

Osler, C. L.

Working Paper

Currency orders and exchange-rate dynamics: Explaining the success of technical analysis

Staff Report, No. 125

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Osler, C. L. (2001) : Currency orders and exchange-rate dynamics: Explaining the success of technical analysis, Staff Report, No. 125, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60582>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

CURRENCY ORDERS AND EXCHANGE-RATE DYNAMICS: EXPLAINING THE SUCCESS OF TECHNICAL ANALYSIS

C. L. Osler

Abstract

This paper provides a microstructural explanation for the success of two familiar predictions from technical analysis: (1) trends tend to be reversed at predictable support and resistance levels, and (2) trends gain momentum once predictable support and resistance levels are crossed.

The explanation is based on a close examination of stop-loss and take-profit orders at a large foreign exchange dealing bank. Take-profit orders tend to reflect price trends, and stop-loss orders tend to intensify trends. The requested execution rates of these orders are strongly clustered at round numbers, which are often used as support and resistance levels. Significantly, there are marked differences between the clustering patterns of stop-loss and take-profit orders, and between the patterns of stop-loss buy and stop-loss sell orders. These differences explain the success of the two predictions. (Key words: Microstructure; Exchange Rates; Orders; Technical Analysis; High-frequency.) (JEL codes G1, F3.)

March 2001

Please direct correspondence to: Federal Reserve Bank of New York/33 Liberty Street/New York, NY 10045. Telephone: (212) 720-1717. E-mail: carol.osler@ny.frb.org. The views expressed in the paper are those of the author and do not necessarily reflect views at the Federal Reserve Bank of New York or the Federal Reserve System. The author thanks Gijoon Hong for excellent research assistance. Errors or omissions are the responsibility of the author.

CURRENCY ORDERS AND EXCHANGE-RATE DYNAMICS: EXPLAINING THE SUCCESS OF TECHNICAL ANALYSIS

Exchange-rate research has brought to light a robust but unexplained result: technical analysis is useful for predicting short-run exchange rate dynamics. This result is especially striking when compared to the lack of success of standard, macro-based exchange-rate models. Such macro models are not useful for predicting short-run exchange rate dynamics themselves (Meese and Rogoff 1983), and are also unable to explain the success of technical analysis. This paper uses an alternative approach to understanding exchange rates, the microstructure approach promoted by Lyons (2001), to explain the success of two common predictions from technical analysis. This relatively new, finance-theoretic approach takes order flow to be the proximate cause of short-run exchange-rate dynamics (Goodhart et al. (1996); Evans and Lyons (1999); Rime (2000); Lyons (2001); Evans (2001)). Consistent with this view, I examine the foreign exchange order book of a large dealing bank. I show that patterns in the placement of stop-loss and take-profit orders provide an empirically-based microeconomic foundation for the success of these two predictions from technical analysis.

Virtually all the research on technical trading rules in the foreign exchange (FX) market concludes that these tools are successful at predicting exchange-rate moves. This research began soon after the advent of floating rates with studies of simple algorithms for catching trends, such as filter rules (Dooley and Shafer (1984)). Follow-up work on similar trend-following strategies showed that such predictive power was not transient (Levich and Thomas (1993)). More recent work shows that there is also predictive power for exchange rates in complicated trend-following algorithms such as the head-and-shoulders pattern (Chang and Osler (1999); Lo et al. (2000)) and in genetic algorithms (Neely et al. (1997); Neely and Weller (2000)).

The starting point for the present paper is research on certain trend-reversal signals, called "support" and "resistance" levels by technical analysts. The concepts of support and resistance are used

so commonly in the FX market that they are embedded in the language of every-day discourse among participants.¹ Technical analysts associate these signals with two main predictions:

- (1) Down-trends (up-trends) tend to be reversed at predictable support (resistance) levels;
- (2) Trends gain momentum once these predictable support and resistance levels are crossed.²

Rigorous empirical support for these two predictions is provided in Osler (2000, 2001), which analyze minute-by-minute quotes for three exchange rates—dollar-mark, dollar-yen, and dollar-U.K. pound—during New York trading hours from January 1996 through April 1998. The support and resistance levels tested include levels distributed on a daily basis by major market players to their clients (Osler 2000), as well as round numbers picked solely on the basis of their proximity to current rates (Osler 2001).

Standard exchange rate models, like the monetary model, suggest one source for such exchange rate behavior: central bank intervention. However, the evidence does not support this suggestion. With respect to trend-following trading rules, some preliminary evidence consistent with the intervention hypothesis was provided by LeBaron (1998), which highlighted a strong contemporaneous correlation between central bank intervention and the profitability of such rules. However, Neely (2000) uses time-disaggregated data to show that the strong correlation does not imply causation.³ With respect to trend-reversal signals such as support and resistance levels, Osler (2001) shows that central bank intervention could not be the source of their predictive success, either.

Standard exchange-rate models based on economic fundamentals have also had a remarkable lack of success at predicting exchange-rate movements over short horizons such as weeks or months, in sharp contrast to the success of technical trading rules which are based on nothing but recent price behavior. As early as 1983, Meese and Rogoff showed that exchange rates were predicted out-of-sample as accurately by the random walk model as by models based on fundamentals. This result has since been confirmed repeatedly. Indeed, the most familiar model of exchange rates, the monetary model, does not explain short-term exchange-rate changes well in-sample (Goldberg (2001)).

The difficulties faced by standard models do not necessarily mean that standard models have highlighted the wrong fundamentals, or the wrong set of structural relationships. Osler (1998) and Rossi (2000) both show that theoretical models could highlight the correct fundamentals while nonetheless failing empirically. Intuitively, it still seems likely that these fundamentals are critical determinants of long-run exchange-rate dynamics. The difficulties do suggest, however, that short-run dynamics are dominated by other forces, and that a new approach might be helpful in analyzing such dynamics.

This paper adopts the microstructure approach to exchange rates (Lyons (2001)), which builds on current finance theory by focusing on order flow. In particular, the paper uses one of the first available data sets of customer orders to explain the predictive success of support and resistance levels. The orders data examined here include detailed information on almost 9,700 stop-loss and take-profit orders placed at a large dealing bank over September 1, 1999 through April 11, 2000. Discussion focuses on orders placed in three markets: dollar-yen, dollar-pound, and euro-dollar. A take-profit order instructs a dealer to buy (sell) a currency at the market rate if its value falls (rises) to a certain level. A stop-loss order instructs the dealer to buy (sell) a currency at the market rate if its value rises (falls) to a certain level. (For convenience, take-profit and stop-loss orders will be referred to jointly as "conditional" orders.⁴)

Heretofore, very little detailed information on dealers' customer orders has been available to exchange-rate researchers. The only other available information on customer orders is analyzed by Lyons (2001), who finds that the information content of orders varies according to the type of agent placing the order. Some limited information about interdealer FX orders is used in Lyons (1995), which examines the order flow to and from one particular dealer over the course of a week, and which focuses on characterizing the dealer's behavior rather than the orders themselves. Otherwise, the best FX microstructure data have been limited to transactions data, and even these data have been scarce (Goodhart and Payne (1996); Goodhart et al. (1996), (1997); Evans and Lyons (1999); Rime (2000);

Evans (2001)). Neither orders nor transactions data have yet been connected empirically to the technical analysis puzzle.

The paper's central finding is based on the cross-section distribution of requested execution rates for stop-loss and take-profit orders. The sample distribution examined here is far from smooth across the range of possible rates. On the contrary, execution rates tend to cluster, with particularly strong clusters at round numbers. Most critically, the pattern of clustering differs in important ways between take-profit and stop-loss orders.

The difference in clustering patterns between stop-loss and take-profit orders, which is well-known to market participants, is potentially important for our understanding of exchange-rate dynamics. This is because short-term exchange-rates are responsive to order flow (Goodhart et al. (1996); Evans and Lyons (1999); Rime (2000), Lyons (2001), Evans (2001), Osler (2001)). If exchange rates respond to order clusters the way they respond to orders on average (Osler 2001), the two types of orders should have different effects on exchange rates. Take-profit orders should tend to reflect or reverse existing price trends, while stop-loss orders should tend to propagate or intensify them.

To explain the first prediction of technical analysts regarding support and resistance levels – that rates tend to reverse at these levels – I note that take-profit orders are more strongly clustered at round numbers than stop-loss orders. This, combined with the reflecting effect of take-profit orders, suggests that round numbers should act as partially reflecting barriers.

To explain the second prediction of technical analysts regarding support and resistance levels – that rates trend strongly after crossing such levels – I note that stop-loss orders have a pronounced tendency to be placed at rates just beyond the round numbers. In particular, stop-loss buy orders tend to be clustered just above round numbers (for example, a stop-loss buy order might be placed at 1.6605, rather than 1.6600 or 1.6595), and stop-loss sell orders tend to be clustered just below round

numbers. This, combined with the trend-intensifying effect of stop-loss orders, suggests that rates should trend strongly after crossing round numbers.

Since comparable orders data is generally unavailable to exchange-rate researchers, it is not possible to prove that these patterns are truly representative of all conditional orders received by all FX dealers. However, there is no reason to suppose that these patterns are not representative. The bank in question deals with the full spectrum of customers—bank and non-bank financial institutions, central banks, non-financial corporate customers, wealthy individuals, hedge funds, commodity trading advisors (CTAs)—from all over the world. Further, similar clustering patterns have been found in large samples of limit orders for stocks: when stocks were typically priced in eighths (a practice that ended completely in April, 2001 with the advent of decimal pricing), there was a strong tendency for limit order prices to cluster at even eighths, and a stronger tendency to cluster at whole numbers (Osborne (1962); Neiderhoffer (1965), Neiderhoffer and Osborne(1966); Harris (1991); Cooney et al. (2000)).

This paper is complemented by a closely related piece, Osler (2001), which shows that actual exchange rates behave at round numbers in ways that are consistent with the two predictions of technical analysis on which I focus here. Together, these papers support the central proposition of the microstructure approach to exchange rates, that short-run exchange-rate movements are heavily influenced by order flow (Goodhart et al. (1996); Osler (1998); Evans and Lyons (1999); Rime (2000); Lyons (2001); Evans (2001)). They also indicate that order flow can change for reasons that are unconnected to changes in the amount of private information reaching the market.

Further, the two papers provide evidence against three key assumptions of standard exchange rate models, which are relaxed in the microstructure approach (Lyons (2001), p. 5). First, the papers show that information in the foreign exchange market is not symmetrically available to all agents: on the contrary, each bank's order book constitutes valuable private information. This private information can be used for managing risk, as well as for improving exchange-rate forecasts.

Second, the papers show that agents in the foreign exchange market are not homogenous in their utility functions and beliefs: on the contrary, there are often offsetting stop-loss and take-profit orders at identical rates. Third, the papers show that institutional factors are not irrelevant for pricing outcomes. Indeed, a trading institution as seemingly innocuous as a quotation convention (should it be dollars per pound or pounds per dollar?) can have a measurable influence on high-frequency exchange-rate dynamics.

The main thesis of the paper—technical analysis works because orders are clustered—is not new. In fact, this has been the explanation provided all along by technical analysts for the success of their predictions. One respected source on technical analysis, Pring (1985), is explicit about the relationship between support and resistance and clustered order flow: "A support zone represents a *concentration of demand*, and a resistance zone a *concentration of supply*" (p. 185, italics in the original). However, these claims have never been substantiated rigorously. In the absence of such substantiation, economists have been skeptical of their validity, since the clustering of order flow is not predicted by any of our theoretical models.

The paper is organized as follows. Section I describes why and how conditional orders are used in the foreign exchange market, and provides basic descriptive statistics about the orders data. Section II documents the clustering patterns of execution rates among FX conditional orders, shows that the patterns are statistically significant, and explains heuristically how the clustering could lead exchange rates to behave as predicted by technical analysts and as documented in Osler (2000, 2001). Section III uses exchange-rate simulations to demonstrate more rigorously that the clustering of conditional orders can indeed tend to generate such exchange-rate dynamics. Section IV concludes.

I. FOREIGN EXCHANGE CONDITIONAL ORDERS: BACKGROUND

Take-profit orders instruct dealers to buy (sell) a currency at the market rate if its value falls (rises) to a certain level. Stop-loss orders instruct dealers to buy (sell) a currency at the market rate if

its value rises (falls) to a certain level. The levels at which dealers are instructed to begin trading will be referred to here as "requested execution rates," since it is the customer's hope that deals will be executed as close to these rates as possible.⁵

Intuitively, it is helpful to think of a stop-loss order as an attempt by a customer to stop losses before they get too big: for example, if the customer owns a currency, and its value falls, it will wish to sell. Likewise, a take-profit order can be viewed intuitively as an attempt to capture profits before they have a chance to disappear. It is not necessary, however, for a customer to have an existing currency position for an order to be defined as a stop-loss or a take-profit. The defining criteria are simply the direction of the trade (buy/sell) and whether the requested execution rate is above or below market rates when the order is placed.

Stop-loss and take-profit orders are widely used in the foreign exchange market. Indeed, placing such orders is a strongly recommended practice for traders in all financial markets: "Stops placement is one of the most essential aspects of successful commodity trading. Protective stops are a highly recommended trading technique" says one familiar technical analysis manual (Murphy (1986), p. 342). For the bank whose conditional order book we analyze later, conditional orders are informally estimated to represent 5 to 10 percent of total deal flow.

Dealers actively seek conditional order flow, which is useful to them in three ways. First, it generates income via bid-ask spreads. Second, it provides private information about likely exchange-rate movements, as demonstrated below. Third, take-profit orders can be used by the dealers to manage risk. For example, suppose the rate is 1.6000 and the dealer has on his/her book a large take-profit sell order at 1.6010. The dealer can with relative safety open a short position at 1.6000, knowing that his/her losses would be capped at the 1.6010 point. In Lyons (2001) taxonomy of the transparency of information, a dealer's conditional order book constitutes pre-trade order flow information that is available only to the dealer (p. 48).

Despite their common use in the FX market, stop-loss and take-profit orders are not included in standard theoretical models of that market.⁶ This section next explains the FX market imperfections that motivate the use of conditional orders, and then presents basic information about the orders explored in the rest of the paper.

I.A. Why Use FX Conditional Orders?

Existing analysis suggests that the use of stop-loss and take-profit orders is not consistent with rationality. In particular, Dybvig (1988) shows that stop-loss and take-profit orders represent an inefficient portfolio strategy in financial markets with frictionless trading, information that is instantly, costlessly, and symmetrically available to all agents, and perfectly rational trading. This section points out how the FX market fails to fit these assumptions.

Non-Frictionless Trading: The process of placing and executing foreign exchange transactions is not instantaneous. It can take a minute or more, which is a lot of time in the foreign exchange market.

To be concrete, consider a U.S. goods and services importing firm that must acquire a certain amount of U.K. pounds for business purposes within the day. In essence, the firm is short pounds. The corporate treasurer could simply buy the pounds at the current market rate of \$1.65/£. Suppose she waits instead, hoping to improve her rate and profit from intraday volatility, until the rate reaches some desirable level below \$1.65/£, say \$1.64/£. Suppose she is lucky, and the rate reaches her target of \$1.64/£. The treasurer then contacts the corporate sales desk at her foreign exchange dealing bank, requesting a two-way quote on pounds. The corporate salesperson receiving the call asks the U.K. pound spot trader nearby for a two-way quote; the salesperson takes the quote from the spot trader, adds a spread, and answers the initial question from the corporate treasurer. While this is all happening, the rate could have moved substantially, and the treasurer might not be able to trade at a rate as favorable as \$1.64/£ after all. Even if the market rate had not changed, the particular quote

from this dealing bank might not be attractive, and the treasurer might have to call another bank to get a more satisfactory rate. This adds further delay and further risk of price slippage.

Information Imperfection #1: Too Much Information Every Minute: Exchange rates change from minute to minute, sometimes from second to second; since the market generates out such a large quantity of price information, only specialists devoted to following the market can truly be "on top" of it. Consider once again the corporate treasurer short U.K. pounds. Even if the trading friction described above were not an issue, the information overload could be prohibitive. She and her staff already have well-defined responsibilities that do not include following the market attentively. They might not know when the market reaches a better rate than \$1.65/£. There are potential gains from trade if the importing firm hires a dealing firm to follow the market in its stead. The dealer, which already follows the market closely as part of its bread-and-butter business of making markets, gains the information and risk-management value of the order, as well as the bid-ask spread. The corporate treasurer gains by avoiding the costs of market monitoring.

Information Imperfection #2: Too Many Minutes: Because the FX market is open 24 hours per day, five days per week, any individual foreign exchange dealer could not successfully execute all conditional orders. Instead, it is the spot traders of a given foreign exchange bank *as a group* that have the full advantage in following markets, and they typically coordinate their efforts to ensure continuous market monitoring in service of their customers.

The information overload problem in the foreign exchange market causes conditional orders to be placed within dealing institutions. For example, a notable minority of the conditional orders examined later were placed by currency option dealers within same firm.⁷

Information Imperfection #3: Asymmetric Information Between Bank and Dealer: There is no necessary congruity between the objective function of an individual foreign exchange dealer and that of his/her employing institution. Since the institution cannot constantly know what the dealer is doing, there could be principal-agent problems in this context (Bensaid and DeBandt (2000)). To protect

themselves, dealing banks typically assign overnight stop-loss limits to each dealer—that is, each dealer is told that his/her overnight losses should never exceed a certain amount. To comply with such limits, dealers must place stop-loss orders when they leave positions open overnight.

Self-Discipline and the "Disposition Effect:" When asked why stop-loss orders are used, dealers consistently mention that they often serve as a self-discipline device for speculators. In essence, agents use these orders to control personal tendencies to time inconsistency. One such tendency has been documented so frequently in the finance literature it has its own name, the "disposition effect:" individuals tend to hold onto losing positions longer than they hold winning positions, which is a money-losing strategy (Shefrin and Statman (1985), Odean (1998), Grinblatt and Keloharju (2000)).⁸

I.B. Basic Descriptive Statistics

The data set examined here includes 9,667 conditional orders, of which 43 percent were in dollar-yen, 33 percent were in euro-dollar, and the remaining 24 percent were in dollar-U.K. pound. The average order size was \$5.8 million; median order size was \$3.0 million. These figures are somewhat higher than the average deal size of \$2.5 million estimated by Goodhart et al. (1996), and roughly span the average trade size figure of \$3.9 million reported in Evans and Lyons (1999).

On an average day, 92 new conditional orders were placed. Dollar-yen order placements peaked during the Asia afternoon/London morning, and placement of orders for the other two currency pairs peaked during the London afternoon/New York morning. Orders remained in the firm's order book until they are executed or deleted. When an office closed at the end of its trading day, the order book was passed on to other offices of the same firm around the globe, where orders continued to be monitored by currency specialists.⁹ Over 70 percent of orders were deleted or closed within one business day, and the median amount of time orders remained open was 0.4 days. Some very large orders remained open for months, however, and the average amount of time orders remained open was 3.4 days.

The number of open orders typically peaked during the London afternoon/New York morning. At 10 am New York time, there were on average 159 open orders for the three currency pairs. There was an overall peak for order execution at around the same time.¹⁰ About one third of all orders were executed, with most orders executed during the roughly 10-hour interval spanning the Asian afternoon through the London afternoon/New York morning.

II. DISTRIBUTIONS OF FX CONDITIONAL ORDERS AND IMPLICATIONS FOR EXCHANGE-RATE DYNAMICS

This section first presents the two predictions of technical analysts that will be the focus of the rest of the paper. It then provides an informal explanation of how conditional orders would generate exchange rate dynamics consistent with those predictions.

II.A. Two Predictions of Technical Analysis

Prediction 1: Down-trends (up-trends) tend to be reversed at support (resistance) levels. According to one major technical analysis manual, "support is a level or area on the chart under the market where buying interest is sufficiently strong to overcome selling pressure. As a result, a decline is halted and prices turn back again. . . Resistance is the opposite of support." (Murphy (1986), p. 59). In essence, technical analysts claim there are predictable partially-reflecting barriers for exchange rates.

Osler (2000) provides evidence consistent with this prediction using support and resistance levels distributed daily to clients by six FX market firms covering dollar-mark, dollar-yen, and dollar-U.K. pound during January, 1996 through April, 1998. These levels were tested against contemporaneous minute-by-minute quotes. Published support and resistance levels are, indeed, able to predict turning points for intraday exchange rate trends. Osler (2001) provides comparable results for round numbers.

Prediction 2: Trends gain momentum once support and resistance levels are crossed. "A breakout occurs when the stock moves outside the trendline or channel. It indicates that there's been an important shift in supply and demand. When confirmed, it should be acted on immediately" (Hardy (1978), p. 54). Though this claim does not relate exclusively to support and resistance levels, it does apply to such numbers.

Osler (2001) provides support for this prediction, using the minute-by-minute data set described above. There, it is shown that the average exchange-rate move during the 15 minutes subsequent to crossing a round number is larger than corresponding moves after crossing arbitrary levels, and the difference is highly statistically significant in all cases.

II.B. Round Numbers and Technical Analysis

In this paper I focus specifically on round numbers as support and resistance levels. For round numbers to matter with respect to currency prices, exchange rates would have to be quoted in just one way, consistently throughout the market, which is indeed the case. In order to ensure the fast execution of trades in the rapidly changing FX market, agents in the dealer/wholesale market consistently quote rates in the dollar-yen, dollar-pound, and euro-dollar markets as follows: ¥/\$, \$/£, and \$/O.

There are two reasons why I focus on round numbers in the present paper. First, the importance of round numbers as technical trading signals is stressed by technical analysts themselves: "There is a tendency for round numbers to stop advances or declines ... [R]ound numbers ... will often act as "psychological" support or resistance levels. . . . Traders tend to think in terms of important round numbers ... as price objectives and act accordingly" (Murphy 1986, p. 67). Indeed, the vast majority of support and resistance levels used in practice are round numbers (Osler (2000)).

Second, round numbers work as technical trading signals. Osler (2001) shows that round numbers function as support and resistance levels, for predictive purposes, as well as most published support and resistance levels.

II.C. The Distribution of Requested Execution Rates

Technical analysts indicate that these two predictions work because of unevenness in the flow of orders. Edwards and Magee, the most familiar authorities in the field of technical analysis, define support as "buying, actual or potential, sufficient in volume to halt a downtrend in prices for an appreciable period," and resistance as "selling, actual or potential, sufficient to satisfy all bids and hence stop prices from going higher for a time" (1997, p. 211).

Orders are indeed clustered around certain execution rates. Figure 1A plots the distributions of the last two significant digits of requested execution rates for all executed orders in each of the three currency pairs.^{11,12} (In this chart, orders are not weighted by their currency values.) Under the null hypothesis that requested execution rates are not clustered, these distributions would be approximately uniform, with each frequency at about 1 percent. In actuality, the pattern of frequencies has numerous high spikes, as well as a disproportionately high number of values under 1 percent. For all three currencies, the hypothesis that the distribution is consistent with the uniform distribution can be rejected at (far better than) the 0.01 percent level using the Anderson-Darling test (D'Agostino and Stephens (1986)).¹³

The high spikes conform to a clear pattern: requested execution rates are strongly clustered at round numbers. The single greatest cluster is at levels ending in 00 (such as ¥123.00/\$, \$0.9100/0, and \$1.6700/£) where roughly 8.7 percent of orders are placed, on average. There are also clusters, at declining levels of intensity, at levels ending in 50, at other levels ending in 0, and at levels ending in 5. There is no special tendency for rates to cluster at levels ending in 25 or 75. Requested execution rates at levels that do not end in 0 or 5 are quite infrequent. Nonetheless, within this group of there are still some intriguing preferences: there is a tendency to favor 2 and 3 over 1 and 4, and to favor 7 and 8 over 6 and 9.

The reader can see in Figure 1A that there are no qualitative differences in these distributions across the three currency pairs. The reader can also see, from Figure 1B, that these patterns are present if orders are weighted by their values.

The existence of predictable order flow clusters at round numbers is not consistent with most theoretical exchange rate models, in which the central economic variable is the log of the exchange rate. These models perform identically (with appropriately changed signs) whether the exchange rate is defined as currency A per currency B or vice versa. The models' implicit assumption is that quotation conventions are irrelevant for price determination. The clustering of orders at round numbers suggests, to the contrary, that quotation conventions matter for high-frequency price movements.

Similar clustering of order prices has long been evident in limit orders for U.S. stocks. In those markets, where prices were generally quoted in eighths until quite recently, prices for limit orders tended to cluster at "even eighths." More specifically, these prices were most strongly clustered at whole numbers, with a secondary tendency for clustering at the halves, and a tertiary tendency to cluster at the odd quarters (Osborne (1960); Neiderhoffer (1965, 1966); Harris (1991)).¹⁴

The parallels between clustering of conditional currency orders and stock limit orders are striking. This is especially true since limit orders are most closely related to take-profit orders, which show the tiering of clusters most prominently. The parallels are intensified by the fact that, within the category of whole numbers, stock prices tend to cluster at numbers ending in 10 and 5, with stronger clusters at numbers ending in 10 (Harris (1991)).

In the stock markets, clustering is stronger for higher-priced stocks (Harris (1991)). In the present setting, there are only two currency pairs for which values are appropriately comparable—dollar-pound and euro-dollar—and these two do conform to this observed regularity (see Figure 1). Clustering is also stronger for stocks traded in dealer markets, whose structure is closest to the foreign exchange market, than for exchange-traded stocks.

Despite similarities in their clustering patterns, it is important to note that the NYSE limit orders, which have been the focus of work on order clustering in stock markets, differ structurally from the FX conditional orders examined here, partly because of institutional differences between the two markets. A limit order indicates that the customer is *willing* to buy (sell) at a particular price, or better. For example, a limit order to buy 100 shares of corporation ZZZ at \$40.00 means that the customer is willing to buy up to 100 shares at any price less than or equal to \$40.00. Other than price, no other condition is attached to such an order. If the market reaches the price specified in the order, the NYSE market specialist is not automatically required to execute a transaction for the customer (if sufficient counterparties are unavailable); if a transaction is executed, however, the execution rate cannot be worse than the specified rate.

The conditional orders examined here, by contrast, each represent a *requirement* that the dealer execute a transaction if the market reaches a certain price. A stop-loss order to buy pounds at \$1.6000 means that the dealer must buy if the price hits \$1.6000, even if the transaction must be conducted at a less favorable (greater) price. Indeed, the stop-loss orders in the data set examined here are often executed at a price worse than the one specified in the original order. (It is worth noting, however, that failure to execute at the requested rate can compromise the dealer's standing with the customer.)

This structural difference between limit orders and conditional orders has important implications for asset price behavior. As shown by Neiderhoffer and Osborne (1966), stock prices should tend to linger at cluster points for limit orders: these same authors document such a tendency in actual stock prices. By contrast, exchange rates should tend to move away quickly from cluster points for conditional orders; consistent with this, De Grauwe and Decupere (1992) document that actual exchange rates tend to spend less time at round numbers than at other rates.

Why would agents cluster their orders at round numbers? This question has been previously addressed with regard to stock limit orders. Gottlieb and Kalay (1985) discuss discreteness of prices due to minimum tick sizes. In the wholesale FX market, where prices are traded in decimals, these

minimum tick sizes are essentially the fifth significant digit of the exchange rate as it is conventionally quoted.¹⁵ In the presence of minimum tick sizes, an otherwise uniform set of underlying prices would appear clustered at the available ending rates for the given commodity. However, this is not the source of the clustering documented here, since these conditional FX orders are clustered at subsets of the available decimal ending rates.

Harris (1991) suggests that a subset of the available prices is used by market participants "to simplify their negotiations:"

A small set [of prices] limits the number of different bids and offers that can be made. Negotiations may therefore converge more rapidly since frivolous offers and counteroffers are restricted. A small set also limits the amount of information that must be exchanged between negotiating traders. This reduces the time it takes to strike a bargain and it decreases the probability that two traders will believe that they have traded at different prices. These savings can be significant if trading is active (p. 390).

This motivation could well describe why FX dealers might choose round numbers in making their quotes to other dealers. However, the requested execution rates for conditional orders are placed by principals who do not negotiate themselves: their agents, the dealers, carry out that responsibility for them. Since those placing the orders will not directly bear the burden of negotiating prices, it is not immediately obvious that this explanation would apply to the conditional orders under examination.

Goodhart and Curcio (1991) suggest that different integers have different attraction levels to people. If so, this variation in basic attraction among integers could possibly represent an internal attempt by individual agents to strike a balance between the cognitive difficulty of managing large amounts of information and the potential gains from managing more information rather than less.

There may also be a self-fulfilling element to the choices of conditional orders' execution rates: given that some agents cluster their orders at round numbers, it may be rational for others to do so, as well. This is discussed below, once the distributions of requested execution rates for stop-loss and take-profit orders have been examined separately. However, this could not be a complete answer to

the question "why round numbers?" because one would still need to specify why clustering at round numbers began in the first place.

II.D. Distribution of Requested Execution Rates For Stop-loss and Take-profit Orders

To explain the predictive power of support and resistance levels, we next examine the balance between stop-loss and take-profit orders. Since take-profit orders reflect trends, and stop-loss orders intensify trends, it is the balance between these two that determines exchange-rate dynamics. The distribution of requested execution rates for each of four order types—stop-loss purchases, stop-loss sales, take-profit purchases, take-profit sales—is depicted visually in Figures 2A-D. (In these figures, all currencies have been aggregated, and are weighted by value so the distributions correspond to those used in the simulations reported in Section III.) The most important tendencies observable in the figures are summarized numerically in Table I.

The tendency to cluster at round numbers is shared across the four types of orders, but there are some critical differences that can potentially explain the success of two predictions of technical analysis on which I focus.

Explaining Prediction 1: Down-trends (up-trends) tend to be reversed at support (resistance) levels. Take-profit orders show a stronger tendency to cluster at numbers ending in 00 than stop-loss orders (Table I). In fact, 9.3 percent of take-profit orders are executed exactly at 00, while the corresponding percent for stop-loss orders is only 4.4.

This difference will be important for later analysis, so I test it for statistical significance. In choosing a null hypothesis and a test methodology, a difficulty immediately presents itself: the underlying frequency distributions do not conform to any standard, easily parameterized distribution. In consequence, some non-parametric procedure seems mandated. My choice is to bootstrap the population distribution of frequencies for rates ending in 00, based on some null hypothesis, from the observed distribution of frequencies (for more on the bootstrap, see Efron (1979, 1982)).

My null hypothesis has two components: (1) the frequencies for rates ending in 0 have the same distribution for stop-loss and take-profit orders, and (2) 00 is not special. The measure on which I choose to focus is the difference of 9.9 percentage points between the sum of the actual frequencies at 00 for take-profit orders and the sum of actual frequencies at 00 for stop-loss orders. I need to find the likelihood of observing a value that large under the null.

To implement the bootstrap test, I first choose at random 4,000 frequencies for rates ending in 0 from the original set of 40 such values, disregarding whether they correspond to stop-loss or take-profit orders. I then sum 2,000 randomly chosen pairs of these frequencies, and take 1,000 randomly chosen differences across pairs. Finally, I take the absolute values of these 1,000 differences, rank them, and use the resulting distribution to calculate the marginal significance of the observed difference. That marginal significance is 1.1 percent, which indicates that take-profit orders are indeed more strongly concentrated than stop-loss orders at 00.

To understand the implications of this difference, suppose rates rise to a round number such as \$1.6600£. Since rates are rising, the conditional orders triggered by this rise would be stop-loss buys and take-profit sells. On average, according to the data behind Figures 2A-D, roughly 2.8 percent of all stop-loss buy orders are triggered when the rate hits a round number ending in 00, while 10.5 percent of take-profit sales orders are triggered. The currency amount of stop-loss and take-profit orders is not greatly different, on average; altogether this means that conditional order flow after arriving at \$1.6600£ should tend to be dominated by take-profit sales orders. In that case, rates would have to fall in order to draw out corresponding buy orders from the rest of the market. In short, this difference between stop-loss and take-profit orders in the amount of clustering at 00 seems to provide a natural, though heuristic, explanation for technical analysts' first prediction.

Explaining Prediction 2: Trends gain momentum once support and resistance levels are crossed. A second striking asymmetry concerns the clustering just above and just below round numbers ending in 00 and 50. Stop-loss buy orders are relatively infrequent at rates just below these

two levels, and have a pronounced tendency to cluster at rates just above these two levels (Table I). Consider the frequency of orders at rates ending near 00: a total of 7.4 percent of all stop-loss buy orders are placed at rates ending between 90 and 99, inclusive; in contrast, almost twice as many stop-loss buy orders, 14.4 percent, are placed at rates ending between 01 and 10, inclusive. A similar asymmetry can be observed above and below the 50 level.

By contrast, stop-loss sell orders tend to cluster at rates just below 00 and 50, and are relatively infrequent at rates just above these two levels. Take-profit buy and sell orders exhibit no pronounced asymmetries around 00 and 50 (Table I).

These asymmetries in stop-loss orders are common knowledge among market practitioners. Nonetheless, it is worth testing whether they are significant, since they are critical to later results. As before, the fact that the underlying distributions do not conform to the uniform prevents any simple parametric analysis. Instead, I turn again to the bootstrap technique (Efron (1979), (1982)). My null hypothesis is that agents' preferences for round numbers are tiered into four roundness categories: (1) all rates ending in 0; (2) all rates ending in 5; (3) all rates ending in 2, 3, 7, and 8; (4) all rates ending in 1, 4, 6, and 9. Beyond this, I assume that the differences between frequencies within roundness categories are random. The alternative hypothesis is that the asymmetries evident in stop-loss orders are meaningful, while those in take-profit orders are not.

For each type of order, I create 2,000 sums of 10 frequencies, where each set includes:

1. One of the original frequencies for rates ending in 0;
2. One of the original frequencies for rates ending in 5;
3. Four of the original frequencies for rates ending in either 2, 3, 7, or 8;
4. Four of the original frequencies for rates ending in either 1, 4, 6, or 9.

All of these frequencies are selected at random, and each has an equal chance of being selected. These sums correspond to the 10 frequencies between, say, rates ending in 90 and 99, or the 10 frequencies

for rates ending between 01 and 10. I then take 1,000 differences between these sums. The ranked differences provide the distribution from which marginal significance levels are calculated.¹⁶

Consistent with the alternative hypothesis, marginal significance levels for three of the four stop-loss asymmetries fall below 5 percent; the marginal significance of the fourth is 6.4 percent. Meanwhile, all four the marginal significance levels for the take-profit asymmetries exceed 10 percent by comfortable margins.¹⁷

Given the trend-intensifying effects of stop-loss orders, these statistically significant asymmetries suggest that rates will often trend strongly after crossing a round number. Suppose once again that rates rise through \$1.6600/£ and reach \$1.6610£. Since rates are rising, conditional order flow will be determined by stop-loss buy orders and take-profit sell orders. At levels ending in 10, roughly 4.5 percent of all stop-loss buy orders are executed and roughly 2.7 percent of all take-profit sales orders are executed, according to the data underlying this analysis. Since the currency amounts of stop-loss and take-profit orders are roughly equal, on average, conditional order flow would likely be dominated by buy orders at this rate. With such an imbalance of conditional orders, the price would likely rise to find offsetting sales orders from other agents in the market. In short, the asymmetric distributions of stop-loss buy and sell orders near round numbers, when compared with the fairly symmetric distributions of take-profit orders, seem to provide a natural heuristic explanation for the technical analysts' second prediction.

As mentioned above, the asymmetric placement of stop-loss orders suggests the possibility of a rational self-fulfilling dynamic between order placement and exchange-rate dynamics. If take-profit orders tend to be placed just at round numbers, then the likely point of maximum gain for someone with a position will be the round number. That could make it rational, given the market imperfections discussed in Section I, to place take-profit orders at the round number. Likewise, if stop-loss orders tend to be placed just beyond round numbers, there would be a tendency for exchange-rates to trend strongly at such points. Given such a tendency, it could be rational for individuals to place stop-loss

orders just beyond round numbers. Indeed, market practitioners recommend placing stop-loss orders just beyond support and resistance levels, basing the recommendation at least in part on the assumption that the two predictions on which we focus here are true. In short, the tendency to place stop-loss orders at certain points could conceivably be creating the conditions necessary for that same tendency to be (conditionally) rational.

This potential for rationally self-fulfilling conditional order-placement is inconsistent with the analysis of Dybvig (1988) which concludes that the use of conditional orders is inefficient. However, his analysis assumes, in addition to perfect markets, that prices follow a standard lognormal diffusion process, which does not seem to be true of the FX market at the high frequency examined here. On the contrary, the behavior of exchange-rates at this horizon is path dependent, since prices tend to be reflected at round numbers or to trend rapidly if they cross round numbers (Osler 2001). This observation raises questions about certain standard forms of analysis. For example, it has been noted that "mean-variance analysis is not valid when the portfolio return is nonlinearly related to market returns, as it will be under [dynamic portfolio] strategies" such as conditional orders (Dybvig and Ingersoll (1982), p. 67).

III. CONDITIONAL ORDERS AND THE DYNAMICS OF SIMULATED EXCHANGE RATES

To provide a more rigorous demonstration that exchange rates affected by conditional orders could indeed display the tendencies claimed by technical analysts, I examine some simulated exchange rates. These simulations are intended to illustrate the likely behavior of exchange rates dominated by conditional order flow. They are not intended to represent the behavior of actual exchange rates: on the contrary, they are intentionally structured so that the effects of conditional order flow exceeds those one might expect in reality. The intended inference is that, since actual exchange rates are affected by

conditional order flow but not completely dominated by it, tendencies similar to those observed in the simulated data will likely be present in actual rates though to a lesser degree.

The simulations are a relatively rigorous demonstration of the influence of order clusters, compared with the heuristic discussion in Section II. This is because the simulations incorporate the distribution of conditional orders at all levels, while the heuristic discussion focused on distributions at just a few levels. In the simulations all four types of conditional orders are calibrated to have requested execution rates distributed according to the actual distributions shown in Figures 2A-D.

III.A. Exchange-Rate Simulations

Overview: The simulated exchange rates are calculated using an iterative procedure. Exchange-rate changes are driven by current order flow, $O(t)$. In turn, the current exchange rate (log) level, $x(t)$, and the most recent exchange-rate change, $x(t+1) - x(t)$, affect the next period's order flow.

From Orders to Exchange Rates: We know that order flow affects exchange rates, as argued by Osler (1998) and as documented by Goodhart et al. (1996), Evans and Lyons (1999), Rime (2000), Lyons (2001), Evans (2001), and Osler (2001). To a first approximation, orders in the simulations determine the change in the log exchange rate via some function $F(\cdot)$: $x(t+1) - x(t) = F(O(t))$. Here $x(t)$ is the log of the exchange rate, and $F(\cdot)$ is the function relating order flow to exchange-rate changes.

Order flow is defined here roughly as it is in Evans and Lyons (1999), Lyons (2001), and Evans (2001): the difference between orders initiated by buyers and orders initiated by sellers. However, here it is defined as the value of buy-initiated orders minus the value of sell-initiated orders, while in the other papers it is defined as the number of buy-initiated trades minus the number of sell-initiated trades. When order flow is positive (negative), prices must rise (fall) to draw additional sellers into the market. Beyond the basic facts that $F(0) = 0$ and $F'(\cdot) > 0$, the empirical exchange rate literature provides little further guidance about the best functional form for $F(\cdot)$, so I set the relationship to be

linear: $x(t+1) - x(t) = \mathbf{b} O(t)$, $\mathbf{b} > 0$. This can be viewed as a linear approximation to the "true" functional form.

Total Orders: Total order flow comprises fundamentals, arbitrage orders, and conditional orders: $O(t) = F(t) + A(t) + C(t)$. Fundamentals are represented schematically as a constant multiple, \mathbf{m} of a standard normal variable with realizations $v(F,t)$: $F(t) = \mathbf{m}v(F,t)$. These orders would, other things equal, cause the exchange rate to follow a pure random walk. This is consistent with results in Evans and Lyons (1999) and Evans(2001), which find long-lasting effects of order flow on exchange rates.

Arbitrage orders are placed by agents attempting to arbitrage the gap between the exchange rate's current value and its long-run equilibrium value, \bar{x} . These orders are set to be a fixed proportion, δ , of that gap: $A(t) = \delta(x(t) - \bar{x})$.

Conditional Orders: The amount of conditional orders executed each period, $C(t)$, is determined by the events leading up to that period plus random factors. One can think of the simulated order book as being refreshed each period, rather than having orders open that were initiated at varying times in the past, as is true in reality. The simulated order flow is calibrated to match the distribution of execution rates in the original data set on average across periods, with a certain amount of randomness from period to period.

Denote the four types of orders—stop-loss buy, stop-loss sale, take-profit buy, take-profit sale—as sb , ss , tb , and ts , respectively. In each period, only two of these four types are executed. If the exchange rate has risen (fallen), then stop-loss buy (sell) and take-profit sell (buy) orders are executed. For convenience, let d be a dummy variable associated with direction of exchange-rate movement from $t-1$ to t : $d = 1$ for a rising exchange rate, $d = 0$ for a falling exchange rate. Then total conditional orders are:

$$C(t+1) = d[sb(t) - ts(t)] + (1-d)[tb(t) - ss(t)].$$

The value of each type of order is determined by three factors: first, the particular set of two-digit exchange rate levels through which the exchange rate has just passed; second, the proportion of total orders of that type in the original data set at each of those particular two-digit levels; and third, a random variable. The random variable captures the fact that, in reality, orders are determined by more than the two-digit levels through which the exchange rate has just passed. Because of the random variable, the value of orders at a given two-digit point such as 05 will not be a constant, though the value of orders at 05 will, on average, represent the correct fraction of all such order value.

More formally, let $\Theta(i,j)$ be the fraction of all executed orders in the original data set represented by type $i \in [sb,ss,th,ts]$ with last two digits $j \in [00,99]$. For example, $\Theta(sb,00)$ is the percent of all stop-loss buy orders in the original data set with requested execution rate ending in 00, which happens to be 0.0284.

Let $J(t)$ be the set of two-digit final numbers through which the exchange rate passed between $t-1$ and t . For example, if the rate rose from 1.6000 in $t-1$ to 1.6002 in t , then $J(t)$ would be (01,02). The value of orders of type i triggered at each individual level j is the product of $\Theta(i,j)$ and a random draw associated with level j , $v(i,j)$. The $v(i,j)$ are drawn from an underlying standard normal distribution. The total value of orders of type i triggered by rate changes going in to period $t+1$ is thus:

$$i(t) = \sum_{j \in J(t)} \Theta(i,j)v(i,j).$$

Total orders: Total orders are simply the sum of the three types of orders. To enhance computational simplicity, conditional orders, $C(t)$, are also given a coefficient, λ :

$$O(t) = \mathbf{I} C(t) + \mathbf{m}(F,t) + \delta(\bar{x} - x(t)) \quad (1)$$

From Orders to Exchange Rates, A Closer Look: Most major exchange rates are not quoted with more than five significant digits. Though the reason for this convention has not been explored empirically, the convention is central to the way orders tend to cluster at round numbers, so it must be incorporated into these simulations. Once order flow is calculated, therefore, the final exchange rate for

the next period can not be calculated directly. Instead, the next step is to calculate a *preliminary* value for $x(t+1)$, denoted $x'(t+1)$:

$$x'(t+1) - x(t) = \mathbf{b}O(t) . \quad (2)$$

This preliminary log exchange rate is then used as a basis for calculating a preliminary non-log exchange rate, $X'(t+1)$: $X'(t+1) = \exp\{x'(t+1)\}$. The preliminary exchange rate, $X'(t+1)$, is then rounded off to the appropriate number of digits to create the final exchange rate, $X(t+1)$. This, in turn, is used to calculate the final log exchange rate, $x(t+1) = \ln(X(t+1))$.

After $X(t+1)$ is calculated, the next round of calculations, associated with period $t+2$, can begin.

III.B. Testing The Simulated Exchange Rates

The next step in this analysis is to test whether these simulated exchange rates display the properties claimed by technical analysts. The null hypothesis is that the behavior of the simulated exchange rates at round numbers is no different than its behavior at arbitrary numbers. The test I use is based on the bootstrap (Efron (1979), (1982)), and involves a number of steps which I summarize here before providing details below. This methodology closely parallels that used in Osler ((2000), (2001)), to show that actual exchange rates conform to the same two predictions.

The simulated exchange rates series included 250,000 observations each. Informally, each observation may be considered equivalent to 15 minutes of active trading time.¹⁸ Assuming there are about 18 hours of active trading per day, and 260 active trading days per year, each simulated series corresponds to about 13 years of exchange rate observations.

I divide the sample period into intervals equivalent to about 10 trading days each, parallel to the two-week trading periods examined in Osler (2001). There are 347 of such intervals.¹⁹ Within each interval I examine the behavior of the simulated exchange rates at round numbers, calculating a number of factors including the frequency with which exchange rate trends are reversed at such numbers. Next, I examine the behavior of the simulated exchange rates at 10,000 sets of arbitrary

numbers, calculating all the same factors. In the penultimate step, I compare each factor for round numbers with the corresponding average value of that factor for the 10,000 sets of arbitrary numbers, interval by interval. I calculate the number of intervals during which the factor's value for round numbers exceeds its value for arbitrary numbers. This number should be distributed according to the binomial distribution, which provides marginal significance levels. These individual steps are now described in detail.²⁰

Step 1, Evaluating Simulated Exchange Rates Around Round Numbers: For each subinterval I first calculate all the round numbers between the period's maximum and minimum. I next examine whether the exchange rate tends to reverse course at such numbers, and also examine the size of exchange rate movements subsequent to reaching round numbers. For all calculations I use horizons of one and two periods following the hit.

I define the exchange rate to have "hit" or reached a round number when it comes within 0.01 percent of it. The exchange rate is defined to have "bounced" off a round number used as a support (resistance) level if the rate rises (falls) over the next period (or two) after a hit. The "bounce frequency" is defined as the ratio of the number of times the rate bounces to the total number of times it reached a round number.

Step 2, Bootstrapping: If these simulated exchange rates behave as predicted by technical analysts, then the bounce frequency at round numbers should be "high," and the average move conditional on a failure to be reflected by the level should be "large." To establish benchmarks for "high" bounce frequencies and "large" exchange-rate changes I examine the average behavior of the same simulated exchange rates around arbitrary exchange-rate levels.

To calculate the benchmarks, I first calculate 10,000 sets of arbitrary exchange rates for each sub-interval. Each set includes the same number of levels as the set of round numbers examined for that sub-interval. These arbitrary levels must satisfy two requirements: they must fall within the

observed range for that sub-interval, and they must have the required number of significant digits.²¹ I then examine the behavior of the simulated exchange rates at each of these arbitrary numbers.

Step 3, Comparing Round Numbers to Arbitrary Numbers: As remarked above, the null hypothesis is that the behavior of the simulated exchange rate series at round numbers is no different than its behavior at arbitrary numbers. Under this hypothesis, each sub-interval is equivalent to a Bernoulli trial in which the results for the round numbers have an even chance of exceeding those for the arbitrary numbers. The number of sub-intervals in which there are hits for both round numbers and arbitrary numbers is the total number of such trials, denoted n ($n \neq 347$). Given the structure of the simulations, these trials can be viewed as independent for the purposes of this test. The number of sub-intervals in which the results for round numbers exceed the results for arbitrary numbers should be distributed according to the binomial distribution with parameters $(n, 1/2)$.

III.C. Parameters and Starting Values

To implement the simulations, it is necessary to select values for four parameters. From equation (1), I need the coefficient on conditional orders, \mathbf{I} , the coefficient on purely random orders, \mathbf{m} , and the coefficient on arbitrage forces, \mathbf{d} . From equation (2) I need the effect of orders on exchange rates, \mathbf{b} .²² These parameter values are set according to three criteria: First, conditional orders should represent the dominant source of order flow. This simply ensures that the simulations highlight the effects of conditional orders per se, and not the other, ancillary components of simulated order flow. Second, average exchange-rate changes should not be less than about one point. Otherwise, there are too many periods of zero change and the results are not meaningful. Finally, average exchange-rate changes should not exceed a three or four points. This ensures that the effects of conditional orders at specific values can be cleanly and clearly delineated. If average exchange rate changes are large, for example, the effects of take-profit orders at round numbers would often be mingled with the effects of stop-loss orders at close-by numbers.

Many sets of parameters meet these three criteria, and the associated results are quite consistent with each other. In consequence, the results presented below reflect one representative set of parameters in which the weights attached to conditional and random-variable order flow (λ and m) are both set at 0.5, the measure of arbitrage strength or mean reversion (d) is set at 3.0, and the effect of orders on log exchange-rate changes (b) is set at 0.00015. The long-run equilibrium exchange rate, necessary for arbitrage order flow, is set equal to the series' starting value.

The presence of conditional orders makes the exchange rate path-dependent. This observation has potentially interesting implications for exchange-rate dynamics (see the conclusion); for the present it simply means that the results of this analysis depend slightly on the exchange rate's starting value. The tests were therefore run on simulated exchange rates beginning at ten different levels: 1.6000, 1.6011, 1.6022, 1.6033, 1.6044, 1.6055, 1.6066, 1.6077, 1.6088, and 1.6099. In these ten simulations, the average share of conditional orders in all orders is about 75 percent (see Table II). The associated exchange rate series all have average absolute exchange rate changes of about 2 points. Neither of these values varies substantially with the chosen starting point.

The "round numbers" tested as support and resistance levels are defined in two alternative ways. In one, any exchange rate with final two digits 00 or 50 are included as candidate support and resistance levels. In the other, round numbers are defined more narrowly, including only exchange rates ending in 00. Not surprisingly, there are more "hits" when round numbers are defined relatively broadly. The number of periods in the simulations is set at 250,000 in order to ensure statistically reliable results when round numbers are defined narrowly.

III.D. Results

The tests indicate that exchange rates dominated by conditional order flow will behave as predicted by technical analysts upon reaching round numbers. The simulated rates have a clear tendency to be reflected at such numbers and, if they are not reflected, they tend to continue trending

rapidly. The following discussion focuses on round numbers defined to include exchange rates ending in both 00 and 50, since results are similar when they are defined more narrowly.

Prediction #1: Down-trends (up-trends) tend to be reversed at support (resistance) levels.

Simulated exchange rates dominated by conditional orders do tend to be reflected at round numbers (Table III). At both the one-period and the two-period horizons, the bounce frequency for round numbers exceeds the bounce frequency for arbitrary numbers in each of the 20 cases (10 simulations, 2 time horizons). The results are statistically significant at the 0.01 percent level in all cases. The increase in bounce frequency associated with round numbers averages about 3.4 percentage points. This is comparable to the increase in bounce frequency associated with round numbers in actual exchange rates, which averages 4.6 percentage points across three currencies (dollar-mark, dollar-yen, dollar-pound; Osler (2001)). There are no noticeable differences between the results at the short and long horizons.

Prediction #2: Trends gain momentum once support and resistance levels are crossed.

Simulated exchange rates dominated by conditional orders tend to trend strongly after crossing round numbers (Table IV). In all 20 cases, the average exchange-rate move after a hit, conditional on a failure to bounce, is higher for round numbers than for arbitrary numbers, and the difference is statistically significant at the one percent level or better. The average move over two periods after a hit and failure to bounce is about 0.0037 percent larger for round numbers than for arbitrary numbers. This is quite a bit bigger than its average move after hitting and failing to bounce at arbitrary numbers, which is about 0.0045 percent. For actual exchange rates, the comparable (30 minute) average move following a hit and failure to bounce is about 0.0014 percent larger than its average move following a hit and failure to bounce at arbitrary rates, which is 0.0063 percent (Osler (2001)).

V. CONCLUSIONS

This paper focuses on two predictions of technical analysis: (1) down-trends (up-trends) tend to be reversed at support (resistance) levels, which are often round numbers, and (2) trends gain momentum once support and resistance levels are crossed. These are arguably the predictions most commonly used in the FX market, where technical analysis is so ubiquitous as to be part of the language of daily conversation.

Previous research has shown that actual exchange rates conform to these predictions on an intraday basis (Osler (2000), (2001)). The present paper provides a microstructural explanation for this behavior based on a close analysis of the complete conditional order book of a reasonably large foreign exchange dealing bank over September 1, 1999 through April 11, 2000.

The paper shows that requested execution rates for the stop-loss and take-profit orders included in the book are clustered at round numbers. Importantly, the pattern of clustering differs between take-profit orders, which reflect price trends, and stop-loss orders, which intensify price trends. The differences in these clustering patterns can produce the behaviors predicted by technical analysts.

The paper first connects conditional orders and the two predictions of technical analysts heuristically. It then uses illustrative exchange rate simulations to make the connection more rigorous. In the simulations, exchange rate dynamics are dominated by the flow of conditional orders, the requested execution rates of which are calibrated to match the actual conditional order book.

The results here suggest the possibility of a self-fulfilling dynamic between order placement and exchange-rate dynamics. This is because the very exchange-rate behaviors implied by the order clusters suggest that it might be rational to place orders in the very places we find clusters. The possibility that technical analysis might be a rational, self-fulfilling phenomenon is explored by Goldbaum (1999), and Shareen et al. (1998).

The overall results of this paper, together with the fact that the two predictions on which it focuses are consistent with the reality of exchange rate dynamics (Osler 2001), support three hypothesis about order flow. The first, promoted by academics such as Goodhart et al. (1996), Osler (1998), Evans and Lyons (1999), Rime (2000), Lyons (2001), and Evans (2001), proposes that order flow is a critical determinant of short-run exchange-rate movements. The second, promoted by technical analysts, proposes that technical analysis derives its predictive power from patterns in the underlying order flow. The third is that order flow may rise and fall for reasons unconnected with changes in private information.

The results of this paper also imply that many assumptions of standard exchange rate models are not consistent with reality: contrary to those assumptions, there *is* private information in the foreign exchange market, agents in that market *are* heterogeneous, and institutional features of the market, such as quotation conventions, *can* matter for prices.

The results of this paper are silent on a the closely related question of market efficiency, since no attempt is made to evaluate whether the patterns documented here create opportunities for risk-adjusted profits.

The analysis suggests at least two additional lines of research. First, the clustering of conditional orders implies that exchange rates are path dependent, as mentioned in Section III. This suggests a possible explanation for the well-known "exchange-rate disconnect" problem, that is, the difficulty economists and others have encountered in using standard structural models to forecast exchange rates over short horizons (Meese and Rogoff (1983)). In addition, it can be shown that the clustering of conditional orders induces excess kurtosis (relative to the normal distribution) in conditional order flow. If the microstructure approach to exchange rates is correct, and short-run exchange rate dynamics are dominated by the influence of order flow, then excess kurtosis of conditional order flow could contribute to observed excess kurtosis in high-frequency exchange rate changes.

REFERENCES

- Ball, Clifford A., Walter N. Torous, and Adrian E. Tschoegl, "The Degree of Price Resolution: The Case of the Gold Market," *The Journal of Futures Markets* 5 (1985): 29-43.
- Bensaid, Bernard, and Olivier De Bandt, "Les Strategies 'Stop-Loss': Theorie et Application au Contrat Notionnel du Matif," *Annales d'Economie et de Statistique* 0 (April-June 2000): 21-56.
- Chang, P.W.K., and C.L. Osler, "Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts," *Economic Journal* 109 (1999): 636-661.
- Cooney, John W., bonnie van Ness, and Robert Van Ness, "Do investors Avoid Odd-Eighths Prices? Evidence from NYSE Limit Orders," Mimeo (December 2000).
- D'Agostino, Ralph B. and Michael A. Stephens, *Goodness-of-Fit Techniques* (New York: Marcel Dekker, Inc., 1986).
- De Grauwe, Paul, and Danny Decupere, "Psychological Barriers in the Foreign Exchange Market," *Journal of International and Comparative Economics* 1 (1992): 87-101.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *The Journal of Finance* 54 (June 1990): 379-95.
- 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.
- Dybvig, Philip H., and J.E. Ingersoll, Jr., "Mean-Variance Theory in Complete Markets," *Journal of Business* 55 (1982): 233-51.
- Edwards and Magee, *Technical Analysis of Stock Trends*, Fifth edition (Boston: John Magee Inc: 1997).
- 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).

Goldbaum, David, "Profitability and Market Stability: Fundamentals and Technical Trading Rules," Rutgers University Working Paper (October 1999).

Goldberg, Michael, "Do Monetary Models of the Exchange Rate With R.E. Fit the Data?" Mimeo, University of New Hampshire (November 2000).

Goodhart, Charles, and Richard Curcio, "The Clustering of Bid/Ask Prices and the Spread in the Foreign Exchange Market," LSE Financial markets Group *Discussion Paper No. 110* (January 1991).

Goodhart, Charles, T. Ito, and R. Payne, "One Day in June 1993: A Study of the Workings of Reuters D2000-2 Electronic Foreign Exchange Trading System," In *The Microstructure of Foreign Exchange Markets*, eds. J. Frankel, G. Galli and A. Giovannini (University of Chicago Press, Chicago: 1996).

Goodhart, Charles, and Richard Payne, "Microstructural Dynamics in a Foreign Exchange Electronic Broking System," *Journal of International Money and Finance* 15 (1996): 829-52.

Goodhart, Charles, Yuanchen Chang, and Richard Payne, "Calibrating an Algorithm for Estimating Transactions from FFX Exchange Rate Quotes," *Journal of International Money and Finance* 16 (1997): 921-30.

Gottlieb, Gary, and Avner Kalay, "Implications of the Discreteness of Observed Stock Prices," *The Journal of Finance*, 40 (March 1985): 135-53.

Grinblatt, Mark, and Matti Keloharju, "What Makes Investors Trade?" Yale International Center for Finance *Working Paper No. 00-02*.

Hardy, C. Colburn, *The Investor's Guide to Technical Analysis* (New York: McGraw-Hill Book Company, 1978).

Harris, Lawrence, "Stock Price Clustering and Discreteness," *The Review of Financial Studies* 4 (1991): 389-415.

Krugman, Paul, and Marcus Miller, "Why have a target zone?" *Carnegie Rochester Conference Series on Public Policy* 38 (1993): 279-314.

LeBaron, Blake, "Technical Trading Rule Profitability of Foreign Exchange Intervention" *Journal of International Economics* Vol. 49 (October 1999): 125-43.

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.

Lo, Andrew, Harry Mamaysky, and Jiang Wang, "Foundations of Technical Analysis," *Journal of Finance* 54 (August 2000): 1705-65.

Lyons, Richard K., "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics* 39 (1995): 321-51.

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

Meese, Richard, and Kenneth Rogoff, "Empirical Exchange Rate Models of the Seventies: Do They Fit Out-of-Sample?" *Journal of International Economics* 27 (1983): 3-24.

Murphy, John J., *Technical Analysis of the Futures Markets: A Comprehensive Guide to Trading Methods and Applications*, (New York Institute of Finance, 1986).

Neely, Christopher J., "The Temporal Pattern of Trading Rule Returns and Central Bank Intervention: Intervention Does Not Generate Technical Trading Rule Profits," Mimeo, Federal Reserve Bank of St. Louis (December 2000).

Neely, Christopher J., and Paul A. Weller, "Predictability in International Asset Returns: A Re-Examination," *Journal of Financial and Quantitative Analysis* 35 (December 2000): 601-20.

Neely, C.J., Paul A. Weller, and R. Dittmar, "Is Technical analysis in the Foreign Exchange market Profitable? A Genetic Programming Approach," *Journal of Financial and Quantitative Analysis* 32 (1997): 405-26,

Niederhoffer, V., "Clustering of Stock Prices," *Operations Research* 13 (1965): 258-65.

Niederhoffer, V., "A New Look at Clustering of Stock Prices," *Journal of Business* 39 (1966): 309-313.

Niederhoffer, V., and M.F.M. Osborne, "Market Making and Reversals on the Stock Exchange," *Journal of the American Statistical Association* 61 (1966): 897-916.

Odean, Terrance, "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance* 53 (October 1998): 1775-98.

Osborne, M.F.M., "Periodic Structure in the Brownian Motion of Stock Prices," *Operations Research* 10 (1962): 345-79.

Osler, C.L., "Information, Order Flow, and High-Frequency Exchange Rate Dynamics," Mimeo, Federal Reserve Bank of New York (2001).

Osler, C.L., "Support for Resistance: Technical Analysis and Intraday Exchange Rates," Federal Reserve Bank of New York *Economic Policy Review* 6 (July 2000): 53-65.

Osler, C. L., "Short-Term Speculators and the Puzzling Behavior of Exchange Rates," *Journal of International Economics* 43 (June 1998): 37-58.

Odean, Terrance, "Are investors reluctant to realize their losses?," *Journal of Finance* 53 (1998): 1775-1798.

Pring, Martin J., *Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points*, Third Edition (New York: McGraw-Hill, Inc., 1985).

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

Rossi, Barbara, "Testing Out-of-Sample Predictive Ability With High Persistence: An Application to Models of Nominal Exchange-Rate Determination," Princeton University Mimeo (April 2000).

Shareen, Joshi, Jeffrey Parker, and Mark Bedau, 1998, "A Prisoner's Dilemma Causes Technical Trading," Mimeo, Santa Fe Institute.

Shefrin, Hershey, and Meir Statman, "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *The Journal of Finance* 40 (July 1985): 777-791.

Stevenson, Richard A., and Bear, Robert M., "Commodity Futures: Trends or Random Walks," *Journal of Finance* 25 (1970): 65-81.

Table I: Requested Execution Rates Near Round Numbers

The table summarizes asymmetries in the distribution of requested execution rates for stop-loss buy and stop-loss sell 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, and sum them.

The marginal significance represents the likelihood of observing the actual difference between orders above and below the round number, under the following null hypothesis: Agents have overall preferences for round numbers consistent with those in Figures 2A-D, with ranked preferences for rates ending in (0), (5), (2,3,7,8), and (1,4,6,9), but their preferences do not differ significantly within categories of round numbers. For each type of order, 2,000 sums of 10 frequencies were generated, where each set included one observed frequency from rates ending in 0, one observed frequency from rates ending in 5, four observed frequencies from rates ending in 2, 3, 7, or 8, and four observed frequencies from rates ending in 1, 4, 6, or 9. From these, 1,000 differences were calculated and ranked. This ranking provided the distribution for calculating marginal significance levels.

	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
Difference	7.0	-4.9	1.2	-1.5
Marg. Sig.	0.029	0.064	0.363	0.356
Around 50				
40-49	17.3	8.0	7.9	6.6
51-60	6.3	16.9	7.8	7.7
Difference	11.0	-8.9	-0.1	1.2
Marg. Sig.	0.002	0.004	0.485	0.365

Table II: Background Information on Simulated Exchange Rates

The table shows baseline information concerning the 10 simulated exchange rates, where each simulated rate began at one of the following values: 1.6000, 1.6011, 1.6022, 1.6033, etc. The average (absolute) exchange-rate change is measured in points. The share of conditional orders in all orders is measured as the sum of the absolute value of executed stop-loss orders plus the absolute value of executed take-profit orders, as a percent of that quantity plus the absolute value of pure random orders plus the absolute value of arbitrage orders. The number of 10-day intervals refers to the number of 10-day equivalent intervals in which the exchange rate actually hit round numbers. The maximum possible number of such intervals was 347.

		00 and 50	00 Only
Number Relevant 10-Day Intervals	Avg.	310	115
	Min	305	107
	Max	314	127
Number of Hits at Round Numbers	Avg.	2,761	544
	Min	2,607	452
	Max	3,069	601
Share of Conditional Orders	Avg.	74.4	74.4
(In Percent)	Min	74.0	74.0
	Max	74.9	74.9

Table III: Round Numbers As Exchange Rate Turning Points.

The table shows a comparison of the bounce frequency at round numbers (BF_{RN}) with bounce frequencies at arbitrary numbers (BF_{AN}) for 10 simulated exchange rate series, with each series beginning at one of the following values: 1.6000, 1.6011, 1.6022, etc. The bounce frequency is the ratio of the number of times the exchange rate trend is reversed after hitting a level to the number of times it hits such levels. In the one- (two-) period results, the exchange rate is determined to have bounced if it has risen in the one (two) period(s) following the period in which it "hit" a level. The exchange rate was defined to "hit" a level if it came within 0.01 percent of it. Round numbers were defined as exchange rates ending in 00 and 50 or, more narrowly, exchange rates ending in just 00.

The first row gives the number of simulated exchange rate series in which the bounce frequency for round numbers exceeds the bounce frequency at arbitrary numbers. The next four rows give the number of results significant at the 0.01%, 1%, 5%, and 10% levels, respectively. The last row gives the average gap between bounce frequencies for round numbers and bounce frequencies for arbitrary numbers.

Round Numbers Defined As:		00 and 50	00 Only
One-Period Results			
$BF_{RN} > BF_{AN}$ (Number of Series)		10	10
Statistical Significance (Number of Series)	0.01%	10	1
	1%	--	3
	5%	--	2
	10%	--	--
Average $BF_{RN} - BF_{AN}$ (Percentage Points)		3.47	2.15
Two-Period Results			
$BF_{RN} > BF_{AN}$ (Number of Series)		10	7
Statistical Significance (Number of Series)	0.01%	5	1
	1%	5	1
	5%	--	2
	10%	--	1
Average $BF_{RN} - BF_{AN}$ (Percentage Points)		3.41	1.80

Table IV: Strong Trends After Crossing Round Numbers

The table compares the average exchange-rate move after hitting a round number (MV_{RN}) with the average move after hitting an arbitrary number (MV_{AN}), in each case conditional on the rate not bouncing over the next one or two periods. The comparison was made for 10 simulated exchange rate series, with each series beginning at one of the following values: 1.6000, 1.6011, 1.6022, etc. In the one- (two-) period results, the exchange rate is determined to have bounced if it has risen in the one (two) period(s) following the period in which it "hit" a level. The exchange rate was defined to "hit" a level if it came within 0.01 percent of it. Round numbers were defined as exchange rates ending in 00 and 50 or, more narrowly, exchange rates ending in just 00.

The first row gives the number of simulated exchange rate series in which the average move after hitting a round number, conditional on a failure to bounce, exceeds the corresponding average move for arbitrary numbers. The next three rows give the number of results significant at the 0.01%, 1%, and 5% levels, respectively. The penultimate row gives the average difference between MV_{RN} and MV_{AN} . The last row gives the average size of MV_{AN} .

Round Numbers Defined As:		00 and 50	00 Only
One-Period Results			
$MV_{RN} > MV_{AN}$ (Number of Series)		10	10
Statistical Significance (Number of Series)	0.01%	10	1
	1%	--	7
	5%	--	2
Average $MV_{RN} - MV_{AN}$ (Points)		4.2	7.4
Average MV_{AN} (Points)		4.7	8.1
Two-Period Results			
$MV_{RN} > MV_{AN}$ (Number of Series)		10	10
Statistical Significance (Number of Series)	0.01%	5	
	1%	5	10
	5%	--	--
Average $MV_{RN} - MV_{AN}$ (Points)		3.7	6.1
Average MV_{AN} (Points)		4.5	8.3

Figure 1A: Requested Execution Rates by Currency, Distribution of Final Two Digits

(orders not weighted by value)

The figures examine all executed stop-loss and take-profit orders for these currency pairs for a major dealing bank over September 1, 1999 through April 11, 2000. They show the frequency distribution of these orders' requested execution rates, looking only at the final (right-hand) two digits. Orders are aggregated across order types (stop-loss, take-profit), but disaggregated across the three currencies (dollar-yen, dollar-pound, and euro-dollar).




Figure 1B: Requested Execution Rates by Currency, Distribution of Final Two Digits (orders weighted by value)

The figures examine all executed stop-loss and take-profit orders for these currency pairs for a major dealing bank over September 1, 1999 through April 11, 2000. They show the frequency distribution of these orders' requested execution rates, looking only at the final (right-hand) two digits. Orders are aggregated across order types (stop-loss, take-profit), but disaggregated across the three currencies (dollar-yen, dollar-pound, and euro-dollar).






Figure 2: Execution Rates by Order Type, Distribution of Final Two Digits

The figures examine all executed stop-loss and take-profit orders for these currency pairs for a major dealing bank over September 1999 through April 11, 2000. They show the frequency distribution of these orders' requested execution rates, looking only at the final (right-hand) two digits. Orders are aggregated across three currencies (dollar-yen, dollar-pound, and euro-dollar) but disaggregated according to type of order.


(A) Stop-loss




(C) Take-profit



(B) Stop-loss



(D) Take-profit



NOTES

¹ To illustrate the way support and resistance levels have become part of the lexicon of foreign exchange markets, consider this recent internet briefing on exchange rates, which began this way: "The EUR/USD began the European session at 0.9005, but quickly found support with reports of US fund buying of the EUR/JPY cross. This helped lift the pairing ..." (Currency Network.com, March 20, 2001).

² Henceforth, I will drop the word "predictable" for brevity.

³ Indeed, causation may work the other way, since profitable exchange-rate moves generally precede intervention episodes, and may trigger them.

⁴ The term "conditional orders" is my own. Neiderhoffer (1965) labels both stop-loss and take-profit orders as "stops." Since it might be confusing to refer to both stop-loss and take-profit orders as "stops," I use the descriptor "conditional" to distinguish these orders jointly from market (or "at-best") orders.

⁵ Technically, foreign exchange dealers should not undertake a foreign exchange transaction until there has been a trade at the specified price, and the dealer should carry out the desired transaction at the best possible price thereafter. By convention, take-profit orders are filled exactly at the specified rate, while the rates at which stop-loss orders are filled may differ.

⁶ One paper does analyze the implications of stop-loss trading for exchange rates is Krugman and Miller (1993). DeLong et al. (1990) analyze the impact of stop-loss-type trading strategies for financial markets generally.

⁷ Options dealers typically open a hedge position in the spot or forward markets at the same time they open an options position, in order to neutralize their directional risk. The size of the perfect hedge varies over time, because the sensitivity of options prices varies with current rates. Because transactions costs are non-zero, options dealers typically adjust their hedges discretely. Everything else equal, the appropriate time for adjusting a hedge can be mapped to the level of the underlying spot or forward rate. This, combined with the fact that options dealers do not follow the spot market attentively, provides a natural motivation for the use of conditional orders in adjusting hedges for standard options. Likewise, certain "exotic" currency options, such as knock-out options (for which the exposure disappears if the exchange rate reaches a certain level), are best hedged with conditional orders (typically huge orders that are left outstanding for months on end).

⁸ Note that the disposition effect does not explain the use of take-profit orders. One might wonder why placing a stop-loss order would have any self-discipline effect, since a speculator could always call and cancel the order before it was executed. The answer to this question presumably lies in the behavioral realm.

⁹ Some banks keep a single office open all the time, and the orders are simply passed across shifts of dealers, but this practice is not common.

¹⁰ The order execution time in the database is the time at which the execution information is entered into the electronic order book. In order to ensure that orders are not mistakenly executed more than once, this entry is typically within just a few minutes after the actual execution time.

¹¹ For example, if the execution rate was \$1.4321/£, the last two significant digits were counted as 21. The last two digits were also 21 for execution rates of ¥143.21/\$ or \$0.9821/0.

¹² I use requested rates for executed orders, rather than for all orders, to ensure that I only describe tendencies that would actually have affected prices. However, the distribution for all orders does not differ greatly from that presented in Figure 1.

¹³ The Anderson-Darling test statistics are 65.8 for dollar-yen, 137.6 for dollar-pound, and 63.5 for euro-dollar. The critical value for the 0.01 percent significance level is 8.0.

¹⁴ Clustering of actual prices has also been observed in other markets, including the London gold market (Ball et al. 1985) and the market for corn and soybean futures (Stevenson and Bear 1970). However, these results are not directly applicable to the analysis of this paper, since they concern prices rather than orders, and those markets function quite differently from the foreign exchange markets. Clustering is also evident in bid-ask spreads in the FX market (Goodhart and Curcio 1991).

¹⁵ For the euro, which had been expected to be quoted at around 1.0 but has since tended to be quoted below that figure, the convention is to quote the rate at the fifth significant digit of the rate as it was expected to be quoted.

¹⁶ Note that this analysis implicitly violates the results of the previous analysis of take-profit orders, by assuming that there is no stronger preference for rates ending in 00 than other rates ending in 0. However, these results are only trivially different if rates ending in 00 are excluded from the set of rates ending in 0. Under this alternative approach, marginal significance levels for Table 1 are 0.027 and 0.003 for stop-loss buy orders, 0.067 and 0.003 for stop-loss sell orders, 0.337 and 0.504 for take-profit buy orders, and 0.323 and 0.439 for take-profit sell orders.

¹⁷ For completeness, I also assess the significance of these asymmetries under the null hypothesis that the frequencies in Figures 2A-D are generated from uniform distributions over [0,1]. I first create 2,000 sums of 10 random numbers taken from such a distribution. I then take 1,000 differences between these sums, and rank them. The maximum observed difference in this sample derived from the uniform distribution is 4.36 (in absolute value). Since all four of the relevant differences for stop-loss orders exceed this figure, those differences are statistically significant at better than the 0.001 percent level under this null. By contrast, all the take-profit asymmetries are not significant by this standard; their marginal significance levels range upward from 0.15.

¹⁸ Given the paucity of high-frequency information about currency market turnover, it would be difficult, and perhaps impossible, to justify empirically my assertion of rough equivalence between one simulated period and 15 minutes of trading time, and I make no attempt to do so. The perception of rough equivalence is drawn solely from my own experience studying these markets.

¹⁹ $347 \text{ intervals} = 250,000 / (18 * 10 * (60/15))$, where 250,000 is the number of observations, 18 is the number of active trading hours per day, 10 is the number of trading days per intervals, and (60/15) is the number of 15-minute intervals per hour.

²⁰ In theory, I could have used the artificial support and resistance levels to estimate the entire distribution of bounce frequencies and other factors under the null. However, the variance and other higher moments of that distribution are sensitive to the number of hits, which is not fully under my control. By simply evaluating, interval by interval, whether the bounce frequency for round numbers exceeds the average bounce frequency for arbitrary numbers, I eliminate the sensitivity of the results to the number of hits.

²¹ These arbitrary numbers were calculated as follows: I calculated the range of simulated rates for a given interval i , $\text{range}(i)$, and also found the maximum rate over that interval, $\text{max}(i)$. I then drew random numbers from a uniform distribution over [0,1], ϖ . Each arbitrary support or resistance level was initially calculated as $\text{sup/res} = \text{max}(i) - \varpi \text{range}(i)$; these preliminary levels were then rounded off to the correct number of significant digits.

²² The simulations could have been set up to eliminate one of these four parameters, without loss of generality. However, it was computationally simpler to include all four.