hara, "Time and the Process of Security Price Adjustment," *Journal of Finance* 47 (1992): 577-605.

Efron, B., "Bootstrap Methods: Another Look at the Jackknife," Annals of Statistics, 7, (1979): 1-26.

________ *The Jackknife, the Bootstrap, and Other Resampling Plans* (Philadelphia: Society for Industrial and Applied Mathematics, 1982).

Evans, Martin, and Richard K. Lyons, "Order Flow and Exchange Rate Dynamics," Institute of Business and Economic Research Working Paper RPF-288 (August 1999).

Evans, Martin, "FX Trading and Exchange Rate Dynamics," *NBER Working Paper* 8116 (February 2001).

Fama, Eugene F. and Marshall E. Blume, "Filter Rules and Stock Market Trading," *Journal of Business* Vol. 39, (1966): 226-241.

Garbade, Kenneth D., Jay L. Pomrenze, and William L. Silber, "On the Information Content of Prices," *The American Economic Review* 69 (March 1979): 50-59.

Garman, Mark, "Market Microstructure," *Journal of Financial Economics* 3 (1976): 257-75.

Goodhart, Charles, Takatoshi Ito, and Richard Payne, "One Day in June 1993: A Study of the working of the Reuters 2000-2 electronic Foreign Exchange Trading System," in Jeffrey A. Frankel, Gimapaolo Galli, and Alberto Giovannini, eds., *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago: 1996): 107-179.

Hasbrouck, Joel, "Trades, Quotes, Inventories and Information" *Journal of Financial Economcs* 22 (1988): 229-52.

Ho, Thomas, and Hans R. Stoll, "Optimal Dealer Pricing Under Transactions And Return Uncertainty," *Journal of Financial Economics* 9 (1981): 47-73.

Hsieh, D., "The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983," *Journal of International Economics,* 24 (1988): 129-45.

Hsieh, D., "Testing for Nonlinear Dependence in Daily Foreign Exchange Rates," *Journal of Business* 3 (1989): 339-68.

Kraus, Alan, and Hans R. Stoll, "Price Impacts of Block trading on the New York Stock Exchange," *Journal of Finance* 27 (1972): 569-88.

Levich, R., and L. Thomas, "The Significance of Technical Trading-Rule Profits in The Foreign Exchange Market: A Bootstrap Approach," *Journal of International Money and Finance* 12 (October 1993): 451-474.

Lyons, Richard K., *The Microstructure Approach to Exchange Rates*, Book in Draft (2001).

Murphy, J. Austin, "Futures Fund Performance: A Test of the Effectiveness of Technical Analysis," *The Journal of Futures Markets* Vol. 6 (1986): 175-85.

Osler, C. L., "Currency Orders and Exchange Rate Dynamics: Explaining the Success of Technical Analysis," Federal Reserve Bank of New York *Staff Report* 25 (April 2001).

Rime, Dagfinn, "Private or Public Information in Foreign Exchange Markets? An Empirical Analysis," Mimeo, April 2000.

Westerfield, Janice Moulton, "An Examination of Foreign Exchange Risk Under Fixed and floating Rate Regimes," *Journal of International Economics* 7 (1977): 181-200.

Table I: Requested Execution Rates Near Round Numbers

The table summarizes asymmetries in the distribution of requested execution rates for conditional orders near round numbers ending in 00 and 50. For each entry, I take the percent of executed orders of each order type with requested execution rates ending in the indicated set of two-digit numbers (weighted by value), and sum them. The underlying data comprise 9,669 stop-loss and take-profit orders in three highly-active currency pairs—dollar-yen, dollar-U.K. pound, and euro-dollar—processed by a major foreign exchange dealing bank during September 1, 1999 through April 11, 2000.

	Stop-loss Orders		**Take-profit Orders**	
	Buy	**Sell**	**Buy**	**Sell**
At 00	**3.5**	**5.2**	**8.1**	**10.5**
Around 00				
90-99	7.4	10.0	10.6	9.1
01-10	14.4	5.1	11.7	7.6
Around 50				
40-49	17.3	8.0	7.9	6.6
51-60	6.3	16.9	7.8	7.7

Table 2: Exchange Rate Behavior At Round Numbers

The table reports tests of the null hypothesis that exchange rates do not exhibit special behaviors at round numbers against two alternative hypotheses: (1): exchange rate trends are more frequently reflected at round numbers than at arbitrary numbers; (2): exchange rate trends are generally stronger after rates cross round numbers. The underlying data are minute-by-minute exchange rate quotes during 9 a.m. to 4 p.m. New York time over January 2, 1996 through April 30, 1998. Round numbers are rates ending in 00, such as \$1.6500/£, ¥123.00/\$, or ∈0.9800/\$, or rates ending in 50.

For 58 non-overlapping 10-trading-day intervals, the exchange rate's average behavior at round numbers was compared with its average behavior at 10,000 sets of arbitrary numbers, 30 numbers per set. For alternative hypothesis (1), I calculated the frequency with which the rate reversed course after hitting a given level (R_{RN} for round numbers, R_{AN} for arbitrary numbers). Hitting a level was defined as coming within 0.01 percent of it; reversing was defined as being above (below) a support level 15 minutes later. For alternative hypothesis (2), I calculated the average exchange-rate move after hitting a level, conditional on a failure to bounce (MV_{RN} for round numbers, MV_{AN} for arbitrary numbers). These moves have a positive sign if the previous trend was continued, and are measured in points. Each interval can be viewed as an independent Bernoulli trial, with probability one half. The final test involved counting the number of intervals in which the exchange rate's behavior at round numbers exceeds its average behavior at arbitrary levels. This number should have a binomial distribution with n = total number of relevant intervals and p = ½.

Hypothesis:			**DEM**	**JPY**	**GBP**	**ALL**
1	Trend Reversals at Round Numbers	Overall Average, R_{RN}	59.4	59.7	57.9	59.8
		Overall Average, R_{AN}	55.0	55.9	53.7	55.3
		Intervals $R_{RN}>R_{AN}$	42	38	37	NA
		Total Intervals	58	57	58	
		Marg Sig.	(0.000)	(0.008)	(0.024)	
2	Strong Trends After Crossing Round Numbers	Overall Average, MV_{RN}	5.85	6.66	4.98	5.83
		Overall Average, MV_{AN}	4.58	4.84	3.69	4.37
		Intervals $MV_{RN}>MV_{AN}$	46	46	46	NA
		Total Intervals	58	57	58	
		Marg Sig.	(0.000)	(0.000)	(0.000)	

Table 3: How Long Do The Effects of Conditonal Order Clusters Last?

The table reports tests of the null hypothesis that exchange rates do not behave differently at round numbers against the alternative hypotheses that exchange rate trends are more frequently reflected at round numbers than at arbitrary numbers. The underlying data are minute-by-minute exchange rate quotes taken over 9 a.m. to 4 p.m. New York time during January 2, 1996 through April 30, 1998. Round numbers are rates ending in 00, such as DM1.5700/$, ¥123.00/$, or $1.6500/£, or rates ending in 50.

For each 10-trading-day interval, the exchange rate's average behavior at round numbers was compared with its average behavior at 10,000 sets of arbitrary numbers, with 30 numbers in each set. I calculated the frequency with which the rate reversed course after hitting a given level (R_{RN} for round numbers, R_{AN} for arbitrary numbers), and the average (log) exchange-rate change subsequent to a hit, conditional on a failure to bounce (MV_{RN} for round numbers, MV_{AN} for arbitrary numbers). Hitting a level was defined as coming within 0.01 percent of it; reversing was defined as remaining above (below) a support level after 30, 60, or 120 minutes. Each interval can be viewed as an independent Bernoulli trial, with probability one half. The final test involved counting the number of intervals in which the exchange rate's behavior at round numbers exceeds its average behavior at arbitrary levels. This number should have a binomial distribution with *n* = total relevant intervals and *p* = ½.

Hypothesis			**DEM**	**JPY**	**GBP**
1	**Trend Reversals at Round Numbers**				
	30 Minute Horizon	Intervals $R_{RN}>R_{AN}$	36	37	31
		Total Intervals	58	57	58
		Marg Sig.	(0.043)	(0.017)	(0.347)
	1-Hour Horizon	Intervals $R_{RN}>R_{AN}$	31	38	31
		Total Intervals	58	57	58
		Marg Sig.	(0.347)	(0.008)	(0.347)
	2-Hour Horizon	Intervals $R_{RN}>R_{AN}$	31	30	28
		Total Intervals	58	57	58
		Marg Sig.	(0.347)	(0.396)	(0.448)
	24-Hour Horizon	Intervals $R_{RN}>R_{AN}$	30	24	36
		Total Intervals	58	57	58
		Marg Sig.	(0.448)	(0.145)	(0.043)
2	**Strong Trends After Crossing Round Numbers**				
	30 Minute Horizon	Intervals $MV_{RN}>MV_{AN}$	38	41	39
		Total Intervals	58	57	58
		Marg Sig.	(0.012)	(0.001)	(0.006)
	1-Hour Horizon	Intervals $MV_{RN}>MV_{AN}$	32	40	34
		Total Intervals	58	57	58
		Marg Sig.	(0.256)	(0.002)	(0.119)
	2-Hour Horizon	Intervals $MV_{RN}>MV_{AN}$	26	35	37
		Total Intervals	58	57	58
		Marg Sig.	(0.256)	(0.056)	(0.024)
	24-Hour Horizon	Intervals $MV_{RN}>MV_{AN}$	29	31	30
		Total Intervals	58	57	58
		Marg Sig.	(0.552)	(0.298)	(0.448)

Table 4: Results for JPY excluding months in which central bank intervention was reported or discussed in the business press.

The table reports tests of the null hypothesis that exchange rates do not behave differently at round numbers against two alternative hypotheses: Alternative hypothesis (1): exchange rate trends are more frequently reflected at round numbers than at arbitrary numbers. Alternative hypothesis (2): exchange rate trends are stronger after the rate crosses round numbers. The underlying data are minute-by-minute exchange rate quotes during 9 am to 4 p.m. New York time over January 2, 1996 through April 30, 1998, excluding months in which intervention in JPY was either reported in the press or discussed in the press by Japanese financial authorities. The excluded months were February, 1996; April through June, 1997; December, 1997; January, 1998; and March and April, 1998. Round numbers are rates ending in 00, such as ¥123.00/$, or rates ending in 50.

For each 10-trading-day interval, the exchange rate's average behavior at round numbers was compared with its average behavior at 1,000 sets of arbitrary numbers, with 30 numbers in each set. For alternative hypothesis (1), I calculated the frequency with which the rate reversed course after hitting a given level (R_{RN} for round numbers, R_{AN} for arbitrary numbers). Hitting a level was defined as coming within 0.01 percent of it; reversing was defined as remaining above (below) a support level after 15 minutes. For alternative hypothesis (2), I calculated the average exchange-rate move after hitting a level, conditional on a failure to bounce (MV_{RN} for round numbers, MV_{AN} for arbitrary numbers). These moves have a positive sign if the previous trend was continued, and are measured in points. Each interval can be viewed as an independent Bernoulli trial, with probability one half. The final test involved counting the number of intervals in which the exchange rate's behavior at round numbers exceeds its average behavior at arbitrary levels. This number should have a binomial distribution with n = total relevant intervals and p = ½.

Hypothesis:			**JPY**
1	Trend Reversals at Round Numbers	Overall Average, R_{RN} Overall Average, R_{AN}	64.8 58.8
		Intervals R_{RN}>R_{AN} Total Intervals Marg Sig.	28 40 (0.008)
2	Strong Trends After Crossing Round Numbers	Overall Average, MV_{RN} Overall Average, MV_{AN}	6.88 4.75
		Intervals MV_{RN} > MV_{AN} Total Intervals Marg Sig.	31 40 (0.000)

Figure 1: Requested Execution Rates: Distribution of Final Two Digits

The figures show the distribution of requested execution rates for all executed stop-loss and take-profit orders for three currency pairs--dollar-yen, dollar-pound, and euro-dollar—processed by a major dealing bank between September 1999 through April 11, 2000. Orders are weighted by value and aggregated across three currencies, but disaggregated according to type of order.

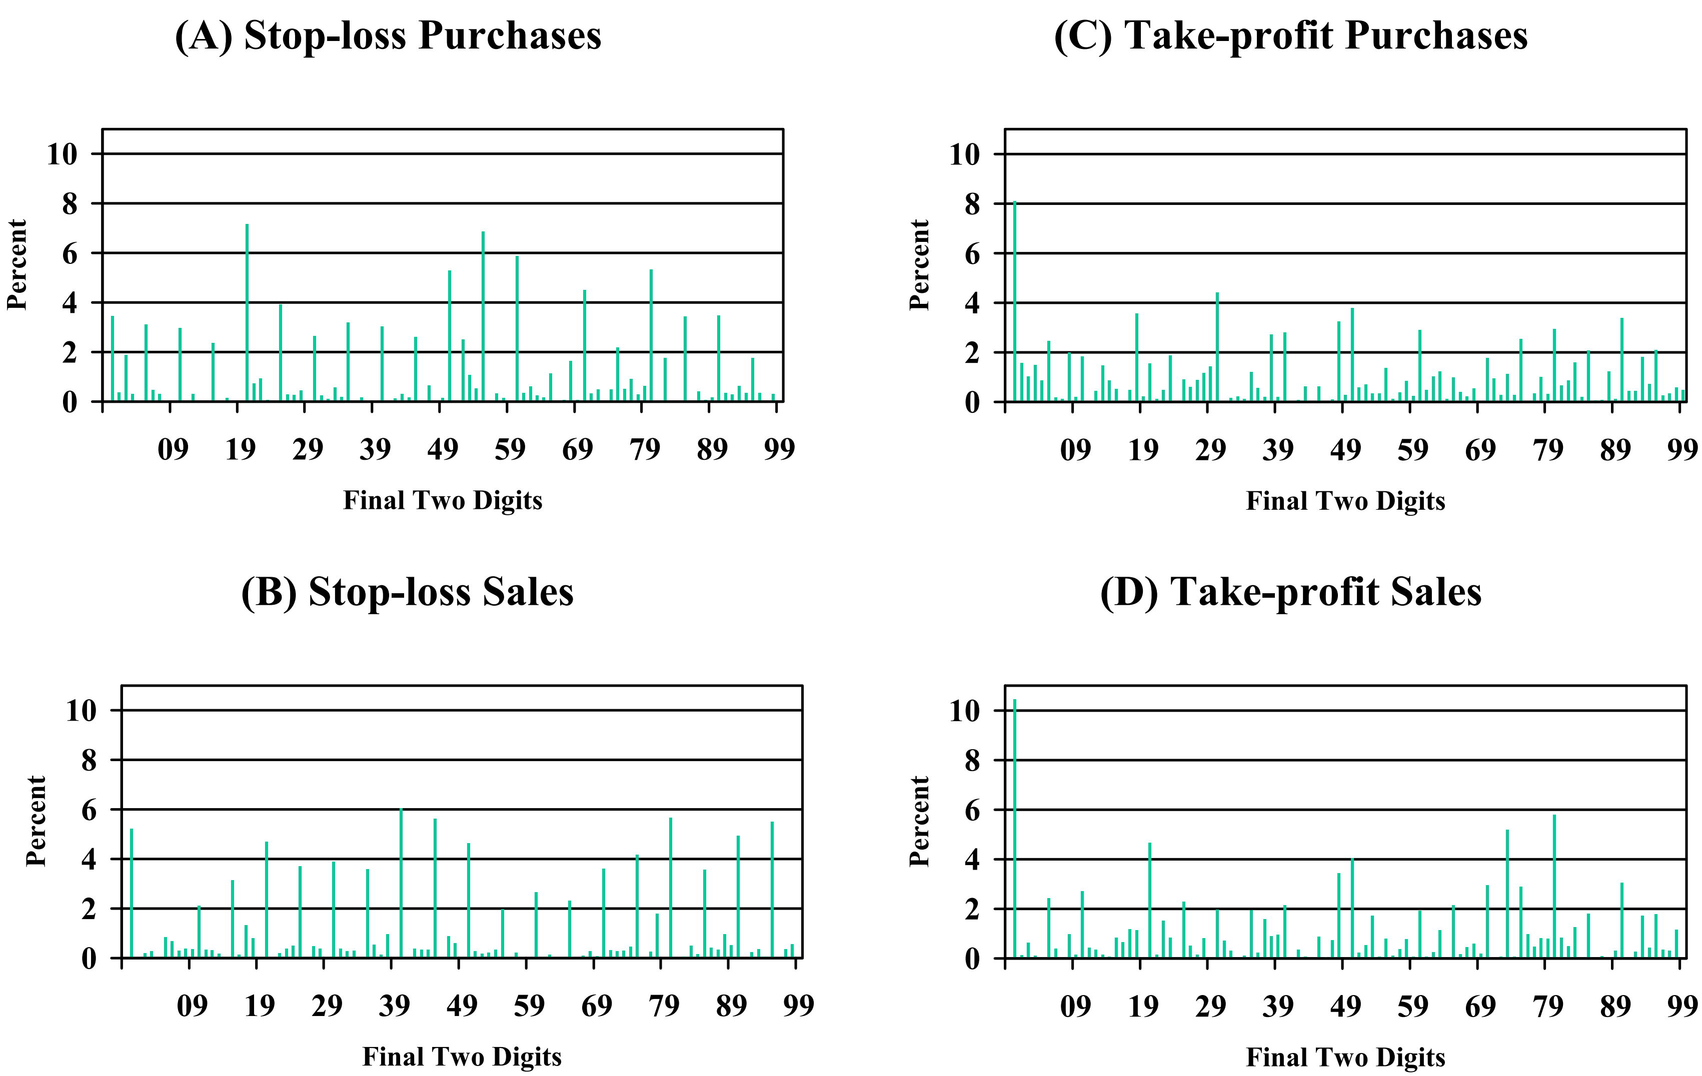

NOTES

[1] Stop-loss and take-profit orders differ structurally from the limit orders that are common in equity markets. A stop-loss or take-profit order represents a market order (buy at the market price) with a condition (only after the exchange rate reaches a particular level. A limit order indicates that the customer is *willing* to buy (sell) at a particular price, or better. Price represents the only condition attached to the order. the NYSE market specialist is not automatically required to execute a transaction for the customer (if sufficient counterparties are unavailable) if the market reaches the price specified in the order; if a transaction is executed, however, the execution rate cannot be worse than the specified rate.

[2] Note that stop-loss and take-profit orders would cause any exchange rate trend to be reversed or to accelerate, but they would not cause the rate to pause at any point. Consistent with this, DeGrauwe and Decupere (1992) note that the dollar-yen exchange rate spends less time at round numbers than at other points, on average.

[3] In a world in which expertise were costless, one might expect importing and exporting firms to consider the likely holding-period returns on their currency positions, as well as the profits to be made from the goods-and-services market transactions facilitated by those positions. In this case, one would still expect trades from financial services firms, for which holding-period returns are the primary concern, to be more informative about future exchange rates than trades from importing and exporting firms. In reality, most importing and exporting firms do not find it cost-efficient to acquire expertise in currency forecasting, so they do not consider holding period returns at all.

[4] Though the three earlier models were primarily intended to reflect the NYSE, in which a single specialist makes markets, the central intuition of these models is likely to carry over to the foreign exchange market. In foreign exchange, there are many dealers but information is dispersed as well, and it is costly for participants to search for the best price, so each dealer has a small amount of market power.

[5] Aside from Evans' (2001) fairly extreme assumption that non-dealer demand is purely stochastic, Evans and Lyons (2001) and other inventory models postulate a world in which inventory considerations co-exist with price-sensitive non-dealer demand. Inventory effects in these models would not exist if customer demand were infinitely elastic. For convenience, however, this paper will continue to refer to these joint effects as "inventory effects."

[6] Similar evidence of inventory effects exists for other markets. Krauss and Stoll (1972) provide evidence for an inventory effect on stock prices when they show that the effects of institutional block trades are largely temporary. Likewise, Garbade et al. (1979) show that inventory effects matter for dealer price-setting in the market for GNMA pass-through securities.

[7] Reuters Dealing 2000-2 is one of the two main electronic broking systems in the wholesale foreign exchange market.

[9] Included here would be a dealer's own decision to sell if the price reaches a certain level, which he never bothers to record in the order book.

[10] Osler (2001) provides further discussion of these market imperfections.

[11] It is known that high-frequency exchange rate returns do not conform to the normal distribution, since evidence consistently shows that such returns are leptokurtotic (Goodhart et al. (1996); Evans and Lyons (1999); Evans (2001)). Formal tests of the applicability of distributions other than the normal to intraday exchange-rates have not been applied, however. Formal tests applied to rates at lower frequencies have been inconclusive (Westerfield (1977); Booth and Glassman (1987); Hsieh (1988)).

[12] Arbitrary round numbers are chosen according to the following formula:

$$A = max - \alpha\, range.$$

where *max* is the maximum exchange rate for the relevant time interval, *range* is the range of rates over that same interval, and α is a random number chosen arbitrarily from a uniform distribution over the unit interval. These numbers are rounded off to the number of significant digits appropriate to each currency.

[13] The actual numbers of hits were 4456 for DEM, 3404 for JPY, and 4145 for GBP.

[14] The excluded months were February, 1996; April through June, 1997; December, 1997; January, 1998; and March and April, 1998.