

# **EXTREME RETURNS WITHOUT NEWS: A Microstructural Explanation**

**Carol Osler<sup>1</sup>**  
Brandeis University

**Tanseli Savaser<sup>2</sup>**  
Williams College

## *Abstract*

What triggers extreme exchange-rate returns? Though news is the source of volatility in standard theoretical models, in reality volatility is often unrelated to news. This paper shows that extreme exchange-rate returns – and, more generally, high kurtosis of returns – are statistically inevitable even in the absence of news. We identify four microstructural sources of return kurtosis in price-contingent order flow: (1) high kurtosis in the distribution of price-contingent order sizes; (2) clustering of price-contingent order executions at certain times of day; (3) clustering of order executions at certain price levels; and (4) the tendency of positive-feedback trading to propagate trends. Using simulations calibrated to price-contingent orders placed at a major foreign exchange dealing bank we show that when each factor operates in isolation, the one that contributes most to kurtosis in returns is kurtosis in the order-size distribution. When the factors operate simultaneously, however, their interactions prove far more important. Extreme returns in the absence of news should be viewed as natural rather than anomalous. (Key words: kurtosis; exchange rates; order flow; high-frequency; microstructure; jump process; value-at-risk; risk management) (JEL codes G1, F3.)

---

<sup>1</sup> Brandeis University, International Business School, Waltham, MA 02454-9110, phone: 781-736-4826, fax: 781-736-2269, email: [cosler@brandeis.edu](mailto:cosler@brandeis.edu).

<sup>2</sup> Williams College, Economics Department, Williamstown, MA 01267, phone: 413-597-4882, fax: 413-597-4045, email: [tsavaser@williams.edu](mailto:tsavaser@williams.edu).

# **EXTREME RETURNS WITHOUT NEWS:**

## **A Microstructural Explanation**

Exchange rates often make huge, abrupt moves. The most dramatic in recent memory was the eleven percent drop in dollar-yen on October 7, 1998, but there are many more examples of startling returns. If we take the normal distribution as our guide, eye-popping exchange-rate moves happen with surprising frequency, as reflected in large values for the kurtosis of returns. The extreme returns in October of 1998 are all the more startling because they do not appear to have been triggered by news (Covrig and Melvin 2005). According to standard macro-based exchange-rate models, extreme returns must accompany news because, in the absence of news, exchange rates move smoothly – by enough to offset interest differentials or to generate a risk premium.

This paper shows that extreme returns are statistically inevitable, even in the absence of news, due to patterns in the placement and execution of price-contingent orders. To arrive at this conclusion we take a micro perspective, rather than a macro perspective, and build our analysis on empirical evidence rather than a model. Our evidence concerns the properties of “stop-loss” orders and “take-profit” orders. Stop-loss orders instruct dealers to buy (sell) when the price is rising (falling); take-profit orders instruct dealers to buy (sell) when the price is falling (rising). Currency dealers manage these orders for customers and for others within their own bank. A large hedge fund, for example, might instruct a dealer to buy €50 million once euro-dollar rises to \$1.35/€ (a stop-loss); or Toyota might instruct a dealer to buy \$25 million after dollar-yen falls to ¥115/\$ (a take-profit). The orders provide efficient execution of trades prompted by technical trading rules, option hedging, and trading for commercial purposes (Osler 2008).

The paper identifies four sources of extreme returns related to these orders: (1) fat tails in the distribution of order sizes; (2) clustering of order executions at certain times of day; (3) clustering of order executions at certain price levels; and (4) the positive-feedback trading generated by stop-loss orders. These factors matter because order flow (the net of buy-initiated and sell-initiated trades) drives returns (Evans and Lyons 2002). Large individual orders (factor 1), for example, trigger large exchange-rate moves, so fat tails in the distribution of these orders can generate a fat-tailed return distribution; the clustering of orders at certain times (factor 2) and

price levels (factor 3) can generate extremes in order flow and returns even without large individual orders.

We show that extreme returns are statistically inevitable even in the absence of news, given the properties of price-contingent order flow, using calibrated simulations. The simulations are calibrated to match the properties of stop-loss and take-profit orders at the Royal Bank of Scotland (formerly NatWest Markets), the world's fifth largest foreign exchange dealing banks (Euromoney 2007).

Our analysis is also relevant to another important question in the literature: Why are extreme returns so much more frequent than predicted by the normal distribution? More generally, what is the source of the high kurtosis in exchange-rate returns? This high kurtosis, which is also observed in equity returns (Fama 1965) and bond returns (Roll 1970), poses significant challenges to those engaged in option pricing – the Black-Sholes model assumes normally distributed returns – and in risk management – since value-at-risk analysis requires an estimate of the likelihood of extreme adverse events. We show that when each of the four factors is taken in isolation the most important single contributor to kurtosis is fat tails in the order-size distribution. When the factors operate simultaneously, however, interactions among them prove far more important than any single factor taken individually. This implies that we can expect extreme returns that cannot be traced, even with perfect information, to an identifiable trigger like a single large order.

Most asset-pricing models assume normally distributed shocks to fundamentals and a linear relationship between fundamentals and returns and therefore conclude that returns should be normally distributed.<sup>3</sup> Even without such assumptions, however, it might be reasonable to expect normally distributed returns given the vast number and variety of shocks that contribute to high-frequency exchange rate dynamics. The central limit theorem might lead one to expect normally distributed returns whatever the shocks' underlying distributions. Our analysis indicates that extreme returns can be expected to occur more frequently than predicted by the normal distribution given the properties of price-contingent order flow. More generally, returns can be expected to have high kurtosis.

---

<sup>3</sup> The list of models that generate normally distributed returns this way is too long to enumerate. One example from the generic asset pricing literature: DeLong et al (1990). One example from the exchange-rate literature: Carlson and Osler (2000).

In identifying microstructural sources of high kurtosis our paper provides a fresh perspective on a familiar topic. We attempt to better understand the forces that generate the return process. Such understanding might ultimately permit us to anticipate changes in the size or frequency of jumps, and adjust option prices or value-at-risk estimates accordingly. Related research has generally focused on describing the return process rather than understanding it. In the analysis of option pricing, the fat tails of returns are typically dealt with by assuming that return generating process includes jumps as well as a smooth process such as Brownian motion. But the sources of jumps are generally left vague, with the implication that the jumps reflect news. Other economists have attempted to identify the statistical distribution that best describes the unconditional distribution of returns. With respect to exchange-rate returns, however, there are still many candidates for this optimal distribution, including: a single continuous distribution with relatively high kurtosis (Westerfield 1973); a mixture of normal and jump processes (Tucker and Pond 1988; Akgiray and Booth 1988); a process that mixes normal distributions with time-varying conditional variances (Andersen *et al.* 2001); and a stretched exponential distribution (Laherriere and Sornette 1998); Researchers with a physics orientation have recently noted that currency returns conform fairly well to a “power law,” with exponent around three (e.g. Gabaix et al. 2003). Daniélsson and de Vries (1997) focus exclusively on estimating the tails of the distribution, and highlight the value of using a tail estimator for value-at-risk analysis.

The sources of extreme returns highlighted here augment the determinants of the behavior of returns previously familiar to exchange-rate theorists. These new determinants include: the biological rhythms of sleeping and eating that influence intraday clustering in order flow; a behaviorally-based preference for round numbers that influences the exchange-rate clustering of orders (Osler 2003); rational price-contingent order placement; and the institutional factors that influence the size distribution of orders, such as the differing needs of corporate and financial customers and the use of barrier options.

The paper has three sections and a conclusion. Section I, which follows, presents the orders data behind our analysis and discusses the connection between extreme returns and kurtosis. Section II shows how kurtosis in returns may be related to kurtosis in the distribution of orders sizes, intraday clustering, and exchange-rate clustering. Section III analyzes how feedback between order-flow and returns can generate price cascades and price halts which further intensify kurtosis in returns. Section IV concludes.

## I. BACKGROUND

This section first describes the orders data used to document the properties of price-contingent order flow and to calibrate our simulations. It then takes a closer look at excess kurtosis in currency returns.

### A. Data

Our primary data set includes all stop-loss and take-profit orders processed by the currency trading desk at the Royal Bank of Scotland in three major currency pairs — dollar-yen, dollar-U.K. pound, and euro-dollar — over two time periods: (i) 1 September, 1999 through 11 April, 2000 and (ii) 1 June, 2001 through 9 September, 2002. For both periods the data include the time and date the order was placed, the currency, the amount, the trigger rate, the direction (buy/sell), the type (stop-loss/take-profit), and whether the order was ultimately deleted, executed, or still open at the end of the sample period. The data also include the execution time and execution rate.

Some basic properties of these orders are depicted in Table 1. Overall, 47,312 orders were placed with the bank in our three currency pairs, worth in aggregate over \$253 billion. About 57 percent of the orders are take-profits. The mean order size is \$5.4 million. About twenty-seven percent of the orders were actually executed; the rest were either deleted (72 percent) or remained open at the end of the sample period. Trigger rates were generally within one percent of the market rate at the time the orders were placed. The majority of orders are executed within one day.

Stop-loss orders involve positive feedback trading, since dealers are instructed to purchase a currency after its value has risen or to sell after its value has fallen. As such, stop-loss orders will tend to propagate trends. Take-profit orders involve negative feedback trading, since dealers are instructed to purchase currency after its value has fallen, and vice versa. As such, take-profit orders will tend to interrupt trends. These orders are conditional market orders, so they absorb liquidity when executed. In this way they differ fundamentally from limit orders; limit orders are price-contingent and, like take-profit orders, buy (sell) orders are executed after price declines (rises). However, limit orders provide liquidity rather than absorb it. Of course, currency order flow in aggregate includes more than just stop-loss and take-profit orders. Order flow between dealers and customers (e.g., asset managers and import-export firms) also includes regular “deals,” in which customers trade at quotes provided by dealers. Order flow among

dealers includes deals and trades carried out over electronic limit-order markets, as well as stop-loss and take-profit orders.

## B. Extreme Returns and Kurtosis

This paper primarily focuses on the origins of extreme exchange-rate returns. A distribution with a high frequency of extreme values is said to have “fat tails,” a property often identified with high kurtosis. Kurtosis is essentially the ratio of a distribution’s fourth central moment to its second central moment.<sup>4</sup> For the normal distribution, this ratio always equals three, and distributions with higher kurtosis are said to have “excess” kurtosis.

Though kurtosis is typically associated with extreme values, it also depends on the relative frequency of observations close to the mean. Consider a symmetric distribution with mean zero. While kurtosis certainly rises with the relative frequency of extreme observations, it also rises with the relative frequency of observations near the mean. For a more formal analysis, consider a random variable  $x$  with mean,  $\bar{x}$ , and take the first derivative of kurtosis with respect to any individual observation:

$$\frac{\partial Kurtosis}{\partial x_i} = \frac{4(x_i - \bar{x})}{(\sigma^2)^2} \left( (x_i - \bar{x})^2 - \sigma^2 Kurtosis \right)$$

This defines two ranges, distinguished by whether the observations’ squared distance from the mean,  $(x_i - \bar{x})^2$ , exceeds the product of kurtosis and the variance,  $\sigma^2$ . For observations “far” from the mean, meaning those for which the squared distance exceeds  $\sigma^2 Kurtosis$ , kurtosis rises when they shift away from the mean. For all other observations kurtosis rises when they shift *towards* the mean. Most of the sources of kurtosis identified in this paper increase the frequency of observations near the mean as well as the frequency of observations in the tails.

Exchange-rate returns have long been known to exhibit excess kurtosis (Westerfield, 1973). Table 2 documents the kurtosis of high-frequency exchange-rate returns. The data are Reuters quotes sampled at five-minute intervals from January, 2000 through November 9, 2002. (These data were kindly provided by the Federal Reserve Bank of New York; they exclude the period September 13 through October 8, 2001 since that bank was closed during that interval.) The table shows that kurtosis at frequencies below an hour can be as high as 24, and that kurtosis

---

<sup>4</sup> Standard measures of kurtosis produced by statistical packages include an adjustment for small sample size, and often represent the difference between kurtosis and its corresponding value (about three) under the normal distribution.

generally declines as the horizon over which returns are measured gets longer. At the one- and two-day horizons kurtosis is close to the value of three that characterizes the normal distribution, though still significantly higher than that benchmark.

## **II. THE DISTRIBUTION OF PRICE-CONTINGENT ORDER FLOW**

This section examines how the size distribution and clustering tendencies of currency orders contribute to high kurtosis in *order flow*. Empirical evidence shows that order flow affects returns: a currency's value tends to rise when buy orders dominate aggregate order flow, and to fall when sell orders dominate (Evans and Lyons 2002). Thus kurtosis in order flow will be reflected in kurtosis in returns. Indeed, if the relationship between order flow and exchange-rate return is linear, then the distribution of exchange rate returns will be isomorphic to the distribution of aggregate order flow.

We focus on three properties of stop-loss and take-profit orders that influence the distribution of price-contingent order flow: (1) the size distribution of individual orders, (2) the clustering of order executions within the trading day, and (3) the clustering of executions at certain exchange rates. We explain how each property can contribute to order-flow kurtosis and provide estimates of the amount of kurtosis that would be contributed to price-contingent order flow from that property taken in isolation (we term this the “direct” contribution). Later in this section we show that kurtosis rises dramatically when these properties are all active together. In the next section we show how feedback from price-contingent order flow to returns produces yet more kurtosis.

### **A. Size Distribution of Stop-Loss and Take-Profit Orders**

A visual inspection of the size distribution of executed stop-loss and take-profit orders suggests that it differs strikingly from the normal. In Figure 1, which shows the sizes of euro-dollar orders, one can see that the distribution has multiple high peaks that are distributed roughly symmetrically on either side of €0 at units of €1 million, €2, million, €3 million, €5 million, €10 million, and €20 million (in absolute value). Distributions for other currencies have a similar shape overall, and have spikes at corresponding magnitudes: for example, there are peaks at £1 million, £2 million, etc. in dollar-pound. These distributions do not differ greatly across order types (stop-loss, take-profit).

Statistical tests confirm that the true, underlying distribution of order sizes in euro-dollar (EUR) orders is highly unlikely to be the normal, and provide similar conclusions for dollar-yen (JPY) and dollar-pound (GBP) orders. The Anderson-Darling statistic for EUR orders is 37.2, which permits us to reject the null easily – the critical value for significance at the 0.0001 level is 8.0. Anderson-Darling statistics for JPY and GBP orders are 35.6 and 45.1, respectively.<sup>5</sup>

High or "lepto-" kurtosis is one source of non-normality of the size distribution of price-contingent orders. For EUR, order size kurtosis is 725; for GBP it is 21, and for JPY it is 26. This high kurtosis in order sizes comes from a relatively high proportion of both small and large orders, as shown in Table 3. For example, in EUR 79.3 percent of order sizes are less than one-half standard deviation from the mean, while under the normal distribution the corresponding frequency would be 38.3 percent. The relative frequency of order sizes far from the mean is also striking. Turning again to EUR orders for illustration, 0.59 percent of observed order sizes fall between 3.5 and 4.5 standard deviations from the mean, while only 0.05 percent would fall in this range under the normal. The middle range, in which order sizes are observed less frequently than they would be under the normal, stretches from 0.5 to 3.5 standard deviations.

The presence of a few extremely large orders is a critical source of this high order-size kurtosis. When we exclude the one largest and one smallest orders from the sample, kurtosis for EUR orders declines from 543 to 27; kurtosis for GBP orders declines from 25 to 22; and kurtosis for JPY orders declines from 35 to 24. The strong influence of large orders on kurtosis of EUR orders does not mean, however, that a random choice of large orders has distorted the sample of executed orders, raising measured kurtosis above its true value. On the contrary, the executed orders in our sample seem to be fairly representative of the overall sample of orders placed at the bank. (the full sample also includes deleted orders and those still open at the end of the sample period). For example, the average size of the largest 5 percent of executed orders was \$34.5 million, quite close to the corresponding figure of \$37.1 million for the full sample of orders. Likewise, the fraction of order sizes greater than ten times the median (using absolute values) was 2.22 percent for executed orders and 2.25 percent for all orders placed with the bank.

If there is any bias to our estimate of order-size kurtosis for executed orders, it would seem more likely to be an error of understatement than overstatement. Kurtosis of all EUR

---

<sup>5</sup> The Anderson-Darling test is based on a weighted average gap between the cumulative distributions for the normal and for the sample. D'Agostino and Stephens (1986) recommend this test as the most powerful test for the uniform distribution.



orders, 985, noticeably exceeds the kurtosis of executed orders, 543. This could be related to the fact that the largest placed order, at €858.25 million, was substantially larger than the largest executed order, at €545.0 million. Extremely large orders are quite infrequent, even in the sample of placed orders, because those placing orders are wary of being exploited by the dealers. Even dealing banks' own options traders will sometimes hide their largest orders from the spot dealers for defensive reasons. Thus price-contingent order flow overall, which would include price-contingent trades planned and executed but not placed with dealers, could well have higher kurtosis than our estimates suggest.

Note that many of the largest orders, which contribute disproportionately to order size kurtosis, were placed by the bank's exotic options desk. It seems likely that these orders are intended as dynamic hedges for barrier options (options that either disappear or appear when the exchange rate crosses a particular level). To dynamically hedge an up-and-out call option (which disappears if the exchange rate rises above a certain level), for example, one must first open a short position in the underlying asset and then repurchase the full amount of the hedge when the exchange rate crosses the knock-out price. This contrasts with the small purchases and sales required to dynamically hedge standard options. The importance of barrier options to order flow and exchange-rate dynamics is widely familiar to market participants, among whom the location and magnitude of barrier options is a daily source of discussion.

To evaluate the contribution of order-size kurtosis to kurtosis in order flow we first assume simplistically that only one order is executed per period. One can imagine that one signed order size is picked at random each period, where a positive (negative) order represents a customer buy (sell). In this case the distribution of order flow is isomorphic to the distribution of order sizes. If multiple orders are executed each period, aggregate order flow becomes the sum of randomly-selected (and signed) order values, and order-flow kurtosis necessarily declines toward a limit of 3.0 because, by the central limit theorem, the distribution of the sum converges to the normal as the number of elements in the sum increases.

To evaluate the speed with which kurtosis declines we turn to Monte Carlo simulations. We create thirty separate series of orders, sampling order sizes at random from the three observed order-size distributions. We take one "period" to represent one half hour, since this appears to represent the time until a transaction has its maximum exchange rate impact (Payne

and Vitale 2003; Evans 2002). Each order series lasts 62,400 periods, which corresponds to five years of 24-hour trading days with 260 trading days per year.

The average kurtosis of individual order sizes for the thirty series based on euro-dollar orders, 513.0, represents the kurtosis of the underlying order samples. As shown in Table 4, order-flow kurtosis falls to 105.3 when five deals are executed each period, with standard error 15.6. It falls to 54.6 at 10 deals per period (standard error 8.3), 13.2 (standard error 1.5) at 50 deals per period, and 8.2 (standard error 0.7) at 100 deals per period. For the other two currencies, kurtosis of order flow begins at a lower value, since the kurtosis of the underlying order size distribution is lower, and falls from there. Kurtosis remains statistically significantly above three in all cases, however.

What is a realistic estimate of price-contingent orders executed per half hour? Since no detailed data on price contingent orders are available beyond those examined here, we construct the following educated guess: On average, during the more recent portion of the orders data, 32 new euro orders were placed with the bank each day. Since 26 percent of these euro orders were actually executed, the bank executed roughly 8.42 euro-dollar orders per day. The bank informally estimates that it captured about 4 percent of the world's currency business during our sample period. This suggests that, market-wide, on the order of 211 euro-dollar price-contingent orders are executed per day. With one "period" representing about 30 minutes, and assuming 24 hours of trading per day, the average number of orders executed per half-hour market-wide should be about 4 for euro-dollar orders. Similar back of the envelope calculations yield 3 and 5 executed orders per half hour for pound-dollar and dollar-yen currency pairs respectively. This would imply an order flow kurtosis of 129.8 for euro-dollar, 9.5 for dollar-yen and 10.2 for pound-dollar. There is, however, a substantial amount of price-contingent trading that is never formalized as orders. Some price-contingent trading is kept private for defensive reasons, as mentioned above. Technical trading is ubiquitous in foreign exchange, and many common technical strategies recommend opening a position after the rate has moved by a certain amount (e.g., momentum strategies), or closing a position according to a similar criterion (e.g., the head-and-shoulders chart pattern's "measuring objective"). We conservatively assume the average of four orders per half-hour for euro-dollar, three orders per half-hour for pound-dollar and five orders per half-hour for dollar-yen currency pairs for the rest of the paper.

## **B. Order Clustering by Time of Day**

The actual number of orders executed varies widely within a day. This variation, which depends on many factors, will also contribute to kurtosis. One source of intraday variation in the number of orders executed is the pronounced intraday pattern in the number of exchange rates crossed per period. Figure 2, which charts this intraday pattern for EUR orders, shows that the average number of exchange-rate levels crossed per half hour varies from a low of three during New York's early evening hours (ten p.m. to midnight G.M.T.) to a high of twelve during London's mid-morning. As has been well-documented within the microstructure literature, the pattern for dollar-yen and sterling-dollar are fairly similar.

This intraday variation in the number of exchange-rates crossed per period should produce predictable variability in the number of orders executed per period. Consistent with this, the intraday frequency of executed EUR orders closely resembles the pattern of exchange-rate levels crossed per period (Figure 3). A moderate fraction of EUR orders are executed during Tokyo morning trading hours, a relatively high fraction are executed during the London trading day, a moderate fraction of orders are executed during the New York trading afternoon, and very few orders are executed between then and the opening of the Tokyo market the next day. The intraday seasonal patterns are qualitatively similar for JPY and GBP orders.

The intraday pattern of order execution does not exactly match the pattern of exchange-rate volatility, however. In particular, the number of executed orders declines sharply and actually reaches zero between the end of trading in London and the beginning of trading in Asia, even though non-trivial exchange rate volatility persists during these hours. Differences in the behavior of exchange-rate volatility and price-contingent order flow probably reflect other factors that influence price-contingent order flow. Most importantly, strong intraday seasonals in the pattern of order placement (Figure 4) are likely to be important, since over one third of executed orders are open less than three hours. Indeed, the sharp decline in order execution after London trading hours finds a clear parallel in a contemporaneous sharp decline in order placement.

Intraday variation in order flow should contribute to kurtosis in aggregate order flow. Periods of high order flow should have high conditional variance, since the order flow will represent the sum of many random numbers. By contrast, periods of low order flow should have low conditional variance. The conditional mean of order flow will not be affected by the number

of orders, of course. It is well known that mixing distributions with the same mean but varying conditional variances creates an overall distribution with high kurtosis.

We estimate the direct contribution of intraday variation to order-flow kurtosis using Monte Carlo simulations calibrated to match our underlying data. We draw  $N$  orders at random each period. We draw each order size from a normal distribution with the mean and standard deviation chosen to match the mean and standard deviation of executed orders in the three currency pairs. In each period, the likelihood that a given order is actually executed is determined by the frequency with which orders were executed during that period in the underlying data. As before, we create 30 different series of 62,400 periods.

Kurtosis of aggregate order flow at the half-hour horizon is estimated to be 4.0, 3.8, and 4.4 for euro-dollar, dollar-yen, and sterling-dollar, respectively. All three figures are different from three at high levels of significance. Nonetheless, these figures are much closer to three than observed values, which range from 11 to 19 across our three currencies. Though one would tend to conclude from these figures that intraday variation in executed order flow is not an important source of exchange-rate kurtosis, our later analysis shows that conclusion to be premature.

### **C. Order Clustering by Exchange-Rate Level**

One important determinant of the number of orders executed in a given period is the exact level of the exchange rate. We focus on the final (right-hand-side) two digits of the requested execution rate -- for example, the last two digits are 45 for both the following exchange rates: \$1.2345/€, ¥123.45/\$. Executed orders cluster strongly at round-numbered rates, meaning those ending in 0: on average roughly three percent of orders have this ending digit. If people had no special preference for any number then these frequencies would be constant at one percent across the one hundred two-digit combinations. They have a secondary tendency to cluster at rates ending in 5, such as \$0.9365, where the average frequency is about two percent. Though requested execution rates end infrequently in other digits, it can still be noted that they end more frequently in 2, 3, 7, and 8 than in 1, 4, 6, and 9 (Osler, 2003).

The implication of this order clustering is that the number of orders executed per period depends on the particular exchange rate(s) just crossed. If the exchange rate has just crossed a round-numbered exchange rate, many orders will be executed; if it has just crossed an exchange rate ending in 41, only a few orders will be executed. This variation in order numbers would contribute to leptokurtosis in aggregate order flow through the same mechanism as time-of-day

clustering. The conditional standard deviation of order flow will be large when many orders are executed and small when few orders are executed. The unconditional distribution of order flow will be a mixture of normal distributions with shared mean but differing conditional variances, so the distribution will have fat tails and high kurtosis.

To evaluate the direct contribution of exchange-rate clustering to order-flow kurtosis we turn again to calibrated Monte Carlo simulations. As before, we assume that order sizes are normally distributed with a mean and standard deviation that matches the moments of the underlying dataset. We begin by supposing that the exchange rate randomly moves one point each period, and that all exchange-rate levels are equally likely to be reached.

Once again we create 30 different series of 62,400 observations for each currency. For each series we first draw at random a two-digit exchange-rate level. We then generate order flow hypothetically “triggered” by that level, calibrating order flow so that its frequency distribution across two-digit levels matches the corresponding distribution in our data. For example, ten percent of all simulated euro-dollar orders are executed after simulated exchange rates reach levels ending in 00, just as ten percent of all executed euro-dollar orders in the dataset have a trigger rate ending in 00.

This calibration is accomplished as follows. Suppose that the exchange rate ends in two-digit combination  $Z$ , and that requested execution rates end in  $Z$  for the fraction  $Y_Z$  of all orders. We randomly select 100 different order sizes from the normal distribution described above. To determine which orders are simulated as "executed," for each order  $k$  we also draw a random variable  $X_k$  from a uniform distribution over  $[0,1]$ . Each order  $k$  is considered executed, and thereby included in order flow, if  $X_k < Y_Z$ ; all other orders are ignored. Aggregate order flow for the period is the sum of (signed) sizes of all included orders. On average, a share  $Y_Z$  of orders is executed after the exchange rate hits a level ending in  $Z$ .

The 30 euro-dollar series generated according to this algorithm have a mean kurtosis of 9.4 with standard error 0.3. Average kurtosis for dollar-yen and sterling dollar are 10.2 and 11.8, respectively. However, these kurtosis levels likely overstate the contribution of this aspect of exchange-rate clustering, because exchange rates usually cross more than one level per half hour. Euro-dollar exchange rates, for example, cross 6.4 levels per half-hour, on average.<sup>6</sup> Each

---

<sup>6</sup> This figure is calculated using all trading hours outside of weekends. Following Baillie and Bollerslev, weekends hours are taken to extend from Friday at four p.m. E.S.T. to Sunday at four p.m. E.S.T.

exchange rate crossed in a given period affects aggregate order flow through the order executions it triggers. Suppose the rate goes from 120.48 to 120.50 in a given period. Few orders are likely to be triggered by rates ending in 49, and many orders are likely to be triggered by rates ending in 50. The share of the day's aggregate order flow executed that period will therefore be closer to its mean of  $1/48$  than it would be if the exchange rate crossed only one level during that period.

If the exchange rates crossed in a given period were chosen at random then, by the central limit theorem, the distribution of order flow would approach the normal as the number of levels crossed per period rises towards infinity. If the exchange rate crossed five hundred levels every period, for example, it would necessarily cross all of them and there would be no recognizable exchange-rate clustering of orders. However, exchange rates are not randomly chosen; instead, they must be adjacent to each other. This constraint slows the convergence of the distribution towards the normal, though the normal remains the limiting distribution.

We next modify our Monte Carlo simulations by supposing that  $N$  exchange-rate levels are crossed per period. An initial level for each period, and the direction of change (up or down) are picked at random, both from uniform distributions. Suppose  $N$  is four, the initial exchange rate level is selected to be something ending in 27, and the direction of motion is chosen to be down; then the exchange rates crossed that period will end in 27, 26, 25, and 24.

To generate each period's order flow we follow a modified version of the scheme described in the previous subsection. If  $N$  exchange rates are crossed in each period, then  $N$  sets of one hundred order sizes are selected at random from the same normal distribution described above, each with an associated variable  $X_k$  randomly selected from a uniform distribution over  $[0,1]$ . Within the hundred orders associated with the first exchange-rate level,  $Z$ , an order  $k$  is once again included in order flow if  $X_k < Y_Z$ . Within the hundred orders associated with the second exchange-rate level,  $V$ , an order  $k$  is included in order flow if  $X_k < Y_V$ ; etc. Order flow for the period is the sum of order flow associated with each of the  $N$  levels crossed that period.

As expected, the kurtosis of aggregate order flow is inversely related to the number of exchange rates crossed per period (Table 5, first two columns), and kurtosis approaches the theoretical limit of three, consistent with the normal distribution, as the number of exchange rates crossed per period gets large. If four exchange rates are crossed per period, for example, order flow kurtosis is 4.2, 4.2, and 4.4 for euro-dollar, dollar-yen, and sterling-dollar, respectively. These are substantially closer to its theoretical limit of three than kurtosis of roughly 10 at  $N=1$ ,

but they still differ significantly from three. Overall, we conclude that the preference for round-numbered trigger rates has a far smaller direct contribution to kurtosis than order-size kurtosis, and its direct contribution is roughly comparable to that of intraday seasonals in order execution.

*Order Flow Depends on Order Type:* Though the analysis has so far assumed that the clustering of trigger rates is independent of order direction (buy/sell) and order type (stop-loss, take-profit), this is not the case. Stop-loss buy orders cluster strongly just above levels ending in 00 or 50, at rates like \$1.4305/£ or ¥125.10/\$, while stop-loss sell orders cluster just below such round numbers. Take-profit orders cluster in the opposite way. These asymmetries are summarized numerically in Table 6.

These differences between the distributions of stop-loss and take-profit orders are important. Any exchange-rate move triggers both stop-loss and take-profit orders, and one of these order types will require purchases while the other requires sales. If the rate rises, for example, it triggers take-profit sell orders and stop-loss buy orders. If both types of orders have the same distribution, the orders will tend to offset each other. When there are many stop-loss buy orders there will likewise be many take-profit sell orders, so the mean of order flow is likely to be small. Since the orders cluster differently, however, there will be times when one order type clusters strongly while the other does not, which could generate large bursts of net order flow and extreme returns. There will also be times when both order types are relatively infrequent, which should generate observations of tiny net order flow and small returns. High net order flow can contribute to kurtosis by creating fat tails, and tiny net order flow can contribute to kurtosis by creating a “tall skinny middle” of the distribution.

To examine the contribution of asymmetric clustering we modify the simulation algorithm above by having not one but four sets of orders: stop-loss buy (SLB), stop-loss sell (SLS), take-profit buy (TPB) and take-profit sell (TPS). The relevant exchange-rate clustering frequencies are taken directly from the underlying sample of executed EUR, GBP, and JPY orders.

As before, in each period, and for each of the  $N$  exchange rates crossed, we randomly select one hundred different order sizes. However, to match the underlying dataset, we take 40 of the orders to be stop-loss orders and the remaining 60 to be take-profit orders. Each order is drawn as a value from a normal distribution with the mean and the standard deviation of actual

order sizes. The sign of the order is then assigned to fit market conditions: If the rate is falling, for example, stop-loss orders take a negative sign and take-profit orders take a positive sign.

Having chosen 40 stop-loss order sizes and signs and 60 take-profit order sizes and signs, we next determine which of these orders will be included in order flow. As above, we draw for each order  $k$  a random number  $X_k$  taken from  $U[0,1]$  and compare that to the frequency with which orders of that type are actually executed at the current exchange-rate level. Suppose the exchange rate has fallen through the two-digit exchange-rate level  $Z$ . Then each stop-loss order will have a negative sign, and each take-profit order will have a positive sign, since falling rates trigger stop-loss sell orders and take-profit buy orders. Suppose that the share  $Y_Z^{SS}$  of stop-loss sell orders typically have trigger levels ending in  $Z$ , along with the share  $Y_Z^{TB}$  of take-profit buy orders. A stop-loss order is included in order flow if  $X_k < Y_Z^S$ ; a similar condition applies for take-profits.

The results of these simulations, reported in columns three and four of Table 5, suggest that asymmetric exchange-rate clustering raises kurtosis, as expected, but not by much. Assuming that four rates are crossed per half hour, order-flow kurtosis for euro-dollar increases from 4.2 to 4.4. This increase, though small, is statistically significant.

Overall, our simulations indicate that the direct contribution of exchange-rate clustering to order-flow and return kurtosis is of the same order of magnitude as the contribution of time-of-day clustering and far smaller than the direct contribution of kurtosis in individual order sizes.

#### **D. Interactions**

We next examine order-flow kurtosis when all three of these sources of kurtosis are active simultaneously, by calibrating the Monte Carlo simulations to all three properties of the underlying data. If only direct effects matter then euro-dollar kurtosis in these simulations would be around 137.7 (137.7= 129.8 from order sizes + 4.0 from time-of-day clustering + 3.9 from exchange-rate clustering). Instead, the kurtosis of half-hour euro-dollar returns in these simulations averages 305.3, over twice as much as the combined direct effects (Table 7, first column). We infer that interactions among the three order-flow factors are as important as the direct contributions of the factors themselves. This, in turn, implies that the time-of-day clustering and exchange-rate clustering are far more important than their direct effects suggest. Without them there would be no interactions and kurtosis in this set of simulations would average only 129.8. Figures for dollar-yen and sterling-dollar paint a similar picture, though



order-flow kurtosis is lower at every time horizon due to these currencies' lower kurtosis of order sizes.

The following example illustrates how these interactions could be so important. Suppose the exchange rate rises and that, by chance, the orders triggered just happen to be particularly large. Suppose as well that the exchange rate specifically rose through  $[00,05]$ , so these large orders are dominated by stop-loss buy orders. Finally, suppose that the price move takes place at ten a.m. London time, the busiest time of the day, so there are many stop-loss orders in this cluster. In this way, the intraday seasonals augment the effect of exchange-rate clustering, which in turn augment the effect of kurtosis in order sizes.

Kurtosis in these simulations decreases monotonically as time horizon lengthens: As we move from a half to a full hour, kurtosis drops by roughly half, and then continues declining. It reaches single digits at the one-day horizon for euro-dollar and within only two hours for dollar-yen and sterling-dollar. Even at the two-day horizon, however, kurtosis remains statistically significantly above three. This monotonic decline in kurtosis is predicted by all three order-flow factors. Consider an increase in our base-line time horizon from a half-hour to an hour. The number of orders executed per period would double, on average, generating a significant reduction in order-flow kurtosis through the central limit theorem. Order-flow kurtosis would also be lowered by a reduction in the variability in the number of orders executed per period across the day and by a doubling in the number of exchange-rates crossed per period (Table 5).

### **III. DYNAMIC INTERACTIONS**

This section examines how dynamic interactions between order-flow and exchange-rate changes increase the kurtosis of each. It also shows that these dynamic interactions affect the relationship between return kurtosis and return horizon.

#### **A. Price-Contingent Orders and Exchange-Rate Trends**

There are three mechanisms through which order flow and returns can interact to generate kurtosis: random exchange-rate volatility, price cascades, and price halts.

*Random Exchange-Rate Volatility:* Randomness in order flow generates randomness in exchange-rate volatility. This should increase the “fat tails” of the return distribution by causing greater variation in the number of exchange-rate levels crossed per period. In previous simulations the exchange rate always crossed a fixed number of levels in a given half-hour. For

example, it always crossed *exactly* twelve levels per half hour around mid-afternoon London time. When exchange-rate changes are determined by random order flow, however, we only require the exchange rate crosses *an average of* twelve levels during that half hour. Sometimes the rate crosses more than eight levels, and sometimes it crosses fewer. Likewise, when in these dynamic simulations we can only ensure that the exchange rate crosses an average of four levels per half hour during the market's overnight hours. As a result, the total range of intraday return volatility will be wider in the dynamic simulations, intensifying the mixture-of-distributions effect associated with time-of-day clustering.

*Price Cascades:* Price cascades are self-reinforcing price moves that should increase the "fat tails" of the return distribution. An exchange-rate decline through a round number could trigger stop-loss sell orders and thus cause the rate to fall further. This additional move might induce the execution of yet more stop-loss sell orders and the rate would continue to fall, creating an extreme return at the two-period horizon. In reality, commentators frequently ascribe large exchange-rate moves to such price cascades, at least in part. Knowledgeable market participants informally estimate that price cascades happen at least weekly. More formal evidence for the existence of price cascades in currency markets is presented in Osler (2003).<sup>7</sup>

*Price Halts:* Price halts should increase return kurtosis by increasing the mass at the middle of the return distribution. For example, a rate decline could trigger a cluster of take-profit buy orders that stops any further decline. With no change in the exchange rate, no price contingent orders will be triggered, so the rate is more likely to stay roughly constant. This dynamic process should raise the frequency of small returns at the two-period horizon, which in turn raises kurtosis at that horizon.

## **B. Dynamic Simulations: Incorporating the Effect of Order Flow on Returns**

To measure the influence of these dynamic interactions on kurtosis we extend our calibrated Monte Carlo simulations to make exchange-rate movements depend on order flow. Our simulations assume that the current return (in logs) is a constant multiple of this period's net order-flow. If  $s$  is the log exchange rate, then

---

<sup>7</sup> If the connection between price moves and stop-loss orders were common knowledge, then rational speculators might step in and stop this "price cascade." However, information about stop-loss orders is held fairly closely by banks, for good strategic reasons: others may "pick off" stop-loss orders by pushing prices through their requested execution rates. In this event the rate is likely to immediately revert to its previous level, making the execution of the stop-loss ex-post inappropriate.

$$(1) \quad \Delta s_t = \text{Constant} * \text{OrderFlow}_t.$$

Order flow, in turn, is largely determined by the number of levels crossed as the exchange-rate changed in the previous period.<sup>8</sup>

Order flow for any period  $t$  is generated as follows: For each exchange-rate level crossed in  $t-1$ , we randomly draw one hundred order sizes from normally-distributed stop-loss and take-profit order sizes, where the mean and standard deviations of each distribution equal the corresponding figures for the underlying data. For instance, if the rate crossed five levels during  $t-1$ , moving from 1.0901 to 1.0906, we pick five hundred orders from these simulated order-size distributions. As before, 40 from each of group of hundred orders are drawn from the stop-loss sample, and the remaining 60 are drawn from the take-profit sample, reflecting the share of stop-loss and take-profit orders in our underlying sample of the Royal Bank of Scotland's executed orders. The exchange-rate's direction of motion in  $t-1$  determines the direction of each order in period  $t$ . If the rate was rising in  $t-1$ , for example, all stop-loss orders in  $t$  are taken to be buys (the order value is given a positive sign) and all take-profit orders in  $t$  are taken to be sells (the order value is given a negative sign).

In setting the new exchange rate for a given period we must ensure that the exchange rate number itself has only five significant digits, consistent with actual practice, so that its last two digits are well defined. To achieve this we invert the log exchange rate and round it off.

We calibrate our dynamic simulations to five critical features of the trading process:

1. *The average number of orders executed per period and its intraday distribution.* As noted in Section II.A, back-of-the-envelope calculations indicate that the average number of price contingent orders per half hour should be between three and five, depending on the currency. We set this number to four for euro-dollar, five for dollar-yen and three for sterling-dollar. With respect to the intraday distribution of executed orders, Figure 5A compares the true intraday frequency distribution for executed euro-dollar orders at RBS with the corresponding distribution in the euro-dollar simulation. They match each other fairly closely. The calibration for dollar-yen and sterling-dollar is about equally accurate.

2. *The average number of exchange-rate levels crossed per period and its intraday distribution.* The average number of exchange rate levels crossed per period was set at 6.4 for euro-dollar, 7.5 for dollar-yen and 7.7 for sterling-dollar in the simulations, matching the figures

---

<sup>8</sup> To initiate the simulations we set the period-one and period-two exchange rates exogenously.

for the underlying dataset. Figure 5B compares the intraday frequency distribution for the number of levels crossed with the corresponding distribution for simulated euro-dollar orders. These match each other fairly well.

3. *The share of exchange-rate changes equal to zero.* The number of times that there was no exchange rate move in the simulated exchange rates is calibrated to be close to the average of this share in the three true exchange-rate series.
4. *The frequency distribution of order execution with respect to two-digit exchange-rate levels, across the four order types (slb, sls, tpb, tps).*
5. *The share of take-profits within all executed orders.*

We evaluate the impact of dynamic factors on kurtosis in two steps. First we analyze the impact of the feedback from order flow to returns without the influence of the three order-flow factors considered previously. We then reintroduce those order-flow factors.

*Dynamics Alone:* For each randomly selected order  $k$ , we draw a random size from the normal distribution described earlier and a random variable  $X_k$  from  $U[0,1]$ . The likelihood that any individual order is executed is set at the overall likelihood for the entire sample, 27 percent. Thus it is independent of the time of day and the exchange-rate levels crossed. As before, each order  $k$  is included in “executed” order flow if  $X_k$  is less than the execution likelihood,  $Y$ ; all other orders are ignored. Aggregate order flow for period  $t$  is the sum of (signed) orders sizes for all these hypothetically executed orders.

These simulations generate simulated euro-dollar return kurtosis of 12.5 at the half-hour horizon, with standard error 0.16 (Table 8, column 1). The existence of kurtosis at this time horizon cannot be attributed to either price cascades or price halts, both of which are multi-period phenomena. Instead, kurtosis at the half hour horizon must uniquely reflect volatility randomness. Kurtosis for dollar-yen and sterling-dollar are close, at 11.0 and 14.0 respectively.

The contribution of price cascades and price halts becomes apparent at the one-hour horizon, where kurtosis in euro-dollar returns is 10.0 (standard error 0.09). Though this is lower than kurtosis at the half-hour horizon, it has not fallen very rapidly: in the earlier, static simulations, kurtosis fell by roughly half between the half-hour and one-hour horizons. As before, similar patterns are observed for the other two currency pairs (Table 8, columns 3 and 5)

*Dynamics Plus Order-Flow Factors:* When we reintroduce the three order-flow factors into these dynamic simulations, average kurtosis of half-hour euro-dollar returns balloons to 946

with standard deviation 241 (Table 8, column 2). This not only exceeds kurtosis due to the dynamics alone (12.5), it also exceeds kurtosis due to order-flow factors alone, which averaged 305, and it exceeds the sum of those two values. This underscores the critical importance interactions among the various sources of kurtosis noted earlier. Similar results apply for our other two currencies, where kurtosis in these simulations is 99 for dollar-yen and 157 for sterling-dollar.

Figure 6 plots an illustrative euro-dollar exchange-rate path generated by these simulations. The price cascades that fatten the tails of the return distribution are apparent as occasional vertical lines, while the “price halts” that increase the distribution’s skinny middle are visible as stretches where the rate seems to stay roughly constant.

#### **IV. NON-LINEAR EFFECTS OF ORDER FLOW ON RETURNS**

The kurtosis figures from our simulations substantially exceed those of actual exchange rates. This can be traced to the assumption that order flow and exchange-rate returns are linearly related. This assumption has been employed in much empirical research (e.g., Evans and Lyons 2002), but is intended exclusively to serve as a first approximation. There are good reasons to believe, and some evidence to indicate, that the effect should be non-linear, and more specifically that large imbalances should have proportionately smaller effects, at least at the intraday horizons.

At the theoretical level, a non-linear relation would be predicted based on the analysis of Bertsimas and Lo (1998), who show that traders can minimize the price impact of a given-sized trade by splitting it into smaller transactions and distributing those transactions over time. This strategy is known to be used ubiquitously in foreign exchange markets, as well as in stock markets such as the London Stock Exchange (Reis and Werner date?). Evidence for a non-linear effect is provided in Berger *et al.* (2006), which shows that the proportionate price impact of minute-by-minute interdealer order flow, while positive in all cases, declines with the amount. Hasbrouck (1991) similarly finds that the proportionate impact of equity order flow declines with the amount.

To capture this, we create a new set of simulations assuming that the effect on returns of a given amount order flow is proportionate to the square root of the order flow’s absolute value. The results, presented in Table 9 show a dramatically lower return kurtosis than resulted from the

linear simulations. EUR kurtosis at the half-hour horizon is now only 7.8, far lower than the value of 945.8 found when order-flow's effect on returns was linear. Kurtosis at the half-hour is also lower in the non-linear specification for JPY and GBP, though the decline is not as dramatic since the values were not as high under the linear simulations. In JPY, for example, half-hour return kurtosis falls from 98.9 to 4.9.

The non-linear specification also differ strikingly from the linear specification in the relation between time horizon and kurtosis. With the linear specification, kurtosis declined monotonically with horizon. In the non-linear specification this is no longer true. Kurtosis initially rises from the half-hour to the hour horizon, and only then declines monotonically. In EUR, kurtosis rises from 7.8 at the half-hour horizon to 10.5 at the one-hour horizon, and then declines only to 8.5 at the two-hour horizon. As before, the pattern is not as dramatic for JPY and GBP but it is still quite marked: In JPY, for example, kurtosis rises from 5.1 at the half-hour horizon to 6.8 at the one-hour horizon, and then declines only to 6.1 at the two-hour horizon. This rise in kurtosis between the one- and two-period horizons is necessarily attributable to the price cascades and trading halts associated with dynamic interactions between returns and order flow, as noted earlier.

Even with the non-linear specification, kurtosis remains significantly above the benchmark value of three for all three currencies at all time horizons. This confirms one's natural intuition that the kurtosis-enhancing forces highlighted here are should apply quite generally.

This paper was originally motivated by the prevalence in financial markets of extreme asset returns unaccompanied by news. Table 10 presents the frequency with which half-hour returns of varying sizes (relative to the standard deviation) would be observed in our simulations. As predicted by the high kurtosis of returns, the frequency of very small returns and very large returns are both higher than under the normal distribution, while the frequency of medium-sized returns is lower than under the normal. Returns of 3.5 to 4.5 standard deviations can be expected to occur between 4.6 and 7.6 times more frequently than under the normal; and returns of 4.5 to 5.5 standard deviations could occur between 24 and 64 times more frequently than under the normal. Returns larger than 5.5 standard deviations would occur between 422 and 5,044 times more frequently than under the normal distribution. To be more concrete, a 5.5-standard-deviation half-hour return in EUR would be 0.6 percent or larger (which, if compounded, would be 14.9 percent over a day). Returns of this or larger magnitude could be expected 39 times per

year. Similarly, a 5.5-standard-deviation hourly return in JPY would be 0.7 percent (or 18.7 percent daily); returns of this or larger magnitude could be expected 115 times per year.

## V. CONCLUSION

This paper shows that extreme returns are statistically inevitable in currency markets, even in the absence of news, due patterns in the placement and execution of price-contingent orders. We highlight four properties of stop-loss and take-profit orders that contribute to a high frequency of extreme returns – and, more generally, to high kurtosis of returns: (1) the distribution of the order sizes, which itself has high kurtosis; (2) time-of-day clustering in order execution; (3) exchange-rate clustering of order execution; (4) feedback from order execution to returns, which generates price cascades and price halts. Using simulations calibrated to the properties of orders at the Royal Bank of Scotland, the world’s fifth largest dealing bank, we evaluate the relative contributions of these factors. When the factors operate in isolation, the single most important factor is leptokurtosis in the distribution of order sizes. When the factors interact with each other, however, their interactions generate far more kurtosis than any single factor in isolation.

By enhancing our understanding of the sources of extreme currency returns, this research may help risk managers' anticipate the likelihood and magnitude of “tail events.” It may also improve the pricing of options by illuminating the forces behind those ambiguous “jump processes” typically invoked to describe returns to the underlying asset. Future research could profitably be directed towards identifying the properties of the rest of order flow and identifying how that non-price-contingent order flow might affect the likelihood of extreme returns.

## REFERENCES

- Akgiray, Vedat, and G. Geoffrey Booth, 1988. Mixed Diffusion-Jump Process Modeling of Exchange Rate Movements *Review of Economics and Statistics* 70: 631-637.
- Andersen, T., and T. Bollerslev, and Das. 2001. Variance-ratio statistics and high-frequency data: Testing for changes and in intraday volatility patterns. *Journal of Finance* 56: 305-328
- Berger, David, Alain Chaboud, Sergey Chernenko, Edward Howorka and Jonathan Wright (2006), "Order flow and exchange rate dynamics in electronic brokerage system data," *International Finance Discussion Papers* 830, Board of Governors of the Federal Reserve System (U.S.).
- Bertsimas, Dimitris, and Andrew W. Lo (1998), "Optimal Control of Execution Costs," *Journal of Financial Markets* 1: 1-50.
- Boothe, Paul, and Debra Glassman (1987), "Comparing Exchange Rate Forecasting Models: Accuracy versus Profitability," *International Journal of Forecasting* 3: 65-79.
- Carlson, John, and Carol Osler (2000), "Rational speculators and exchange rate volatility," *European Economic Review* 44, 231-53.
- Covrig, Vicentiu, and Michael Melvin (2005), "Tokyo Insiders and the Informational Efficiency of the Yen/Dollar Exchange Rate," *International Journal of Finance and Economics* 10: 185-193
- D'Agostino, R. B., and Stephens, M. A. (1986), *Goodness-of-fit Techniques*, New York: Marcel Dekker.
- Danielsson, Jon and Casper de Vries (1997), "Extreme Returns, Tail Estimation, and Value-at-Risk," *Financial Markets Group Discussion Papers* 1997-07 (London School of Economics).
- Delong, Bradford, Andrei Shleifer, Lawrence Summers, and Robert J. Waldmann (1990), "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance* 45: 379-395.
- Easley, David, and Maureen O'Hara (1991), "Order Form and Information in Securities Markets," *Journal of Finance* 46: 905-27.
- Euromoney Survey, 2007. "FX Poll 2007: Overall Market Share" (<http://www.Euromoney.com/Article/1330390/Article.html>)
- Evans, Martin (2002), "FX Trading and Exchange Rate Dynamics," *Journal of Finance* 57.
- Evans, Martin, and Richard Lyons (2002), "Order Flow and Exchange-Rate Dynamics," *Journal of Political Economy*.
- Fama, Eugene F. (1965), "The Behavior of Stock Market Prices," *Journal of Business* 38: 34-105.



- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley (2003), "A Theory of Power-Law Distributions in Financial Market Fluctuations," *Nature* 423: 267-270.
- Genotte, Gerard, and Hayne Leland (1990), "Market Liquidity, Hedging, and Crashes," *American Economic Review* 80: 999-1021.
- Hasbrouck, Joel (1991), "Measuring the information content of stock trades," *Journal of Finance* 46.
- Laherrere, J., and D. Sornette (1998), "Stretched Exponential Distributions in Nature and Economy: "Fat Tails" With Characteristic Scales," *European Physical Journal B* 2: 525-539.
- Osler, Carol L. (2002), "Stop-Loss Orders and Price Cascades in Currency Markets," *Federal Reserve Bank of New York Staff Papers* No. 150.
- \_\_\_\_\_ (2003), "Currency Orders and Exchange-Rate Dynamics: Explaining the Predictive Success of Technical Analysis," *Journal of Finance* 58: 1791-1819.
- \_\_\_\_\_ (2008), "Foreign Exchange Microstructure: A Survey," Forthcoming, Springer Encyclopedia of Complexity and System Science.
- Payne, Richard, and Paolo Vitale (2003), "A Transaction Level Study of the Effects of Central Bank Intervention on Exchange Rates," *Journal of International Economics* 61: 331-52.
- Reiss, Peter C., and Ingrid M. Werner (2004), "Anonymity, adverse selection, and the sorting of interdealer trades," *Review of Financial Studies* 18, 599-636.
- Roll, Richard (1970), *The Behavior of Interest Rates: An Application of the Efficient Market Model to U.S. Treasury Bills* (Basic Books, New York).
- Tucker, Alan L., and Lallan Pond (1988), "The Probability Distribution of Foreign Exchange Price Changes: Tests of Candidate Processes," *Review of Economics and Statistics*: 638-647.
- Westerfield, Janice M. (1977), "An Examination of Foreign Exchange Risk Under Fixed and Floating Rate Regimes," *Journal of International Economics* 7: 181-200.

**Table 1: Descriptive Information on Stop-Loss and Take-Profit Orders**

The table describes all stop-loss and take-profit orders for the euro-dollar, pound-dollar, dollar-yen currency pairs processed by the Royal Bank of Scotland, a major foreign exchange dealing bank, over two periods: (1) 1 September, 1999 through 11 April, 2000 and (2) 1 June, 2001 through 9 September, 2002.

	<b>All Orders</b>	<b>Stop-Loss</b>	<b>Take-Profit</b>
<b>Number Orders</b>	47,312	20,213	27,099
<b>Value (\$ Billions)</b>	253.9	114.6	139.5
<b>Share of Orders (%)</b>	100.0	42.7	57.3
<b>Mean Size (\$ Mill.)</b>	5.37	5.67	5.15
<b>Mean Distance To Mkt. (%)</b>	0.47	0.48	0.47
<b>Median Days Open</b>	0.6	0.4	0.7
<b>Share Executed (%)</b>	27	26	29

**Table 2: Exchange-Rate Kurtosis**

Absolute kurtosis of the exchange rate returns. Underlying data are Reuters quotes sampled at five-minute intervals from January, 2000 through November 9, 2002, exclusive of September 13, 2001 through October 8, 2001.

	EUR		JPY		GBP	
	Kurtosis	S.E.	Kurtosis	S.E.	Kurtosis	S.E.
<b>15 Min.</b>	<b>23.82</b>	0.02	<b>18.13</b>	0.02	<b>13.33</b>	0.02
<b>30 Min.</b>	<b>18.71</b>	0.03	<b>14.35</b>	0.03	<b>10.52</b>	0.03
<b>1 Hour</b>	<b>13.75</b>	0.04	<b>11.88</b>	0.04	<b>8.76</b>	0.04
<b>2 Hours</b>	<b>12.27</b>	0.05	<b>8.81</b>	0.05	<b>8.08</b>	0.05
<b>6 Hours</b>	<b>6.67</b>	0.09	<b>7.33</b>	0.09	<b>5.41</b>	0.09
<b>12 Hours</b>	<b>5.21</b>	0.13	<b>6.59</b>	0.13	<b>5.08</b>	0.13
<b>24 Hours</b>	<b>4.10</b>	0.18	<b>3.90</b>	0.18	<b>3.68</b>	0.18
<b>48 Hours</b>	<b>4.70</b>	0.26	<b>3.45</b>	0.26	<b>4.37</b>	0.26
<b>72 Hours</b>	<b>3.43</b>	0.32	<b>2.97</b>	0.32	<b>3.22</b>	0.32

**Table 3: Kurtosis in the Distribution of Order Sizes**

Table shows the distribution of order sizes for all executed stop-loss and take-profit orders placed at the Royal Bank of Scotland during August 1, 1999 through April 11, 2000 and June, 2001 through September, 2002 in euro-dollar. In this table, first, the standard deviation of the dollar-denominated order sizes is calculated. Then, the fraction of order sizes differing from the mean by various multiples of the standard deviation is calculated and compared with the corresponding fraction for the normal distribution.

<b>A</b>	<b>B</b>	<b>C</b>	<b>D</b>
<b>Standard deviations</b>	<b>Normal Distrib. (%)</b>	<b>Actual Distrib. Euro-dollar (%)</b>	<b>Ratio, C/B EUR</b>
<b>&lt; 0.5</b>	<b>38.29</b>	<b>79.28</b>	<b>2.07</b>
<b>0.5 to 1.5</b>	<b>48.35</b>	<b>14.75</b>	<b>0.31</b>
<b>1.5 to 2.5</b>	<b>12.12</b>	<b>4.38</b>	<b>0.36</b>
<b>2.5 to 3.5</b>	<b>1.20</b>	<b>0.53</b>	<b>0.44</b>
<b>3.5 to 4.5</b>	<b>0.04585</b>	<b>0.59</b>	<b>12.87</b>
<b>4.5 to 5.5</b>	<b>0.000676</b>	<b>0.13</b>	<b>192.31</b>
<b>5.5 to 6.5</b>	<b>0.00000380</b>	<b>0.088067</b>	<b>23175.53</b>
<b>&gt; 6.5</b>	<b>0.0000000081</b>	<b>0.251933</b>	<b>31102839.51</b>

**Table 4: Average order-flow kurtosis with calibrated order size distribution**

The table shows kurtosis of aggregate order flow with varying numbers of orders executed per period. Thirty series of 62,400 periods each were created. In each period, order sizes were drawn at random from a distribution calibrated to match that of the underlying sample of executed orders. Kurtosis is calculated for 30 separate series for each currency; the average and standard deviation of those kurtosis values is reported below. The underlying sample comprises all euro-dollar, dollar-yen, and dollar-pound orders executed by the Royal Bank of Scotland during September, 1999 through April, 2000 and June, 2001 through September, 2002. This includes approximately 13,000 orders with aggregate value in excess of \$63 billion.

Orders per Period	EUR		JPY		GBP	
	Average Kurtosis	Standard Error	Average Kurtosis	Standard Error	Average Kurtosis	Standard Error
<b>1</b>	<b>512.99</b>	68.04	<b>35.89</b>	3.00	<b>24.88</b>	0.62
<b>2</b>	<b>252.39</b>	41.95	<b>19.43</b>	2.01	<b>13.98</b>	0.46
<b>3</b>	<b>173.06</b>	31.52	<b>14.10</b>	1.12	<b>10.25</b>	0.25
<b>4</b>	<b>129.80</b>	18.04	<b>11.33</b>	0.80	<b>8.47</b>	0.23
<b>5</b>	<b>105.29</b>	15.59	<b>9.54</b>	0.67	<b>7.35</b>	0.18
<b>10</b>	<b>54.55</b>	8.31	<b>6.35</b>	0.34	<b>5.20</b>	0.10
<b>20</b>	<b>28.55</b>	3.86	<b>4.65</b>	0.16	<b>4.08</b>	0.06
<b>50</b>	<b>13.18</b>	1.54	<b>3.66</b>	0.07	<b>3.44</b>	0.03
<b>100</b>	<b>8.16</b>	0.70	<b>3.33</b>	0.04	<b>3.22</b>	0.02

**Table 5: Average order-flow kurtosis with calibrated order clustering at exchange rates**

The table shows kurtosis of aggregate order flow with varying numbers of exchange rate levels crossed per period. Thirty series of 62,400 periods each were created. For the first simulation, order sizes were drawn at random each period from a distribution calibrated to match that of the underlying sample of executed orders. For the second set of simulations, order sizes of four different order types (stop loss buy/sell, take profit buy/sell) were drawn at random in each period from normal distributions calibrated to match that of the underlying sample of executed stop loss and take profit orders. Kurtosis is calculated for each series; the average and standard deviation of those kurtosis values is reported below. The underlying sample comprises all euro-dollar, dollar-yen, and dollar-pound orders executed by the Royal Bank of Scotland during 24 months between September 1999 and September, 2002.

Number of Exchange-Rate Levels Crossed Per Period	No Distinction, Stop-loss vs. Take-profit		Separate Clustering Patterns, Stop-loss vs. Take-profit	
	Average Kurtosis	Standard Errors	Average Kurtosis	Standard Errors
<b>EUR</b>				
1	<b>9.38</b>	0.03	<b>10.65</b>	0.08
2	<b>5.52</b>	0.02	<b>6.11</b>	0.02
4	<b>4.17</b>	0.01	<b>4.36</b>	0.01
6	<b>3.78</b>	0.01	<b>3.86</b>	0.01
8	<b>3.57</b>	0.01	<b>3.63</b>	0.01
10	<b>3.43</b>	0.01	<b>3.45</b>	0.01
<b>JPY</b>				
1	<b>10.24</b>	0.05	<b>10.43</b>	0.06
2	<b>5.66</b>	0.02	<b>5.76</b>	0.03
4	<b>4.21</b>	0.01	<b>4.30</b>	0.01
6	<b>3.77</b>	0.01	<b>3.83</b>	0.01
8	<b>3.51</b>	0.01	<b>3.63</b>	0.01
10	<b>3.37</b>	0.01	<b>3.46</b>	0.01
<b>GBP</b>				
1	<b>11.75</b>	0.05	<b>12.76</b>	0.09
2	<b>6.48</b>	0.02	<b>6.95</b>	0.03
4	<b>4.37</b>	0.01	<b>4.53</b>	0.01
6	<b>3.90</b>	0.01	<b>3.99</b>	0.01
8	<b>3.67</b>	0.01	<b>3.71</b>	0.01
10	<b>3.46</b>	0.01	<b>3.45</b>	0.01

**Table 6: Trigger 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, we 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.

	<b>Stop-loss Orders</b>		<b>Take-profit Orders</b>	
	<b>Buy</b>	<b>Sell</b>	<b>Buy</b>	<b>Sell</b>
<b>At 00</b>	2.8	4.8	8.6	11.3
<b>Just Below 00: 90-99</b>	6.9	10.0	10.9	8.9
<b>Just Above 00: 01-10</b>	14.3	5.0	12.4	8.6
<b>At 50</b>	3.8	4.5	3.9	4.0
<b>Just Below 50: 40-49</b>	6.3	16.3	7.5	7.4
<b>Just Above 50: 51-60</b>	18.1	8.0	8.4	6.4

**Table 7: Average order-flow kurtosis with calibrated order flow**

The table shows kurtosis of aggregate order flow in which the average number of orders varies across periods according to the number and specific levels of exchange rates crossed. All simulations are calibrated to match observed patterns in the data, which comprise all euro-dollar, dollar-yen, and dollar-pound stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months between September, 1999 through September, 2002. Exchange-rate returns influence order flow and we accurately calibrate order sizes, the intraday pattern in the number of levels crossed, and the frequency distribution of order execution across exchange rate levels. Each set of dynamic simulations includes thirty runs of 62,400 periods, corresponding to five years of half-hourly trading days. Kurtosis, shown in bold, is the average across the thirty simulations; the standard error is also taken across the thirty simulations.

Time Horizon (hours)	EUR		JPY		GBP	
	Average Kurtosis	Standard Error	Average Kurtosis	Standard Error	Average Kurtosis	Standard Error
<b>0.5 hr</b>	<b>305.29</b>	14.63	<b>20.29</b>	0.49	<b>23.75</b>	0.40
<b>1</b>	<b>155.94</b>	7.58	<b>11.98</b>	0.26	<b>13.78</b>	0.16
<b>2</b>	<b>79.98</b>	3.88	<b>7.77</b>	0.13	<b>8.94</b>	0.09
<b>6</b>	<b>28.67</b>	1.40	<b>4.85</b>	0.05	<b>5.74</b>	0.07
<b>12</b>	<b>16.28</b>	0.87	<b>4.11</b>	0.06	<b>4.88</b>	0.07
<b>24</b>	<b>9.26</b>	0.49	<b>3.30</b>	0.04	<b>3.41</b>	0.04
<b>48</b>	<b>5.82</b>	0.23	<b>3.15</b>	0.04	<b>3.24</b>	0.03
<b>72</b>	<b>5.09</b>	0.21	<b>3.04</b>	0.05	<b>3.16</b>	0.05



**Table 8: Return kurtosis with linear effect of order flow on returns**

The table shows kurtosis of exchange-rate returns at varying time horizons based on simulations in which the average number of orders and the number of exchange rates crossed are mutually determined. Simulations assume that order flow has a linear effect on returns over the subsequent half hour while returns affect order flow simultaneously. All simulations are calibrated to match observed patterns in the data, which comprise all euro-dollar, dollar-yen, and dollar-pound stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months between September, 1999 through September, 2002. Returns and order flow influence each other with a half-hour lag; returns are a linear function of order flow. When order-flow factors are excluded, the dynamic simulations assume normally distributed orders with no intraday seasonals and no exchange-rate clustering. The “complete” simulations incorporate the true size distribution of orders and the true clustering by exchange-rate level for stop-loss and take-profit orders separately. Further, we calibrate the simulations to match the intraday pattern of order execution frequency. Each set of dynamic simulations includes thirty runs of 62,400 periods, corresponding to five years of half-hourly trading days. Kurtosis, shown in bold, is the average across the thirty simulations; the standard error, shown in parentheses, is also taken across the thirty simulations.

Time Horizon (hours)	EUR		JPY		GBP	
	No Order-Flow Factors	Complete	No Order-Flow Factors	Complete	No OrderFlow Factors	Complete
0.5	<b>12.51</b> (0.16)	<b>945.80</b> (241.14)	<b>11.02</b> (0.18)	<b>98.90</b> (13.09)	<b>14.02</b> (0.29)	<b>156.86</b> (16.72)
1	<b>9.96</b> (0.09)	<b>497.06</b> (82.06)	<b>10.08</b> (0.09)	<b>77.80</b> (6.77)	<b>13.70</b> (0.26)	<b>125.76</b> (15.20)
2	<b>9.93</b> (0.16)	<b>271.20</b> (39.78)	<b>7.87</b> (0.10)	<b>48.94</b> (4.94)	<b>9.49</b> (0.17)	<b>76.08</b> (8.64)
6	<b>6.12</b> (0.09)	<b>107.12</b> (16.41)	<b>5.19</b> (0.06)	<b>18.55</b> (1.13)	<b>5.72</b> (0.13)	<b>32.98</b> (4.04)
12	<b>4.67</b> (0.06)	<b>58.14</b> (8.84)	<b>4.21</b> (0.06)	<b>11.73</b> (0.72)	<b>4.44</b> (0.08)	<b>17.63</b> (1.49)
24	<b>3.90</b> (0.08)	<b>30.76</b> (4.55)	<b>3.63</b> (0.06)	<b>7.16</b> (0.34)	<b>3.85</b> (0.08)	<b>10.49</b> (0.83)
48	<b>3.51</b> (0.06)	<b>16.44</b> (2.15)	<b>3.31</b> (0.06)	<b>5.33</b> (0.21)	<b>3.43</b> (0.06)	<b>6.82</b> (0.51)
72	<b>3.30</b> (0.06)	<b>11.80</b> (1.52)	<b>3.12</b> (0.05)	<b>4.84</b> (0.25)	<b>3.23</b> (0.05)	<b>5.64</b> (0.38)

**Table 9: Return kurtosis with non-linear effect of order flow on returns**

The table shows kurtosis of exchange-rate returns at varying time horizons based on simulations in which the average number of orders and the number of exchange rates crossed are mutually determined. Simulations assume that order flow has a non-linear effect on returns over the subsequent half hour, with the effect proportional to the square root of (absolute) order flow, while returns affect order flow simultaneously. All simulations are calibrated to match observed patterns in the data, which comprise all euro-dollar, dollar-yen, and dollar-pound stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months between September, 1999 through September, 2002. Returns and order flow influence each other with a half-hour lag; returns are a non-linear function of order flow. The simulations incorporate the true size distribution of orders and the true clustering by exchange-rate level for stop-loss and take-profit orders. Further, we calibrate the simulations to match the intraday pattern of order execution frequency. Each set of dynamic simulations includes thirty runs of 62,400 periods, corresponding to five years of half-hourly trading days. Kurtosis, shown in bold, is the average across the thirty simulations; the standard error, shown in parentheses, is also taken across the thirty simulations.

Time Horizon (hours)	EUR		JPY		GBP	
	Returns	Order Flow	Returns	Order Flow	Returns	Order Flow
0.5	<b>7.831</b> (0.11)	<b>506.38</b> (73.07)	<b>4.93</b> (0.03)	<b>33.02</b> (1.20)	<b>5.11</b> (0.03)	<b>35.75</b> (1.12)
1	<b>10.46</b> (0.9)	<b>408.24</b> (82.58)	<b>7.65</b> (0.06)	<b>26.75</b> (0.78)	<b>6.76</b> (0.04)	<b>27.72</b> (1.04)
2	<b>8.51</b> (0.19)	<b>245.49</b> (45.33)	<b>6.56</b> (0.05)	<b>17.87</b> (0.68)	<b>6.09</b> (0.05)	<b>18.28</b> (0.69)
6	<b>5.82</b> (0.09)	<b>92.48</b> (16.27)	<b>5.32</b> (0.07)	<b>9.25</b> (0.25)	<b>4.88</b> (0.05)	<b>9.62</b> (0.31)
12	<b>4.84</b> (0.07)	<b>47.70</b> (7.49)	<b>4.71</b> (0.08)	<b>6.54</b> (0.16)	<b>4.25</b> (0.05)	<b>6.78</b> (0.21)
24	<b>3.71</b> (0.05)	<b>25.00</b> (4.01)	<b>4.10</b> (0.08)	<b>4.66</b> (0.10)	<b>3.52</b> (0.04)	<b>4.63</b> (0.12)
48	<b>3.40</b> (0.05)	<b>13.72</b> (1.90)	<b>3.81</b> (0.07)	<b>3.79</b> (0.09)	<b>3.30</b> (0.05)	<b>3.79</b> (0.06)
72	<b>3.26</b> (0.05)	<b>10.38</b> (1.37)	<b>3.69</b> (0.08)	<b>3.45</b> (0.08)	<b>3.31</b> (0.08)	<b>3.59</b> (0.08)

**Table 10: Kurtosis in the Distribution of Simulated Exchange-Rate Returns**


The table shows the distribution of the exchange-rate returns generated from the non-linear specification (averaged over 30 simulations) in euro-dollar, pound-dollar and yen-dollar. For each currency pair, the standard deviation of the simulated exchange-rate series is determined. Then, the fraction of returns differing from the mean by various multiples of the standard deviation is calculated.

<b>Standard deviations</b>	<b>Normal Distr. (%)</b>	<b>EUR (%)</b>	<b>GBP (%)</b>	<b>YEN (%)</b>	<b>EUR Ratio to Normal</b>	<b>GBP Ratio to Normal</b>	<b>YEN Ratio to Normal</b>
<b>Half-Hour Returns</b>							
< 0.5	<b>38.29</b>	47.21	41.31	41.73	<b>1.23</b>	<b>1.08</b>	<b>1.09</b>
0.5 to 1.5	<b>48.35</b>	40.09	45.97	45.60	<b>0.83</b>	<b>0.95</b>	<b>0.94</b>
1.5 to 2.5	<b>12.12</b>	10.05	10.77	10.58	<b>0.83</b>	<b>0.89</b>	<b>0.87</b>
2.5 to 3.5	<b>1.20</b>	2.26	1.73	1.78	<b>1.88</b>	<b>1.44</b>	<b>1.48</b>
3.5 to 4.5	<b>0.04585</b>	0.35	0.21	0.27	<b>7.63</b>	<b>4.58</b>	<b>5.89</b>
4.5 to 5.5	<b>0.000676</b>	0.04327	0.016026	0.020834	<b>64.01</b>	<b>23.71</b>	<b>30.82</b>
> 5.5	<b>0.00000380</b>	0.00320520	0.0016026	0.019166	<b>843.47</b>	<b>421.74</b>	<b>5043.68</b>
<b>Hourly Returns</b>							
< 0.5	<b>38.29</b>	45.80	41.53	41.49	<b>1.20</b>	<b>1.08</b>	<b>1.08</b>
0.5 to 1.5	<b>48.35</b>	41.46	45.70	45.80	<b>0.86</b>	<b>0.95</b>	<b>0.95</b>
1.5 to 2.5	<b>12.12</b>	10.25	10.83	10.50	<b>0.85</b>	<b>0.89</b>	<b>0.87</b>
2.5 to 3.5	<b>1.20</b>	2.09	1.70	1.87	<b>1.74</b>	<b>1.42</b>	<b>1.56</b>
3.5 to 4.5	<b>0.04585</b>	0.35	0.22	0.29	<b>7.63</b>	<b>4.80</b>	<b>6.32</b>
4.5 to 5.5	<b>0.000676</b>	0.051284	0.012821	0.041668	<b>75.86</b>	<b>18.7</b>	<b>61.64</b>
>5.5	<b>0.00000380</b>	0.0032052	0.007179	0.019231	<b>843.47</b>	<b>1889.21</b>	<b>5060.79</b>

**Figure 1: Size Distribution of Executed Orders**

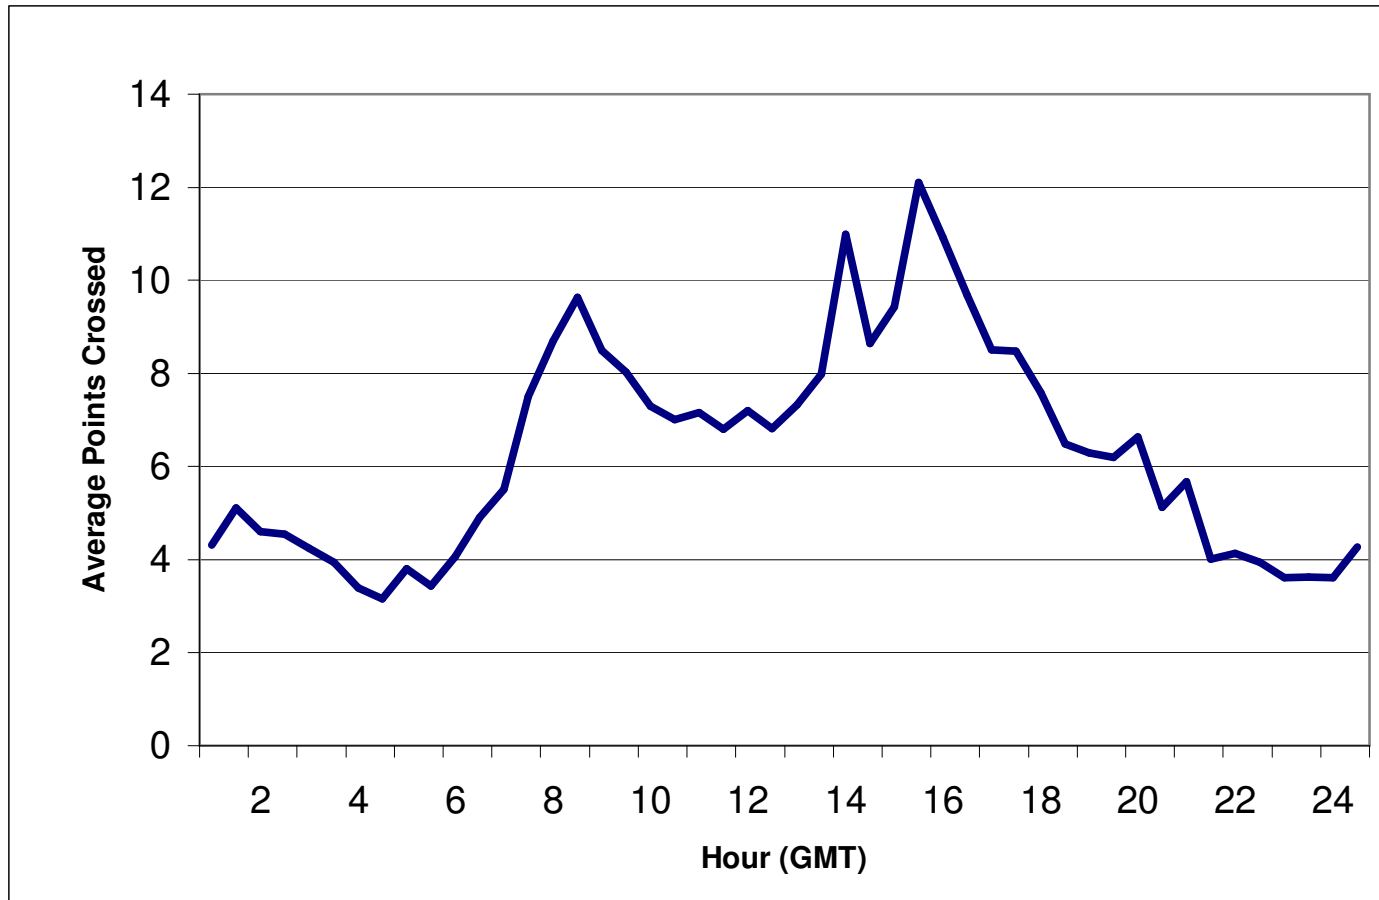
The figure shows the frequency distribution of order sizes for executed price-contingent orders. The underlying data include 4,542 executed orders with aggregate value of 25.8 billion euros, processed by the Royal Bank of Scotland during September 1, 1999 through April 11, 2000 and June 1, 2001 through September 9, 2002. Order sizes are denominated in millions of euros.

**Size Distribution of Orders  
All Executed Euro-Dollar Orders, Unweighted**




**Figure 2: Average Number of Exchange-Rate Levels Crossed, Half-Hour Intervals Across the Currency Trading Day**

The underlying data are euro-dollar exchange rate quotes (mid-rates) sampled at five-minute intervals on Reuters over June 2001 through November 2002 but excluding September 13-October 8, 2001.




**Figure 3: Frequency of Order Execution During Hour-Long Intervals Across the Currency Trading Day**

The underlying data are all euro-dollar stop-loss and take-profit orders executed by Royal Bank of Scotland during September, 1999 through April 11, 2000 and June 1, 2001 through September 9, 2002. This sample includes 4,542 executed orders with aggregate value of 25.8 billion euros.



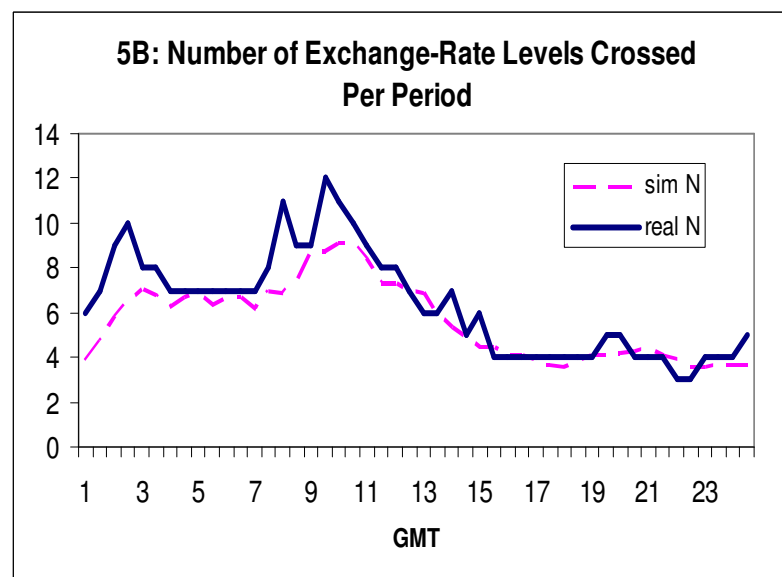
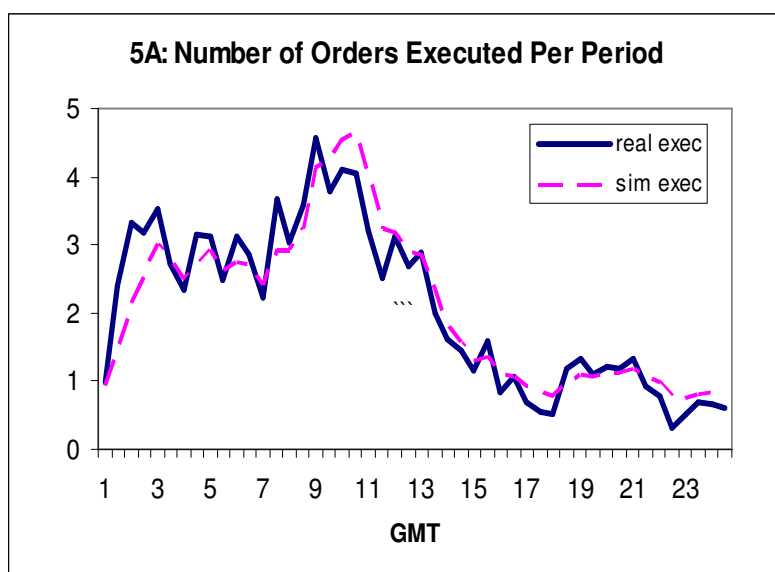
**Figure 4: Frequency of Order Placement During Hour-Long Intervals Across the Currency Trading Day**

The underlying data are all euro-dollar stop-loss and take-profit orders placed at the Royal Bank of Scotland during September, 1999 through April 11, 2000 and June 1, 2001 through September 9, 2002. This sample includes 13,658 orders with aggregate value of 98.3 billion euros.



**Figure 5: Intraday seasonals in real data and dynamic simulations: Euro-dollar only**


The underlying data are all euro-dollar stop-loss and take-profit orders executed by Royal Bank of Scotland over 24 months between September, 1999 and September, 2002. Dynamic simulations allow period- $t$  exchange-rate changes to determine period  $t+1$ 's order flow, which in turn determines period  $t+1$ 's exchange rate change. Order flow is calibrated to match the underlying data in the distribution of order sizes, the share of take-profit orders relative to stop-loss orders, the frequency with which exchange rates remain unchanged from period to period, the intraday pattern of order execution frequency, and the frequency with which orders are executed at different two-digit exchange-rate levels. The figures compare observed and simulated intraday seasonals in the frequency of order execution and in the number of exchange-rate levels crossed per period.





### Figure 6: Simulated Exchange Rate Path: Dynamic Simulations

The figure shows a simulated exchange-rate path, in which period- $t$ 's exchange-rate changes determine period  $t+1$ 's order flow, which in turn determines period  $t+1$ 's exchange rate return according to the linear relationship described in the text. Order flow is calibrated to match the underlying data in the distribution of order sizes, the share of take-profit orders relative to stop-loss orders, the frequency with which exchange rates remain unchanged from period to period, the intraday pattern of order execution frequency, and the frequency with which orders are executed at different two-digit exchange-rate levels. The underlying data are all euro-dollar stop-loss and take-profit orders executed by Royal Bank of Scotland (formerly NatWest Markets) over two time periods: September, 1999 through April 11, 2000 and June 1, 2001 through October 18, 2002. 64,200 observations would correspond to roughly five years of half-hour time periods.



### Figure 7: Simulated Exchange Rate Path: Dynamic Simulations

The figure shows a simulated exchange-rate path, in which period- $t$ 's exchange-rate changes determine period  $t+1$ 's order flow, which in turn determines period  $t+1$ 's exchange rate return according to the non-linear relationship described in the text. Order flow is calibrated to match the underlying data in the distribution of order sizes, the share of take-profit orders relative to stop-loss orders, the frequency with which exchange rates remain unchanged from period to period, the intraday pattern of order execution frequency, and the frequency with which orders are executed at different two-digit exchange-rate levels. The underlying data are all euro-dollar stop-loss and take-profit orders executed by Royal Bank of Scotland (formerly NatWest Markets) over two time periods: September, 1999 through April 11, 2000 and June 1, 2001 through October 18, 2002. 64,200 observations would correspond to roughly five years of half-hour time periods.

