

STOP-LOSS ORDERS AND PRICE CASCADES IN CURRENCY MARKETS

C. L. Osler

Abstract

This paper provides empirical evidence that currency stop-loss orders contribute to rapid, self-reinforcing price movements, or "price cascades." Stop-loss orders, which instruct a dealer to buy (sell) a certain amount of currency at the market price when its price rises (falls) to a pre-specified level, are a natural source of positive-feedback trading. Theoretical research on the 1987 stock market crash suggests that stop-loss orders can cause price discontinuities, which would manifest themselves as price cascades. Empirical analysis of high-frequency exchange rate movements suggests the following: (i) Exchange rate trends are unusually rapid when rates reach stop-loss order cluster points; (ii) The response to stop-loss orders is larger than the response to take-profit orders, which generate negative-feedback trading and are therefore not likely to contribute to price cascades; (iii) The response to stop-loss orders lasts longer than the response to take-profit orders. Most results are statistically significant for hours. Together, these results indicate that stop-loss orders propagate trends and are sometimes triggered in waves, contributing to price cascades. The paper also provides evidence that exchange rates respond to non-informative order flow. (Key words: Positive-feedback, Stop-loss, Order Flow, Portfolio Insurance, Information, Exchange Rates, High-frequency, Currency Market Microstructure) (JEL numbers: F1, G3.)

April 2002

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

STOP-LOSS ORDERS AND PRICE CASCADES IN CURRENCY MARKETS

On October 7, 1998, the dollar-yen exchange rate fell 11 percent. On March 7, 2002, the rate dropped over 3 percent. These moves, which dwarf the 0.7 percent standard deviation of daily returns in dollar-yen since 1990, are symptomatic of a broader phenomenon: the well-known "fat tails" of exchange rate returns. Since 1990, dollar-yen returns above four standard deviations have occurred 85 times more frequently than predicted by the normal distribution; under a normal distribution, daily returns above 3 percent would occur fewer than once every 100 years.¹

Dramatic exchange rate moves are as puzzling to economists as they are disruptive to market participants. According to standard exchange rate models (e.g., Dornbusch (1981)), the main force behind them must be news. Yet Cai et al. (2002) find that the arrival of news was of only "secondary importance" for extraordinary yen volatility throughout 1998 (Conclusion, p. x). Likewise, Evans (2001) finds that "public news is rarely the predominant source of exchange rate movements over *any* horizon" (p. 1, italics in the original). Of greater importance, these authors suggest, is order flow.

Researchers have also turned to order flow to account for the stock market crash of 1987, another dramatic price move that cannot be explained by news (Shiller (1989)). Theoretical analyses have highlighted an important role for portfolio insurance and stop-loss orders, two trading schemes in which sell (buy) orders are triggered by a price decline (rise) to a pre-specified level. Because these schemes involve price contingent, positive feedback trading, they can contribute to market discontinuities—that is, crashes—under imperfect information (Genotte

and Leland (1990); Easley and O'Hara (1991); Jacklin et al. (1992)).² In the most commonly cited scenario, a price decline from any source triggers portfolio insurance sales, which further depress prices, which triggers additional portfolio insurance sales, etc. This type of self-reinforcing price dynamic will be referred to here as a "price cascade." Since information about portfolio insurance and stop-loss orders is not public, rational trading by uninformed agents could intensify such a price cascade (Genotte and Leland (1990); Easley and O'Hara (1991)).

The present paper uses this theoretical analysis of the 1987 stock market crash to help explain the high frequency of large exchange rate moves. In what is, to the author's knowledge, the first empirical attempt to examine the effects of price-contingent positive-feedback trading, the paper asks, Do stop-loss orders contribute to price cascades in currency markets? The evidence presented here suggests that the answer is Yes.

This idea is hardly new. Among market participants it is common knowledge that stop-loss orders contribute to price cascades. With regard to the March 7, 2002 drop in dollar-yen, for example, Deutsche Bank noted the following: "Without any news to trigger the move, Japanese accounts aggressively sold USD/JPY, which in turn triggered successive waves of stop-loss orders. The first wave of stop-loss selling occurred on the break of ¥130.50 and then again on the break of ¥130. Once below ¥129.80, USD/JPY fell within seconds to ¥129.40 . . ." (DB (2002)). Likewise, the currency market newsletter to which this author subscribes reported that "stops were triggered," or some equivalent, on at least 16 of the approximately 190 trading days from December 2000 through August 2001, or at least once every two weeks.³ Market lore suggests that participants sometimes intentionally trigger a series of stop-loss orders, and the activity has

¹ The normal distribution used for comparison here has the same mean and standard deviation as actual returns.

² Positive-feedback trading involves sales (purchases) following price declines (rises).

³ Since currency commentators are not active market traders, and their knowledge of the existence and execution of stop-loss orders is at best second-hand, it is possible that large stop-loss orders were triggered more frequently than this suggests.

its own name: "running the stops." Major movements when stop-loss orders are triggered are characterized by market participants as extremely rapid and "gappy," meaning that individual prices are skipped as the rate moves from one price level to another.

This paper attempts to show that this phenomenon exists in currency markets, where stop-loss orders are commonplace. To identify when such orders are executed, I turn to evidence that stop-loss orders cluster in predictable ways (Osler (2001)): stop-loss sell orders cluster just below round numbers; stop-loss buy orders cluster just above round numbers.⁴ The empirical analysis therefore focuses on high-frequency exchange rate behavior near round numbers. The tests rely on minute-by-minute exchange rate quotes for dollar-mark, dollar-yen, and dollar-U.K. pound during New York trading hours from January 1996 through April 1998. The statistical methodology is a variant of the bootstrap (Efron (1979), (1982)).

The analysis first shows that exchange rates tend to move rapidly after reaching levels where stop-loss orders cluster.⁵ This indicates that a trend can be prolonged by the execution of some stop-loss orders triggered by that trend, consistent with the paper's main hypothesis. However, this result need not demonstrate that stop-loss orders are sometimes executed in "waves," as described by DeutscheBank. That is, it need not indicate that the execution of stop-loss orders at one level sometimes propels rates to new levels, thereby triggering more stop-loss orders.

To evaluate whether stop-loss orders are actually triggered in waves, the paper undertakes two tests in which exchange rate behavior after reaching stop-loss orders is compared with behavior after reaching other orders, called "take-profits." Take-profit orders instruct dealers to buy (sell) a certain amount of currency if the rate falls (rises) to a particular level.

⁴ These stop-loss orders are, to my knowledge, the first positive-feedback trades at the individual trade level available to researchers.

Take-profit orders differ from stop-loss orders in that they generate **negative** feedback trading, as a result of which take-profit orders should not contribute to price cascades and would never be triggered in waves. Furthermore, take-profit orders cluster **on**, rather than near, round numbers (Osler (2001)). The two tests both exploit the observation that, if stop-loss orders are sometimes triggered in waves, the response to stop-loss orders should be larger, and should last longer, than the response to take-profit orders. Results support these implications.

The paper also evaluates the average response of exchange rates to take-profit orders, *per se*. Results indicate that rates reverse course relatively frequently upon reaching clusters of take-profit orders., which is relevant to a separate question within the currency market microstructure literature: Why do exchange rates respond to order flow? There is general agreement that "information effects" are important, meaning that exchange rates respond to the information content of order flow (Lyons (1995), Evans and Lyons (2001), Payne (1999) (informed traders); Bjornes and Rime (2000), Rime (2000)). There is no consensus, however, about whether exchange rates also respond to non-informative order flow. If not, then the average effect of take-profit orders should be zero, because the average surprise component of those orders should also be zero and the response to positive and negative surprises should be symmetric. The observation that the average effect of take-profit order clusters is not zero suggests that exchange rates also respond to non-informative order flow. In this, exchange rates are apparently similar to stock and bond prices, which also seem to respond to non-informative order flow (Shleifer (1986); Harris and Gurel (1986); Messod and Whaley (1996); Lynch and Mendenhall (1997); Simon (1991)).

⁵ The paper will refer interchangeably to "stop-loss order clusters" or, more accurately, "stop-loss dominated orders clusters."

If stop-loss trading in currency markets contributes to price cascades, then it also contributes to the high frequency of large moves relative to that predicted by the normal distribution, a property known familiarly as "fat tails." Fat tails, in turn, contribute to "excess kurtosis," or kurtosis higher than the value of three associated with the normal distribution.⁶ Existing research on excess kurtosis in currency markets has primarily focused on its statistical origins (Westerfield (1977); Andersen et al. (20...)). In addressing instead the economic origins of excess kurtosis, this paper joins Amihud and Mendelson (1987), which shows that the distribution of daily NYSE returns varies with the trading process. It also joins LeBaron (2001), which shows that, for a given trading process, kurtosis is positively related to agent heterogeneity.

Though this paper interprets unusual exchange rate behavior near round numbers as the response to clusters of stop-loss and take-profit orders, statistical analysis cannot prove that the connection is causal. To evaluate whether some other factor might explain the unusual behavior, the paper closely examines both alternative factors suggested in the literature: central bank intervention and chaotic exchange rate processes. Both alternatives seem unlikely to explain the behaviors. The paper also shows that exchange rates seem to respond more strongly to price-contingent order flow when liquidity is low, and that there is no clear asymmetry between responses to buy and sell orders. Finally, the paper documents patterns in the placement of large stop-loss and take-profit orders that foster the likelihood of price cascades.

The paper has four sections and a conclusion. Section I presents the data and discusses excess kurtosis in exchange rate returns. Section II provides evidence that exchange rates trend rapidly after reaching stop-loss clusters. Section III provides evidence that rates reverse course

⁶ "Excess kurtosis" is defined as kurtosis above the level of three, where three is the kurtosis of normally distributed variables. Kurtosis arises from two sources: (i) fat tails, or equivalently a relatively large probability mass at the tails

relatively frequently when they reach take-profit clusters, and interprets this as indicating that rates respond to non-informative order flow. Section IV provides evidence that the average response to stop-loss clusters is larger, and lasts longer, than the average response to take-profit clusters. Section V evaluates possible alternative sources for the unusual exchange rate behaviors near round numbers. Section VI offers concluding remarks.

I. BACKGROUND

This section discusses the clustering tendencies of stop-loss and take-profit orders, describes the exchange rate data used in the empirical analysis, and reviews the fat-tailed property of exchange rate returns.

A. *Order Clustering*

Osler (2002) documents the clustering tendencies of stop-loss and take-profit orders. The data were taken from the complete order book of a major foreign exchange dealing bank during August 1, 1999 through April 11, 2000. They include 9,655 orders, with aggregate face value over \$55 billion, in three currency pairs: dollar-yen, dollar-U.K. pound, and euro-dollar. Further information about these orders is shown in Tables I and II.

Both stop-loss orders and take-profit orders tend to cluster at round numbers (Figure 1).⁷ Almost 10 percent of all such orders are placed at rates ending in 00 (such as ¥123.00/\$ or \$1.4300/£); on average, about 3 percent of orders are placed at each of the other rates ending in 0

of a distribution; and (ii) a relatively large probability mass near the mean of a distribution. The normal is a natural benchmark because most models of financial markets predict normally distributed returns.

⁷ These data are likely to be representative of the market-wide population of orders because the bank in question deals with the full spectrum of financial and non-financial customers and is active in the interbank market.

(such as ¥123.20/\$ or \$1.4370/£); about 2 percent of orders are placed at each of the rates ending in 5.⁸

Nonetheless, there are striking differences between the clustering patterns of the two order types. These can be observed in Figure 2, which disaggregates orders according to type (take-profit or stop-loss) and direction (buy or sell). Table III summarizes two critical asymmetries.⁹ First, stop-loss buy orders cluster just above round numbers (meaning specifically numbers ending in 00 or 50), and stop-loss sell orders tend to cluster at rates just below round numbers. For example, 14.3 percent of executed stop-loss buy orders have requested execution rates ending in the range [01,10], while only 6.9 percent of those orders have requested execution rates ending in the range [90,99]. Second, take-profit orders have a stronger tendency to cluster at rates ending in 00 than stop-loss orders. About 9.9 percent of executed take-profit orders (weighted by value) have requested execution rates exactly at rates ending in 00; the corresponding figure for stop-loss orders is 3.8 percent.

The possibility that clusters of price-contingent orders could affect exchange rates, though widely familiar to market practitioners, is not implied by traditional models of currency markets. It is, however, consistent with the more recent microstructure approach to exchange rates (Lyons (2001)), in which a high-frequency relationship between order flow and exchange rates plays a central role. The constituent elements of this relationship are discussed in depth in Section III. The analytical relationship between aggregate price-contingent order flow and the distribution frequencies in Figure 2 is derived in the Appendix.

⁸ In currency markets, the concept of a round number is predicated on two quotation conventions. First, in the wholesale market exchange rates for a given currency pair are universally quoted with the same currency in the denominator: For dollar-mark and dollar-yen that currency is the dollar, for euro-dollar that currency is the euro. Second, each exchange rate is universally quoted to a fixed number of significant digits: for dollar-mark and dollar-pound, rates are quoted to four decimal places; for dollar-yen rates are quoted to two decimal places.

⁹ Osler (2002) shows that these asymmetries are statistically significant.

B. Exchange Rate Data

The empirical strategy of this paper is to examine exchange rate behavior near round numbers, exploiting these stop-loss and take-profit clustering patterns. The tests use minute-by-minute exchange rate quotes taken from Reuters over January 1996 through April 1998, covering three currency pairs—dollar-mark, dollar-yen, and dollar-U.K. pound—during New York trading hours of 9 a.m. to 4 p.m. The quote for a given minute was taken to be the one posted at or most recently before the exact beginning of the minute.

Though transactions data would be ideal for this purpose, available transactions price series are at most only four months long, and could not provide reliable hypothesis tests. Further, quote levels and transactions prices are generally not widely divergent (Goodhart et al. (1996)), though differences do exist (Danielsson and Payne (1999)). To insulate the results from problems associated with these differences, I compare quote behavior at round numbers with quote behavior at arbitrarily chosen numbers, rather than with any absolute benchmark.

C. Excess Kurtosis and Fat Tails

These quote data exhibit the familiar excess kurtosis of daily exchange rate returns. For all three currencies, kurtosis of daily log changes in mid-rates (taken at 9 am) substantially exceeds the value of three associated with the normal distribution (Table IV). Excess kurtosis can be associated with "fat tails," or a high frequency of large moves, or with a high frequency of tiny moves. Since the present paper is exclusively concerned with fat tails, it is interesting to note that daily returns in excess of three standard deviations occur between 2.5 (dollar-mark) and 4.7 times (dollar-yen) more frequently than would be expected under the normal; changes in excess of four standard deviations occur between 24 (dollar-yen) and 63 times (dollar-pound) more frequently than would be expected under the normal.

Excess kurtosis has also been documented in returns to equity prices (Fama (1965)) and bond prices (Roll (1970)). Attempts to understand the phenomenon have typically approached it from a statistical perspective, rather than an economic perspective. The primary question investigated to date is whether the kurtosis reflects a "mixture of normal distributions" (e.g., Harris (1986); Ane and Geman (2000); Andersen et al. (2000.....)), a mixture of normal and jump processes (Tucker and Pond (1988)), or a single distribution that departs from the normal (e.g., Mandelbrot (1963); Westerfield (1977)). Within this tradition, the strongest evidence tends to support the mixture-of-normals hypothesis.

This paper investigates the economic origins of excess kurtosis, rather than the statistical origins. The literature provides little guidance on this issue. Ane and Geman (2000) provide evidence that trading volume could be the mixing variable in a mixture-of-normals distribution for equity returns. Amihud and Mendelson (1987) show that distribution of daily NYSE returns varies with the trading process. LeBaron (2001) shows that kurtosis is higher in markets with higher heterogeneity among agents. Building on the common denominator of these observations, that excess kurtosis could reflect microstructural factors, this paper suggests that stop-loss orders may be a significant source of excess kurtosis in exchange rates. The central idea is that stop-loss orders contribute to price cascades, as follows: a change in the exchange rate from any source could trigger the execution of stop-loss orders; these would propagate the initial change, thereby triggering the execution of more stop-loss orders, etc. Such a price cascade would be cut short by stabilizing speculation if arbitrage were unlimited and if stop-loss orders were public knowledge. However, stop-loss orders are generally known only to the agents placing them and the dealing bank monitoring them. Asymmetric information models suggest that price cascades could be even more severe than suggested above: rational but uninformed market participants could misinterpret stop-loss trading as the activity of informed investors and trade in parallel with the

stop-loss orders, thereby intensifying price cascades (Genotte and Leland (1990); Easley and O'Hara (1991)).

Though this analysis is based on theoretical research connecting portfolio insurance to the 1987 stock market crash, the cascades to which stop-loss orders might contribute need be neither dramatic nor infrequent. Instead, stop-loss propagated cascades, of varying sizes, might happen as frequently as once per week if rates are particularly volatile.

II. STOP-LOSS ORDERS AND PRICE CASCADES: A FIRST TEST

If stop-loss orders contribute to price cascades, then exchange rate trends should be especially rapid where stop-loss orders cluster. This section provides evidence that this is true. It then uses existing theories of the relationship between order flow and exchange rates to show that such rapid trending is necessary, but may not be sufficient, to support the paper's central hypothesis.

A. Empirical Analysis

To test whether stop-loss-dominated order flow is associated with relatively rapid trends, on average, I look closely at how exchange rates behave after crossing round numbers. Since stop-loss buy (sell) orders cluster just above (below) round numbers, trends after crossing round numbers should be relatively rapid, on average, if these orders occasionally contribute to price cascades. I first find every episode in which the exchange rate reaches a round number, where "reaching" a number is defined as coming within 0.01 percent of it. I partition these episodes into two subsamples, one in which rates cross the round number, and another in which they reverse course.¹⁰ Round numbers are any rates ending in "00" or "50."

¹⁰ When examining round numbers reached by downtrends (uptrends), only bid (ask) prices are used, to avoid complications associated with changing spreads. The analysis of Easley and O'Hara (1992) implies that spreads

The rate is defined to have crossed a number if it is above (below) the number 15 minutes after rising (falling) to the number. For the crossing subsample I calculate the average (log) exchange rate change during the 15 minutes after reaching the rate.¹¹ Movements are signed so that a larger positive number means a faster movement in the direction consistent with the hypothesis that stop-loss orders contribute to price cascades: if the exchange rate reaches a particular number from above (below) as it moves from period $t-1$ level to period t , the corresponding 15-minute move is measured as $s_t - s_{t+15}$ ($s_{t+15} - s_t$), where s_t represents the log of the exchange rate. If the average signed exchange rate change after reaching round numbers tends to exceed the corresponding average for arbitrary numbers, in the rigorous statistical sense described below, I conclude that trends tend to be unusually rapid after rates cross round numbers.

The bootstrap: The bootstrap methodology (Efron (1979), (1982)), which is used throughout this paper, permits researchers to be agnostic about the correct statistical distribution for hypothesis testing. This is advantageous when examining exchange rates, the dynamics of which are not known to fit any parametric distribution.¹² The behavior of exchange rates at 10,000 sets of 30 arbitrarily chosen exchange rates is used to approximate the behavior of exchange rates at round numbers under the null hypothesis that round numbers are not special.

Arbitrary numbers are initially set as follows:

$$A = \max - \alpha \text{ range.}$$

might decrease near round numbers as dealers anticipate a surge in liquidity trades. Hartmann (1999) provides evidence consistent with such an effect in currency markets.

¹¹The interval of 15 minutes is chosen based on the analysis of Yao (1997), who finds that prices reach the level of their sustained response to a trade after roughly 16 minutes. Sixteen minutes is the product of the average 4 minutes between passive trade times and the 5-trade time interval for the maximum price impact.

¹²High-frequency exchange rate returns do not conform to the normal distribution, since they are leptokurtotic. Formal tests of the applicability of distributions other than the normal to intraday exchange-rates have not been applied. Formal tests applied to rates at lower frequencies have been inconclusive (Westerfield (1977); Booth and Glassman (1987); Hsieh (1988)).

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

The sample is divided into 58 intervals of 10 consecutive trading days. For each interval I compare the average signed log exchange-rate change subsequent to crossing round numbers (MV_R) with the average signed change subsequent to crossing the arbitrary numbers (MV_A). Under the null hypothesis, MV_R has an even chance of exceeding MV_A , so each interval can be viewed as a Bernoulli trial with probability 0.5. Under the alternative hypothesis that stop-loss orders contribute to price cascades, MV_R should tend to exceed MV_A . Results for the combined set of trials should conform to the binomial distribution with parameters $(0.5, n)$, where $n \leq 58$ is the number of ten-day intervals in which both round numbers and arbitrary numbers are reached at least once.¹³

Results: For all three currencies, $MV_R > MV_A$, consistent with the alternative hypothesis that rates trend rapidly, on average, after crossing round numbers (Table V). In each case, the null hypothesis that behavior at round numbers is not special can be rejected at a high level of significance. For example, dollar-mark moves an average 0.061 percent during the 15 minutes after crossing a round number, but only 0.054 percent after crossing an arbitrary number. Since the average move is higher for round numbers than for arbitrary numbers in 51 of the 58 relevant 10-day intervals, marginal significance for this test is below 0.001 percent.

¹³ The simplest possible bootstrap test, in which this comparison is undertaken once for the entire sample period, would be statistically unreliable. Under the central limit theorem, the second moment of the distribution used for hypothesis testing depends on the number of times the exchange rate reaches the arbitrary levels. It proved infeasible to ensure in a rigorous way that the exchange rate reached arbitrary numbers roughly the same number of times, on average, that it reached round numbers, so the second moment of the critical distribution could not be appropriately calibrated. The bootstrap test applied here relies only on first moments.

The rapid trending of exchange rates after crossing round numbers certainly indicates that stop-loss orders propagate trends, consistent with the hypothesis that they contribute to self-reinforcing price movements or price cascades. It does not affirmatively indicate, however, that stop-loss orders are sometimes triggered in "waves," as described by Deutsche Bank in their analysis of the events of March 7, 2002. It is possible to provide evidence of such waves, however, by contrasting the behavior of exchange rates after reaching stop-loss orders with their behavior after reaching "take-profit" orders. Section III, which follows, examines what happens after rates reach take-profit orders. This material is interesting in its own right, since it suggests that exchange rates respond to non-informative order flow. Section IV then contrasts exchange rate behavior at take-profit and stop-loss orders.

III. EXCHANGE RATES AND TAKE-PROFIT ORDERS

A take-profit order instructs a dealer to buy (sell) a certain amount of currency if its value rises to a certain level. Since take-profit orders generate price-contingent *negative*-feedback trading, they should not contribute to price cascades. If stop-loss orders are sometimes triggered in waves, then the average response to stop-loss orders should be larger, and should last longer, than the average response to take-profit orders.

A. Average Response to Take-Profit Orders: Theory

Do exchange rates respond to take-profit orders, on average? Before undertaking an empirical analysis of this question, it is useful to pause and examine the implications of existing theoretical research. In the currency microstructure literature, the influence of order flow on exchange rates is commonly modeled as deriving from "information effects," meaning that exchange rates react to the private information conveyed by order flow to dealers (Lyons (1995), Evans and Lyons (2001), Payne (1999) (informed traders); Bjonnes and Rime (2000), Rime

(2000)).¹⁴ If information effects are the only ones active, the average effect of take-profit orders should be zero, according to the following logic: When exchange rates only react to the information content of order flow, then only the unexpected component of such order flow affects rates (Hasbrouck (1988)). These order surprises should be zero, on average, since active market participants are all familiar with the clustering of price-contingent orders. Customers who place orders, currency salespersons who advise customers on placing orders, and dealers who execute orders, can all assess the likely magnitude and direction of price-contingent order flow conditional on the exchange rate's level and direction. Since expectation errors average close to zero, the average effect of order clusters should also be close to zero if positive and negative order surprises have symmetric effects on exchange rates. With take-profit orders, there is no reason to expect any asymmetry between the effects of positive and negative order surprises.¹⁵

Stock and bond prices appear to respond to non-informative order flow. Stock prices rise when shares are listed on the S&P 500 index, an event that appears to be non-informative (Shleifer (1986); Harris and Gurel (1986); Messod and Whaley (1996); Lynch and Mendenhall (1997); see also Kaul et al. (2000)). Yields on specific Treasury bills rise, relative to yields of bills with adjacent maturity, when their supply is announced to increase; there would be no change in relative yields if yields responded solely to the arrival of information (Simon (1991)). The rest of this section tests whether the average exchange rate response to take-profit orders is zero. If not, one might infer that exchange rates respond to non-informative order flow.

¹⁴ In the 1980s many economists accepted the more extreme proposition that non-informative order flow can only affect prices in an inefficient market (e.g., Shleifer (1986)).

¹⁵ There should be an asymmetry between the average effects of positive and negative stop-loss order surprises if these orders are triggered in waves during some price cascades.

B. Average Response to Take-Profit Orders: Evidence

In testing whether there is any average exchange rate response to clusters of orders dominated by take-profits, the null hypothesis will be the same as the previous null: clusters of price-contingent orders have no effect on exchange rates, on average. The alternative hypothesis is that exchange rate reverse course unusually frequently upon reaching clusters of take-profit orders. Using the bootstrap methodology described above, I compare the proportion of times the rate reverses course, or "reversal frequency," at round numbers and at 10,000 sets of 30 arbitrary numbers each. As before, a rate is defined to "reach" a level if it comes within 0.01 percent of it; a rate is defined as "reversing course" if it is not beyond a level 15 minutes after reaching it; the key statistic is the number of 10-day intervals in which the reversal frequency at round numbers exceeds the average reversal frequency for the arbitrary numbers.

Results: For all three currency pairs the overall average reversal frequency for round numbers, RV_R , is higher than the overall average reversal frequency for arbitrary numbers, RV_A , consistent with the alternative hypothesis that rates reverse course frequently where take-profit orders cluster (Table V.B). The null hypothesis can be rejected at high levels of significance for dollar-mark and dollar-yen, though not for dollar-pound. Using dollar-mark to illustrate once again, the average reversal frequency is 59.3 percent at round numbers and 54.8 percent at arbitrary numbers. The reversal frequency at round numbers exceeds the reversal frequency at arbitrary numbers in 46 of the 58 relevant ten-day intervals, so the null hypothesis can be rejected with marginal significance below 0.001 percent.

C. Interpretation: Inventories and Downward-Sloping Demand

If exchange rates do react to take-profit orders, on average, as these results suggest, then one might infer that non-informative order flow affects rates. To explain evidence to the same effect in other financial markets, research has focused on two hypotheses (Harris and Gurel

(1986)). The first hypothesis asserts that the long-run demand curve for financial assets is "downward-sloping," which could be true if agents are risk-averse and the asset has no perfect substitutes. These two conditions seem plausible for currency markets: risk premiums are well documented in forward markets, and it is well-known that the major exchange rates are poorly correlated with each other and with other liquid assets. A downward-sloping demand curve could also exist if arbitrage is limited (Shleifer and Vishny (1987)) and agents are heterogeneous in terms of preferences, tax bases, or views of the future. These conditions are also plausible: the long and familiar list of limits to arbitrage includes many that are relevant to currency markets, such as wealth and credit constraints, position limits, and constraints on portfolio allocations. The heterogeneity of currency market participants, suggested in part by the long list of participants (dealers, importers, exporters, investors, speculators, and hedgers), is also highlighted by research on currency forecasts (Ito (1986); Frankel and Froot (1987); Oberlechner (2001)). The downward-sloping demand curve hypothesis has received substantial support as an explanation for the effect of non-informative order flow on equity prices (Shleifer (1986); Messod and Whaley (1996); Lynch and Mendenhall (1997); see also Kaul et al. (2000)).

The "price pressure" hypothesis suggests, in the context of currency markets, that market participants must be rewarded for taking on risky unwanted inventory (Harris and Gurel (1986)). The potential relevance of this hypothesis is highlighted by the observation that currency dealers are quite averse to holding inventory (Bjonnes and Rime (1999), Lyons (1995), Yao (1997)). Nevertheless, some observers interpret existing evidence as suggesting that inventory effects are unlikely in currency markets. Specifically, it appears that many dealers do not shade their quotes to other dealers in response to unwanted inventory accumulation (see Lyons (1995); (Bjonnes and Rime (2000); Yao (1997)). Direct price shading is an important response to inventory build-up in traditional inventory models, where a monopolistic dealer trades a financial asset (Ho and

Stoll (1981), O'Hara and Oldfield (1986), Garman (1976), Amihud and Mendelson (1980)).

However, inventory imbalances could still affect rates in the absence of direct price shading.

Currency dealers are not monopolists, and can trade with each other. Empirical evidence shows that dealing through brokers is the primary response to inventory accumulation for many currency dealers (Yao (1997); Bjonnes and Rime (2000)).

An alternative "market-wide inventory" effect of inventories on prices can be inferred by combining the tendency of dealers to unload unwanted inventory through brokered trades with the fact that dealers shade prices downward (upward) when they observe a broker deal at the bid (offer) (Goodhart et al. (1996)). Through this connection the market as a whole could adjust prices to unintended inventory accumulation at a single dealer, even if individual dealers do not directly shade their prices.¹⁶

It would be interesting to know whether the downward-sloping demand hypothesis fits currency data better than the inventory hypothesis, or vice versa. However, the issue is not central to this paper. Furthermore, the order clusters examined here do not permit the two hypotheses to be distinguished according to the standard criterion, which is the duration of the effect of permanent order flow on prices.¹⁷

IV. STOP-LOSS ORDERS AND PRICE CASCADES: FURTHER TESTS

So far, this paper's main result suggests that stop-loss orders propagate trends, consistent with the paper's central hypothesis that such orders contribute to price cascades. However, no

¹⁶ Further evidence for the presence of inventory effects comes from bid/ask spreads. Hartmann (1998, 1999) finds that daily spreads in dollar-yen increase with exchange rate volatility, consistent with inventory effects. Bjonnes and Rime (2000) show that bid/ask spreads widen with order size; though they interpret this as an information effect, it is also predicted by pure inventory models such as Ho and Stoll (1981).

¹⁷ Under a downward-sloping demand curve, order flow that is relatively permanent would have permanent effects (Shleifer (1986)), and order flow that is soon reversed would have temporary effects. By contrast, under the price pressure hypothesis all order flow would have only temporary effects (Harris and Gurel (1987); Kraus and Stoll

evidence has been presented so far suggesting that stop-loss orders are actually triggered in waves. This section provides such evidence, using tests motivated by the observation that stop-loss orders can, and take-profit orders cannot, contribute to price cascades. Other things equal, if stop-loss orders are triggered in waves then the average exchange rate response to stop-loss orders should be larger, and should last longer, than the average exchange rate response to take-profit orders. If stop-loss orders are not triggered in waves, then the average response to stop-loss orders should be roughly equivalent to the average response to take-profit orders.¹⁸ This contrast is apt regardless of whether exchange rates respond to non-informative order flow.

A. Relative Size of Response

Is the average exchange rate response to stop-loss orders bigger than the average response to take-profit orders? To examine this question, I calculate two measures: (i) the average excess (log) price movement after rates cross round numbers (MV_{CR}), and (ii) the average excess (log) price movement if rates reverse course at round numbers (MV_{RV}). Both movements are defined to be positive under the assumption that order clusters affect exchange rates as indicated in Sections II and III. Under the null hypothesis that stop-loss and take-profit orders have equal effects on exchange rates, the expected value of $MV_{CR} - MV_{RV}$ is zero. Under the alternative hypothesis that stop-loss orders have a bigger effect, the expected value of $MV_{CR} - MV_{RV}$ is positive. This comparison is undertaken for the 15 minute time interval used earlier and for five longer intervals: thirty minutes, one hour, two hours, one day, and two days.

Results at the 15-minute horizon: For all three currency pairs, $MV_{CR}(15) - MV_{RV}(15)$ is positive, consistent with the alternative hypothesis that stop-loss orders contribute to price

(1972); Hasbrouck (1988)). This distinction is not useful here because it cannot be ascertained whether stop-loss and take-profit order flow is soon reversed.

¹⁸ One important element that may not be equal is the average size of stop-loss and take-profit order clusters. This issue is examined explicitly later.

cascades, and the null can be rejected at a significance level below 0.001 percent (Table VI.A). For example, dollar-yen moves by 0.0130 percent more after crossing a round number than after reversing course at a round number. Since $MV_{CR} - MV_{RV}$ is positive in 46 of the 58 relevant 10-day intervals, the null hypothesis can be rejected with marginal significance below 0.001 percent.

Results at longer horizons: For all three currency pairs $MV_{CR} - MV_{RV}$ remains positive at the thirty-minute, one-hour, and two-hour horizons (Table VI.A). For the yen, the difference remains positive at the one-day horizon. Most of these positive differences are statistically significant. Overall, these results suggest that any excess movement associated with the predictable component of stop-loss orders is fast and intense, an impression that conforms to the general view of the marketplace.¹⁹

One might question the reliability of these results, based on the lack of any control for the absolute size of order clusters. If net price-contingent order flow is typically greater when stop-loss orders dominate than when take-profit orders dominate, one might observe the same result even in if stop-loss orders were not triggered in waves. To investigate this possibility, I use equation (A1) to calculate average net order flow (as a fraction of total executed take-profit order value) at two points: (i) after crossing round numbers, where stop-loss orders should dominate, and (ii) exactly at round numbers, where take-profit orders should dominate. Table VII shows that, if rates are rising (falling), net negative-feedback trading at rates ending in 00 should be equivalent to 9.3 (5.1) percent of all executed take-profit order flow, on average; this is far greater than net positive-feedback trading just beyond rates ending in 00, which averages less than 2.3 percent of all executed take-profit order flow. Thus, the absolute size of order clusters should not be a confounding influence on the results for round numbers ending in 00.

¹⁹ One market participant noted that, when stop-loss orders are involved, price cascades are rapid and "gappy."

For round numbers ending in 50, however, the absolute size of order clusters could be a confounding influence. Consequently, I repeat the analysis above including only round numbers ending in 00. The results are qualitatively unchanged, as shown in Table VI.B. In most cases $MV_{CR} - MV_{RV}$ has the same sign as it did in the previous test, and is of the same order of magnitude. The statistical significance of the tests declines modestly, but this change could reflect the lower power of the tests.²⁰

B. Relative Duration of Response

Does the exchange rate response to stop-loss orders last longer than the response to take-profit orders? To examine this question, I examine how long the special behaviors identified in Sections II and III remain statistically significant. For the crossing subsample I calculate MV_R and MV_A , as defined in Section II.B., at the five horizons beyond 15 minutes. For the reversal subsample I calculate the fraction of reversal episodes in which the exchange rate would still be considered to have reversed course after the longer intervals of time.

Results: For every currency, the tests indicate that the effects of stop-loss order clusters last longer than the effects of take-profit order clusters, consistent with the alternative hypothesis. the tendency to reverse course at round numbers remains statistically significant less than thirty minutes (Table VIII). By contrast, the average exchange-rate movement upon crossing a round number remains statistically significant for at least two hours.

Once again, these conclusions are qualitatively unchanged if we examine only the subsample of round numbers ending in 00. Interestingly, the results vary across subsamples in a manner consistent with differences in the underlying frequency distributions (see Table III). Specifically, the results for trend reversals are stronger for the smaller subsample, consistent with the fact that the dominance of take-profit orders at round numbers is **more** pronounced at levels

²⁰ The power of the tests declines because fewer round numbers are reached when they are defined more narrowly.

ending in 00 than at levels ending in 50. By contrast, results for trend continuations are weaker for the smaller subsample, consistent with the fact that the dominance of stop-loss orders just beyond round numbers is **less** substantial at levels ending in 00 than at levels ending in 50.

C. Extensions

This section examines whether the behavior of exchange rates at round numbers varies according to time of day or to the initiating direction of a trade.

Time of Day: Evidence suggests that exchange rates respond more strongly to aggregate order flow when markets are relatively illiquid (Payne (2001)). Consequently, I run the tests of Sections II and III separately on the New York morning and afternoons.²¹ As shown in Table IX, panel A, the special exchange rate behaviors at round numbers do tend to be more pronounced in the New York afternoon. Most noteworthy in this context is dollar-pound; this rate's tendency to reverse after reaching round numbers is statistically significant for the New York afternoon, though it is not for the New York morning or for the New York trading day as a whole.

Conceivably, this morning/afternoon asymmetry could reflect higher amounts of price-contingent order flow in the afternoons, or intraday differences in the way orders cluster. However, the aggregate face value of open orders is likely to be smaller, not larger, in the New York afternoon (Osler (2002)). Further, clustering patterns should be largely independent of time because orders are typically open for many hours. Thus, the morning/afternoon asymmetry observed here seems consistent with the hypothesis that price-contingent order-flow has a stronger effect on rates during periods of low liquidity.

Direction of Trade: Evidence from equity markets suggests that share purchases affect prices differently than share sales, an asymmetry that could reflect short-sales constraints (Chan and Lakonishok (1993), (1995)). Since there are no short-sales constraints in currency markets, it

should be instructive to learn whether they exhibit a similar asymmetry. Consequently, I run the relevant tests—those of Sections II and III—on buy and sell orders separately, where the commodity currency is taken to be dollars in each case. The absence of any strong or consistent asymmetry (Table IX, panel B) supports the relevance of short-sales constraints for the buy-sell asymmetry in equity prices.

V. FURTHER ANALYSIS

The analysis so far suggests that anticipated stop-loss and take-profit order clusters affect exchange rates, and that stop-loss orders contribute to price cascades and are sometimes triggered in waves. These results suggest at least two additional questions. First, could the unusual behaviors of exchange rates near round numbers reflect something *other* than the rates' responses to clusters of price-contingent orders? Second, can we learn more about the likelihood of price cascades?

A. *Alternative Explanations*

Although it is natural to interpret the unusual exchange rate behaviors documented above as the consequences of order clusters, the statistical analysis is actually silent about the existence of any causal connection. Thus it is worth looking closely at the two other factors suggested by the literature that could conceivably generate those unusual behaviors: central bank intervention and chaotic exchange-rate processes.

Central Bank Intervention: Standard exchange rate models suggest only one potential source for the special behavior of exchange rates at round numbers: central bank intervention.²² However, central bank intervention at round numbers must fit a very specific pattern to generate

²¹ The 9 a.m. to 4 p.m. period was split at 12:30 for these tests.

²² No other source is suggested by these models because they assume that only the *log* exchange rate matters for economic outcomes, in which case the concept of a "round number" cannot be defined.

the special behaviors: When rates arrive at round numbers, central banks must initially trade "against the wind;" when rates cross round numbers, central banks must switch and trade "with the wind." Such behavior seems unlikely to fulfill the standard motivations for central bank intervention, such as calming disorderly markets, smoothing exchange rate paths, or targeting exchange rates.

Empirical evidence also indicates that central bank intervention is unlikely to have generated special behavior near round numbers during the sample period. For dollar-mark and dollar-pound there were no reports of intervention at all during the sample period in *The Wall Street Journal* and the *Financial Times*. Subsequent official reports confirm that the Federal Reserve, at least, did not intervene in these currencies. The Bank of Japan was reported to intervene on a number of occasions during the sample period, and the possibility of intervention was occasionally discussed publicly by Japanese officials. I run the tests of Sections II and III over the subset of 20 months in which intervention was neither reported nor discussed publicly by government officials.²³ As shown in Table X, the yen's special behaviors at round numbers differ only slightly between the full and the restricted samples. Taken as a whole, the evidence indicates that central bank intervention is not a primary source of unusual exchange rate behaviors at round numbers.

Chaos: The exchange rate special behaviors at round numbers identified here have long been familiar to technical analysts (Edwards and Magee (1997); Osler (2002)).²⁴ Clyde and Osler (1997) show that technical trading signals based on visual price patterns, like the "head-and-shoulders," might derive their forecasting power from an underlying chaotic structure in the

²³ The excluded months were 2/96; 4-6/97; 12/97; 1/98; and 3-4/98.

²⁴ Technical analysts try to predict financial prices using information limited to past prices and volumes.

financial price series.²⁵ There are two strong reasons why this hypothesis is unlikely to explain the predictive power of round numbers. First, the conditional distribution of a series generated by a chaotic process is of necessity independent of round numbers. Second, empirical tests have generally provided little support for the hypothesis that exchange rates are chaotic (Hsieh (1989); Cecen and Erkal (1996)).

B. The Likelihood of Price Cascades

While useful, the results presented above do not bring us very close to the actual phenomenon of price cascades. Even if price cascades actually did take place during the sample period, they probably did not occur frequently: market participants suggest informally that large, noticeable cascades happen at most once per week, on average, though smaller cascades—in which the market moves a few points or less—could occur more frequently. Since rates crossed round numbers on the order of five times per day in the sample period, the "average exchange rate move after crossing round numbers" is probably based primarily on episodes without price cascades. Because of data limitations—in particular, insufficient data on large, executed stop-loss orders—we cannot yet look directly at price cascades. However, it is possible to investigate whether conditions in currency markets make large price cascades more likely.

Price cascades would be most likely to occur when individual stop-loss orders are unusually large and clustered together, and when offsetting take-profit orders are not large and not clustered. To investigate whether currency markets foster outcomes of that sort, I look closely at the distribution of "very large" orders, meaning those with a face value of \$50 million

²⁵ Some chaotic processes will form patterns, called "attractors," when charted in phase space (or equivalently, when x_t is plotted against x_{t-k}). If such series are graphed against time, certain visible patterns might have predictive power.

or more.²⁶ For all such stop-loss (take-profit) orders in the original orders dataset, Figure 3A (3B) plots size against the final two digits of their requested exchange rate.

A comparison of Figures 3A and 3B reveals two important points. First, the largest stop-loss orders are far larger than the largest take-profit orders. For example, the face value of the largest stop-loss order in the dataset, almost \$450 million, is over twice the face value of the largest take-profit order, which falls short of \$200 million. Second, very-large stop-loss orders are tightly clustered near rates ending in 00, and very-large take-profit orders are not so clustered. Over 62 percent of the total face value of very-large stop-loss orders had requested execution rates ending in [90, 100] or [01,09], a figure that far exceeds the corresponding figure of 28 percent for very-large take-profit orders.

The data thus suggest that very large stop-loss orders tend to cluster near round numbers, where they are unlikely to be offset by take-profit order clusters of similar magnitude. Since market participants are generally uninformed about stop-loss orders, currency markets do seem to exhibit the conditions required for unanticipated stop-loss orders to contribute to price cascades.

V. CONCLUSIONS

Important empirical work of recent years indicates that order flow is a critical determinant of high-frequency exchange rate movements (Goodhart et al. 1996; Evans and Lyons (1999); Rime (2000); Lyons (2001); Evans (2001)). This paper broadens our appreciation for the potential contribution of order flow to exchange rate dynamics, by showing that stop-loss orders may contribute to price cascades and, thereby, to excess kurtosis. The empirical evidence

²⁶ This definition of a very large order is based on a general perception among market participants that orders must be at least \$50 million significantly affect rates in liquid markets.

is based on a close analysis of over two years of minute-by-minute exchange rate quotes for three currencies: dollar-mark, dollar-yen, and dollar-U.K. pound.

The paper first provides evidence that exchange rates trend rapidly, on average, when they hit clusters of stop-loss orders, consistent with the paper's central hypothesis that stop-loss orders are one mechanism through which price trends can become self-reinforcing. Additional tests examine whether stop-loss orders are sometimes triggered in waves. These tests exploit the fact that stop-loss orders might contribute to price cascades but take-profit orders will not. Results indicate that the response of exchange rates to stop-loss orders is larger, and lasts longer, than the response to take-profit orders, consistent with the view that stop-loss orders are sometimes triggered in waves. This holds true if the sample is split into morning and afternoon periods, and if buy and sell orders are examined separately.

The paper also provides evidence bearing on a separate question of interest to exchange rate researchers: Are information effects the only source of influence from order flow to exchange rates? or, equivalently, Do exchange rates react to non-informative order flow? This question is interesting because equity and bond prices have been shown to react to non-informative order flow, while evidence for inventory effects—specifically price shading—is mixed in currency markets. The question is examined by testing whether the average exchange rate reaction to take-profit order clusters is zero. Since these clusters can be rationally anticipated by market participants, only the surprise component should affect rates if information effects are the only ones operative; since the surprises themselves average to zero, the average influence of the surprise components should also be zero. Results indicate that exchange rates reverse course relatively frequently upon reaching take-profit order clusters, suggesting that rates react to non-informative order flow.

The paper has additional implications for our understanding of exchange rates. First, the results provide further confirmation that order flow is a strong influence, possibly the dominant influence, on high-frequency exchange rate movements. Second, the paper highlights the significance of an institutional feature of the market that is typically ignored in exchange rate research, the quotation convention. Third, the results indicate that order flow need not be monotonically related to the flow of private information into the market. Fourth, the importance of agent heterogeneity is suggested by the fact that stop-loss orders simultaneously exist at different requested execution rates and that stop-loss and take-profit orders can simultaneously be open at the same requested execution rates.

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

APPENDIX

The relationship between the frequency distribution of stop-loss and take-profit orders and net price-contingent order flow can be articulated as follows: Let $SL^b(kl)$ represent the amount of stop-loss buy orders triggered by the exchange rate's arrival at a level ending in the two-digit combination kl , and let $s^b(kl)$ represent the share of all executed stop-loss buy order value triggered at levels ending in kl ; define $TP^s(kl)$ and $t^s(kl)$ similarly for take-profit sales orders. Suppose rates rise one point between periods $t-1$ and t , reaching a rate ending in kl_t and triggering the execution of stop-loss buy and take-profit sell orders in period t . The effect of these orders on exchange rates depends on net price-contingent orders, $SL^b(kl_t) - TP^s(kl_t)$, the expected value of which is:

$$E[SL^b(kl)_t - TP^s(kl)_t] = [s^b(kl)\rho - t^s(kl)]V .$$

Here, V is the value of the population of executed take-profit sell orders and ρ is the ratio of all executed stop-loss buy order value to all executed take-profit sell order value. Substantial differences in sample frequencies, such as those near round numbers, should correspond to substantial amounts of net price-contingent order flow, since ρ is 72 percent in the original orders dataset (Osler (2001)).

REFERENCES

- Amihud, Yakov, and Haim Mendelson, "Dealership Market: Market-Making With Inventory," *Journal of Financial Economics* 8 (1980): 31-53.
- _____, "Trading Mechanisms and Stock Returns: An Empirical Investigation," *Journal of Finance* 42 (July 1987): 533-553.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, "Modeling and Forecasting Realized Volatility," *NBER Working Paper* 8160 (March 2001).
- Ane, Thierry, and Helyette Geman, "Order Flow, Transaction Clock, and Normality of Asset Returns," *Journal of Finance* 55 (October 2000): 2259-2284.
- Boothe, P., and D. Glassman, "The Statistical Distribution of Exchange Rates," *Journal of International Economics*, 22 (1987): 297-319.
- Brennan, Michael J., and Eduardo S. Schwarz, "Portfolio Insurance and Financial Market Equilibrium," *Journal of Business* 62 (1989): 455-72.
- Cai, Jun, Yan-Leung Cheung, Raymond S.K. Lee, Michael Melvin, "Once-in-a-Generation' Yen Volatility in 1998: Fundamentals, Intervention, and Order Flow," *Journal of International Money and Finance*, (date): Pages.
- Cecen, A.A., and Erkal, C., "Distinguishing Between Stochastic and Deterministic Behavior in High Frequency Foreign Exchange Rate Returns: Can Non-linear Dynamics Help Forecasting?" *International Journal of Forecasting* 12 (1996): 465-73.
- Chan, Louis K.C., and Josef Lakonishok, "Institutional Trades and Intraday Stock Price Behavior," *Journal of Financial Economics* 33 (1993): 173-199.
- _____, "The Behavior of Stock Prices Around Institutional Trades," *Journal of Finance* 50 (September 1995): 1147-1174.
- Clyde, W.C., and C.L. Osler, "Charting: Chaos Theory in Disguise?" *Journal of Futures Markets* 17 (1997): 489-514. (NIL)
- Cooney, John W., Bonnie Van Ness, and Robert Van Ness, "Do Investors Avoid Odd-Eighths Prices? Evidence From NYSE Limit Orders," Mimeo (December 2000).
- Danielsson, Jon, and Richard Payne, "Real Trading Patterns and Prices in Spot Foreign Exchange Markets," Mimeo, London School of Economics (March 1999).
- Diamond, Douglas W., and Robert E. Verrecchia, "Constraints on Short-Selling and Asset Price Adjustment to Private Information," *Journal of Financial Economics* 18 (1987): 277-311.

Deutsche Bank, *Forex Daily*, 7:30 GMT 7 March 2002.

Easley, David, and Maureen O'Hara, "Order Form and Information in Securities Markets," *Journal of Finance* 46 (July 1991): 905-27.

_____, "Time and the Process of Security Price Adjustment," *Journal of Finance* 47 (1992): 577-605.

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) (Forthcoming, *Journal of Finance*).

Fama, Eugene F., "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25 (May 1970): 383-417.

Frankel, Jeffrey A., and Kenneth A. Froot, "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review* 77 (March 1987): 133-53.

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

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

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

Grossman, Sanford J., and Merton H. Miller, "Liquidity and Market Structure," *Journal of Finance* 43 (1988): 617-33.

Grossman, Sanford J., "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies," *Journal of Business* 61 (1988): 275-98.

Harris, Lawrence, "Cross-Security Tests of the Mixture of Distributions Hypothesis," *Journal of Financial and Quantitative Analysis*, 21 (March 1986): 39-46.

Hartmann, Philipp, "Trading Volumes and Transaction Costs in the Foreign Exchange Market: Evidence from daily Dollar-Yen Spot Data," *Journal of Banking and Finance* 23 (1999): 801-824.

_____, "Do Reuters Spreads Reflect Currencies' Differences in Global Trading Activity?" *Journal of International Money and Finance* 17 (1998): 757-784.

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

_____, "Measuring the Information Content of Stock Trades," *Journal of Finance* 47 (March 1991): 179-207.

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

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

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

Jacklin, Charles J., Allan W. Kleidon, Paul Pfleiderer, "Underestimation of Portfolio Insurance and the Crash of October 1987," *Review of Financial Studies* 5 (1992): 35-63.

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

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

LeBaron, Blake, "Empirical Regularities From Interacting Long- and Short-Memory Investors in an Agent-Based Stock Market," *IEEE Transactions on Evolutionary Computation* 5 (October 2001): 442-455.

Lyons, Richard K., *The Microstructure Approach to Exchange Rates*, (M.I.T. Press, Cambridge, MA: 2001).

Miller, Edward, "Risk, Uncertainty, and Divergence of Opinion," *Journal of Finance* 32 (September 1977): 1151-1168.

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

- 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.
- Oberlechner, Thomas, "Evaluation of Currencies in the Foreign Exchange Market: Attitudes and Expectations of Foreign Exchange Traders," *Zeitschrift fur Sozialpsychologie* 3 (2001): 180-188.
- O'Hara, Maureen, and George S. Oldfield, "The Microeconomics of Market Making," *Journal of Financial and Quantitative Analysis* 21 (December 1986): 361-76.
- Osborne, M.F.M., "Periodic Structure in the Brownian Motion of Stock Prices," *Operations Research* 10 (1962): 345-79.
- Osler, C. L., "Currency Orders and Exchange Rate Dynamics: Explaining the Success of Technical Analysis," Federal Reserve Bank of New York *Staff Report* 25 (April 2001(a)).
- Rime, Dagfinn, "Private or Public Information in Foreign Exchange Markets? An Empirical Analysis," Mimeo, April 2000.
- Shiller, Robert, "Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence," in *Market Volatility* (M.I.T. Press, Cambridge Massachusetts: 1989): 379-403.
- Shleifer, Andrei, Robert Vishny, "Limits to Arbitrage," *Journal of Finance* 52 (March 1997): 35-55.
- Simon, David P., "Segmentation in the Treasury Bill Market: Evidence from Cash Management Bills," *Journal of Financial and Quantitative Analysis* 26 (March 1991): 97-108.
- _____, "Further Evidence on Segmentation in the Treasury Bill Market," *Journal of Banking and Finance* 18: 139-151.
- Spiegel, Matthew, and Avanidhar Subrahmanyam, "On Intraday Risk Premia," *Journal of Finance* 50 (March 1995): 319-339.
- Westerfield, Janice Moulton, "An Examination of Foreign Exchange Risk Under Fixed and floating Rate Regimes," *Journal of International Economics* 7 (1977): 181-200.
- Yao, Jian, " "Market Making in the Interbank Foreign Exchange Market," New York University, Salomon Center, Working Paper S/98/03 (November 1997).

Table I: Descriptive Information on Stop-Loss and Take-Profit Orders

The table describes all stop-loss and take-profit orders for three currency pairs—dollar-yen, dollar-U.K. pound, and euro-dollar—placed at a major foreign exchange dealing bank over August 1, 1999 through April 11, 2000. There are 9,655 orders with aggregate value in excess of \$55 billion. The symbol * (**) indicates significance at the five (one) percent level.

	All Orders	Stop-Loss	Take-Profit	Z-Stat., SL vs TP
Number Orders	9,655	3,935	5,720	
Share of Orders	100.0	42.6	57.4	
Size (\$ Mill.):				
Mean	5.78	6.35	5.39	3.63**
Median	3.00	3.24	2.12	10.81**
Dist. to Mkt. (%):				
Mean	0.92	0.91	0.93	0.52
Median	0.53	0.43	0.68	8.08**
Days Open:				
Mean	3.79	3.42	4.05	2.39**
Median	0.56	0.46	0.59	9.72**
Share Executed	27.9	28.3	29.9	1.71*
Share Placed by Customers	72.8	64.1	78.8	15.9**

Table II: Further Descriptive Information of Stop-Loss and Take-Profit Orders

The table lists the sources of all stop-loss and take-profit orders for three currency pairs—dollar-yen, dollar-U.K. pound, and euro-dollar—placed at a major foreign exchange dealing bank over August 1, 1999 through April 11, 2000. There are 9,655 orders with aggregate value in excess of \$55 billion. If orders amounts were not originally measured in dollars, the dollar value represents the original order amount adjusted by the requested execution rate. "Other" customer orders are orders from all sources placed in Tokyo and orders intended to hedge customer options positions. "Internal" orders are those placed by agents within the bank.

	Number of Orders	Percent of Orders	Dollar Value of Orders (\$ Billions)	Percent of Order Value
All Orders	9,655	100.0	55.9	100.0
Customer Orders	7,027	72.8	335.8	64.0
Fin. Inst.	5,357	55.5	27.6	49.4
Non-Fin. Inst.	770	8.0	4.9	8.8
Other	900	9.3	3.3	5.9
Internal	2,628	27.2	20.1	36.0

Table III: Requested Execution Rates Near Round Numbers

The table summarizes asymmetries in the distribution of requested execution rates for stop-loss and take-profit orders near exchange rates with far-right digits 00 or 50. For each entry, I take the percent of executed orders of each order type with requested execution rates ending in the indicated set of two-digit numbers (weighted by value), and sum them. The underlying data comprise 9,655 stop-loss and take-profit orders in three highly-active currency pairs—dollar-yen, dollar-U.K. pound, and euro-dollar—processed by a major foreign exchange dealing bank during August 1, 1999 through April 11, 2000.

	Stop-loss Orders		Take-profit Orders	
	Buy	Sell	Buy	Sell
At 00	2.8	4.8	8.6	11.3
Around 00				
90-99	6.9	10.0	10.9	8.9
01-10	14.3	5.0	12.4	8.6
Difference	7.4	-4.9	1.4	-0.2
Marg. Sig.	0.028	0.063	0.330	0.505
At 50	3.8	4.5	3.9	4.0
Around 50				
40-49	6.3	16.3	7.5	7.4
51-60	18.1	8.0	8.4	6.4
Difference	11.7	-8.3	0.9	-1.0
Marg. Sig.	0.002	0.005	0.387	0.403

Table IV: The High Frequency of Large Exchange Rate Moves

The table illustrates the high frequency of large exchange rate moves using daily exchange rate quotes taken at 9 a.m. New York time over January 2, 1996 through April 30, 1998. "Excess Kurtosis" refers to kurtosis above the level of 3 associated with the normal distribution (with a standard small-sample adjustment). The "Frequency Ratio" shows the ratio of (1) the frequency with which absolute log exchange rate moves exceed a given cutoff in the data and (2) the frequency with which log exchange rate moves would exceed a given cutoff under the normal distribution. A number above unity implies that changes of a given magnitude are observed more frequently in the data than predicted by the normal distribution.

	DEM	JPY	GBP
Excess Kurtosis	1.52	3.36	1.71
Frequency Ratio			
Changes Above 2 Std Dev.	0.89	1.02	1.17
Changes Above 3 Std Dev.	2.50	4.71	4.67
Changes Above 4 Std Dev.	39.66	24.25	63.10

Table V: Exchange Rate Behavior At Round Numbers

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

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

A. Strong Trends After Crossing Round Numbers

	DEM	JPY	GBP
Overall Average, MV_R	6.08	6.89	4.92
Overall Average, MV_A	5.37	5.98	4.59
Intervals $MV_{RN} > MV_A$	51	33	36
Total Intervals	58	57	58
Marg Sig.	(0.000)	(0.000)	(0.024)

B. Frequent Trend Reversals at Round Numbers

	DEM	JPY	GBP
Overall Average, R_R	59.3	60.1	58.1
Overall Average, R_A	54.8	57.3	56.6
Intervals $R_R > R_A$	46	40	33
Total Intervals	58	57	58
Marg Sig.	(0.000)	(0.002)	(0.179)

Table VI: Are the Effects of Stop-Loss Orders Larger Than the Effects of Take-Profit Orders?

The table reports tests of the null hypothesis that exchange rate movements after crossing round numbers are equal in size, on average, to exchange rate movements after failing to cross round numbers. Rows associated with "Excess Movement" report the average excess, relative to arbitrary numbers, of the difference between (1) the (absolute) exchange rate movement conditional on crossing a round number (where crossing is defined as ...), and (2) the (absolute) average exchange rate movement conditional on failing to cross round numbers. Figures are in hundredths of a percent, and are of the same order of magnitude as "points." Bold figures highlight all horizons at which the results are positive, consistent with the hypothesis that stop-loss orders can propagate price cascades. Rows associated with "Statistical Significance" report the marginal significance of the same difference, calculated using the bootstrap algorithm described below. Bold figures highlight all horizons at which the results are statistically significant at the 10 percent level. The underlying data are minute-by-minute exchange rate quotes taken over 9 a.m. to 4 p.m. New York time during January 2, 1996 through April 30, 1998. Round numbers are rates ending in 00, such as DM1.5700/\$, ¥123.00/\$, or \$1.6500/£, or rates ending in 50.

For each 10-trading-day interval, the exchange rate's average behavior at round numbers was compared with its average behavior at 10,000 sets of arbitrary numbers, with 30 numbers in each set. Hitting a level was defined as coming within 0.01 percent of it; reversing was defined as remaining above (below) a support level after 15 minutes. Each interval can be viewed as an independent Bernoulli trial, with probability one half. The final test involved counting the number of intervals in which the exchange rate's behavior at round numbers exceeds its average behavior at arbitrary levels. This number should have a binomial distribution with $n = \text{total relevant intervals}$ and $p = \frac{1}{2}$.

VI.A. All round numbers (levels ending in 00 and 50)

		DEM	JPY	GBP
Excess Movement	15 Minutes	1.12	1.30	0.77
	30 Minutes	0.89	1.34	0.80
	1 Hour	0.57	1.26	0.27
	2 Hours	0.85	2.35	0.35
	1 Day	-5.23	1.27	-2.60
	2 Days	-2.53	-2.13	-2.63
Statistical Significance	15 Minutes	0.000	0.000	0.000
	30 Minutes	0.000	0.001	0.012
	1 Hour	0.024	0.017	0.179
	2 Hours	0.119	0.017	0.347
	1 Day	0.119	0.500	0.074
	2 Days	0.552	0.500	0.179

VI.B Levels ending in 00, only

		DEM	JPY	GBP
Excess Movement	15 Minutes	1.92	1.27	0.63
	30 Minutes	1.06	1.06	0.88
	1 Hour	0.98	0.16	0.40
	2 Hours	1.08	1.61	0.67
	1 Day	-1.12	-2.40	-4.67
	2 Days	-1.64	-9.74	-1.87
Statistical Significance	15 Minutes	0.000	0.004	0.043
	30 Minutes	0.024	0.056	0.179
	1 Hour	0.179	0.396	0.256
	2 Hours	0.256	0.500	0.179
	1 Day	0.448	0.396	0.119
	2 Days	0.448	0.500	0.256

Table VII: Net Order Flow Near Round Numbers

The table shows net price-contingent positive-feedback trading as a percent of all executed take-profit order value. "Just below (above) round numbers" means at all rates ten points or fewer below (above) the number. The underlying data comprise 9,655 stop-loss and take-profit orders in three highly-active currency pairs—dollar-yen, dollar-U.K. pound, and euro-dollar—processed by a major foreign exchange dealing bank during August 1, 1999 through April 11, 2000.

Order flow triggered	Rates ending in 00	Rates ending in 50
At round numbers, rising rates (SLB-TPS)	-9.32	-0.64
At round numbers, falling rates (SLS-TPB)	-5.13	-1.22
Just below round numbers, falling rates (SLS-TPB)	-0.12	1.84
Just above round numbers, rising rates (SLB-TPS)	2.27	4.63

Table VIII: Do the Effects of Stop-Loss Orders Last Longer Than the Effects of Take-Profit Orders?

The table reports tests of the null hypothesis that exchange rates do not behave differently at round numbers against two alternative hypotheses. For the rows under "Trend Reversal," the alternative hypothesis is that, if exchange rates reverse course after hitting a round number, they are relatively unlikely to return to average behavior. For the rows under "Trend Continuation," the alternative hypothesis is that exchange rates will trend relatively rapidly if they fail to reverse course after hitting round numbers. The underlying data are minute-by-minute exchange rate quotes taken over 9 a.m. to 4 p.m. New York time during January 2, 1996 through April 30, 1998. Round numbers are rates ending in 00, such as DM1.5700/\$, ¥123.00/\$, or \$1.6500/£, or rates ending in 50.

For each 10-trading-day interval, the exchange rate's average behavior at round numbers was compared with its average behavior at 10,000 sets of arbitrary numbers, with 30 numbers in each set. Hitting a level was defined as coming within 0.01 percent of it; reversing was defined as remaining above (below) a support level after 15 minutes. Each interval can be viewed as an independent Bernoulli trial, with probability one half. The final test involved counting the number of intervals in which the exchange rate's behavior at round numbers exceeds its average behavior at arbitrary levels. This number should have a binomial distribution with n = total relevant intervals and $p = \frac{1}{2}$.

The figures represent the marginal significance of the results under the null hypothesis. Figures highlighted in bold represent horizons at which the results are statistically significant at the 10 percent level or better.

VIII.A. All round numbers (levels ending in 00 and 50)

		DEM	JPY	GBP
Trend Reversal	30 Minutes	0.119	0.214	0.256
	1 Hour	0.043	0.500	0.347
	2 Hours	0.043	0.396	0.043
	1 Day	0.552	0.500	0.179
	2 Days	0.552	0.092	0.347
Trend Continuation	30 Minutes	0.001	0.001	0.006
	1 Hour	0.001	0.031	0.074
	2 Hours	0.074	0.145	0.448
	1 Day	0.074	0.145	0.179
	2 Days	0.552	0.298	0.074

VIII.B Levels ending in 00, only

		DEM	JPY	GBP
Trend Reversal	30 Minutes	0.004	0.004	0.256
	1 Hour	0.448	0.500	0.119
	2 Hours	0.552	0.500	0.119
	1 Day	0.552	0.396	0.347
	2 Days	0.552	0.298	0.347
Trend Continuation	30 Minutes	0.074	0.017	0.024
	1 Hour	0.043	0.145	0.024
	2 Hours	0.256	0.214	0.179
	1 Day	0.552	0.396	0.179
	2 Days	0.256	0.298	0.179

Table IX: Robustness Tests

The table examines whether the special exchange rate behaviors at round numbers are more pronounced when markets are relatively liquid (the New York morning) or when they are less liquid (New York afternoon). For each test, the null hypothesis that exchange rates do not behave differently at round numbers. There are two alternative hypotheses: (1): exchange rate trends are more frequently reflected at round numbers than at arbitrary numbers. (2): exchange rate trends are stronger after the rate crosses round numbers. (3) trending after round numbers are crossed is stronger than trending after rates reverse at round numbers. The underlying data are minute-by-minute exchange rate quotes during 9 am to 4 p.m. New York time over January 2, 1996 through April 30, 1998, excluding months in which intervention in JPY was either reported in the press or discussed in the press by Japanese financial authorities. The excluded months were February, 1996; April through June, 1997; December, 1997; January, 1998; and March and April, 1998. Round numbers are rates ending in 00, such as ¥123.00/\$, or rates ending in 50.

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

IX.A. Time of Day

Hypothesis:	Currency New York time	DEM		JPY		GBP	
		am	pm	am	pm	am	pm
1 Strong Trends After Crossing Round Numbers	$MV_R - MV_A$ Marginal Significance	0.64 (0.000)	0.79 (0.000)	0.70 (0.000)	1.29 (0.017)	0.22 (0.347)	0.62 (0.003)
2 Trend Reversals at Round Numbers	$RV_R - RV_A$ Marginal Significance	3.03 (0.003)	7.15 (0.000)	1.26 (0.145)	6.93 (0.000)	0.17 (0.347)	4.91 (0.000)
3 Trends After Crossing > Trends After Reversing	Excess Movement Marginal Significance	0.84 (0.000)	1.48 (0.000)	1.06 (0.000)	1.56 (0.001)	0.55 (0.012)	1.42 (0.000)
4 Trends After Crossing Last Longer Than Reversals	Trends Sig. Through Reversal Sig. Through	15 min. 30 min.	30 min. 60 min.	0 min. 30 min.	15 min. 30 min.	0 min. 0 min.	15 min. 120 min.

IX.B: Direction of Trade

Hypothesis:	Currency New York time	DEM		JPY		GBP	
		Sell\$	Buy\$	Sell\$	Buy\$	Sell\$	Buy\$
1 Strong Trends After Crossing Round Numbers	$MV_R - MV_A$ Marginal Significance	0.80 (0.000)	1.34 (0.000)	1.19 (0.000)	0.56 (0.000)	0.22 (0.074)	0.43 (0.012)
2 Trend Reversals at Round Numbers	$RV_R - RV_A$ Marginal Significance	4.28 (0.012)	4.65 (0.012)	3.47 (0.092)	2.10 (0.145)	1.32 (0.552)	1.74 (0.119)

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

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

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

Hypothesis:			JPY
1	Strong Trends After Crossing Round Numbers	$MV_R - MV_A$ Marginal Significance	1.05 (0.000)
2	Trend Reversals at Round Numbers	$RV - RV_A$ Marginal Significance	3.7 (0.001)
3	Trends After Crossing > Trends After Reversing	Excess Movement Marginal Significance	1.23 0.000
4	Trends After Crossing Last Longer Than Reversals	Trends Significant Through Reversal Significant Through	60 min. 30 min.

Figure 1. Requested Execution Rates for Stop-loss and Take-Profit Orders: Distribution of Final Two Digits

The figure shows the distribution of requested execution rates for all placed stop-loss and take-profit orders for three currency pairs—dollar-yen, dollar-pound, and euro-dollar—processed by a major dealing bank between August 1, 1999 through April 11, 2000. Orders are aggregated across the three currencies.






Figure 2. Requested Execution Rates: Distribution of Final Two Digits

The figures show the distribution of requested execution rates for all executed stop-loss and take-profit orders for three currency pairs--dollar-yen, dollar-pound, and euro-dollar—processed by a major dealing bank between August 1, 1999 through April 11, 2000. Orders are weighted by value and aggregated across three currencies, but disaggregated according to type of order.


(A) Stop-loss Buy



(C) Take-profit Buy



(B) Stop-loss Sell



(D) Take-profit Sell




Figure 3A: Very Large Stop-Loss Orders

The figure plots the size and final two significant digits for every stop-loss order with face value of \$50 million or more. The data include stop-loss orders for three currency pairs--dollar-yen, dollar-pound, and euro-dollar--processed by a major dealing bank between August 1, 1999 through April 11, 2000




Figure 3B: Very Large Take-Profit Orders

The figure plots the size and final two significant digits for every take-profit order with face value of \$50 million or more. The data include stop-loss orders for three currency pairs--dollar-yen, dollar-pound, and euro-dollar--processed by a major dealing bank between August 1, 1999 through April 11, 2002

