

EXTREME RETURNS:

The Case of Currencies

Carol Osler^a
Tanseli Savaser^b

November, 2010

Abstract

This paper investigates how active price-contingent trading contributes to extreme returns even in the absence of news. Price-contingent trading, which is common across financial markets, includes algorithmic trading, technical trading, and dynamic option hedging. The paper highlights four properties of such trading that increase the frequency of extreme returns, and then estimates the relative of these properties using data from the foreign exchange market. The four key properties we consider are: (1) high kurtosis in the distribution of order sizes; (2) clustering of trades within the day; (3) clustering of trades at certain prices; and (4) positive and negative feedback between trading and returns. Calibrated simulations indicate that interactions among these properties are at least as important as any single one. Among individual properties, the orders' size distribution and feedback effects have the strongest influence. Price-contingent trading could account for over half of realized excess kurtosis in currency returns.

JEL classification: G1, F3.

Key words: Crash; Fat Tails; Kurtosis; Exchange Rates; Order Flow; High-Frequency; Microstructure; Jump Process; Value-At-Risk; Risk Management.

^aCorresponding author. Brandeis International Business School, Brandeis University. Mailstop 32, 415 South Street, Waltham, MA 02454-9110. cosler@brandeis.edu. Tel: 781-736-4826. Fax: 781-736-2269.

^bEconomics Department, Williams College, Williamstown, Massachusetts 01267 USA.
Tanseli.Savaser@williams.edu.

EXTREME RETURNS:

The Case of Currencies

1. Introduction

What triggers crashes in financial prices? According to standard models, the trigger for extreme returns can only be new information. In reality, however, extreme returns are often unaccompanied by financial-market news in both currency markets (Covrig and Melvin, 2005) and stock markets (Cutler, Poterba and Summers, 1991). There was no news of obvious relevance, for example, to trigger the flash-crash of May 2010, the more sustained stock-market crash of 1987 (Shiller, 1989), or the sustained 15 percent drop in dollar-yen in October 1998 (Covrig and Melvin, 2005).

Eye-popping financial returns happen much more frequently than predicted by the normal distribution. Fat tailed returns and, more generally, high kurtosis of returns have been documented for equities (Fama, 1965), bonds (Roll, 1970), and currencies (Westerfield, 1977). This contrasts with another prediction of standard models, that returns should be normally distributed. This prediction arises because the models typically assume normally distributed shocks and a linear relation between returns and fundamentals. Even without such assumptions, however, one might infer normally distributed returns in light of the central limit theorem coupled with the frequency and variety of financial market shocks.

This paper investigates how the frequency of extreme exchange-rate returns is influenced by price-contingent trading – meaning trading that is triggered when the price arrives at a specific level. Price-contingent trading is common in most financial markets because it arises from algorithmic trading, technical trading, and dynamic option hedging. We highlight four properties of price-contingent trading that contribute to fat-tails in the return distribution: (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 (Osler, 2003); and (4) positive and negative feedback between trading and returns (Genotte and Leland, 1990; Morris and Shin, 2004; Osler, 2005). We also show that interactions among these factors, which cannot be measured in any straightforward way, also contribute to fat-tails.

While some of these factors have been mentioned previously in the literature, we are the first to evaluate their relative importance and to point out the importance of their interactions, which we do using calibrated simulations. When each factor operates in isolation, the orders' size distribution and feedback effects make substantial contributions while the contributions of the other factors are economically small. When all the factors operate simultaneously, interactions among the factors prove more important than any single factor taken in isolation.

We also use the simulations to investigate the overall contribution of price-contingent trading to the excess kurtosis of returns. When we assume that order flow affects returns linearly, consistent with most of the empirical literature (e.g., Evans and Lyons, 2002), the kurtosis associated with price-contingent order flow is not only substantial, it exceeds observed kurtosis. When we assume, more realistically, that the effect of order flow on returns is concave (Hasbrouck, 1991; Berger *et al.*, 2008; Osler and Vandrovysh, 2009), simulated kurtosis accounts for between half and two-thirds of excess kurtosis associated with 1-hour returns. We conclude that price contingent orders may be an important source of kurtosis in exchange-rate returns. A further implication of our analysis is that extreme exchange-rate returns unaccompanied by news are statistically inevitable.

Our dataset comprises the complete record of stop-loss and take-profit orders processed by the Royal Bank of Scotland, the world's fifth largest foreign exchange dealing bank

(Euromoney, 2007), in euro-dollar, dollar-yen, and sterling-dollar during 24 months over the period September, 1999 through September, 2002. Stop-loss orders instruct dealers to buy (sell) currency when its value rises (falls) to a certain level; take-profit orders instruct dealers to buy (sell) after the currency's value falls (rises) to a certain level.

In identifying microstructural sources of high kurtosis, our paper provides a fresh perspective on a familiar topic. Related research generally tries to describe the return process, rather than to explain it. There is no consensus on the true distribution of returns, though there is a long list of candidates: a single continuous distribution with relatively high kurtosis (Westerfield, 1977); a mixture of normal and jump processes (Tucker and Pond, 1988; Akgiray and Booth, 1988); a stretched exponential distribution (Laherrere and Sornette, 1998); and a process that mixes normal distributions with time-varying conditional variances (Andersen et al., 2001), as suggested by the well-known statistical success of GARCH models of returns. Researchers familiar with the physical sciences 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 (1998), who focus exclusively on estimating the tails of the distribution, highlight the value of using a tail estimator for value-at-risk analysis.

Our analysis suggests that these distributions are at best approximations. Underlying the distribution of returns are the distributions of order sizes and trigger rates, and these are spiky because people have preferences for round numbers. This suggests that, though the distribution of returns should converge to the normal, it may not conform to any parsimonious function.

The existing literature also assumes implicitly that the distribution of returns is fixed, a hypothesis that is difficult to verify. Indeed, the lack of consensus on the appropriate distribution

may reflect shifts in that distribution over time. Our focus on the underlying determinants of kurtosis may help identify when and why the return distribution shifts.

Our analysis is also related to currency option pricing and the practical estimation of value-at-risk. In the empirical analysis of options prices, it is often assumed that the return generating process includes jumps as well as a smooth process such as Brownian motion. The sources of jumps are not specified, however, and the implication that the jumps just reflect news is inconsistent with the poor empirical correspondence between extreme returns and news. Our research suggests that jumps are inevitable in markets with significant amounts of price-contingent trading and suggests reasons why jump distributions might vary across time and across assets.

Our analysis of the mechanism behind tail events has an important implication for the literature's developing understanding of forward exchange rates. The attractive average returns to high-interest currencies (Burnside et al., 2010a) may compensate, in part, for skewness associated with occasional crashes (Brunnermeier et al., 2008; Plantin and Shin, 2009), though it may also reflect imperfect rationality (Burnside et al. 2010b). The market is definitely aware when carry-trade exposures build up, raising the conditional probability of a crash, yet it is difficult to know when fellow traders will pull out (Abreu and Brunnermeier, 2003; Levich and Pojarliev, 2010). Because they protect traders in the event of a crash, placing stop-loss orders can be a rational hedging strategy. But the orders themselves increase the likelihood that some arbitrary price shock turns into a crash. In equilibrium, traders place stop-loss orders, and high-interest currencies appreciate, in response to the risk of crashes that might not occur without the orders.

The four properties of price-contingent orders we highlight augment the list of return drivers familiar to exchange-rate theorists. These four properties in turn depend on features of the financial landscape on which researchers typically do not focus: the use of barrier options and other influences on the size distribution of orders (property 1); the biological rhythms of eating and sleeping that influence intraday clustering of order flow (property 2); and a behaviorally-based preference for round numbers (properties 1 and 3).

Our analysis is potentially relevant to equity and futures markets, as well as currency markets. Stop-loss orders have been identified as a significant contributor to the May 6, 2010 flash-crash in Canadian stocks (IIROC, 2010) and are suspected as a contributor in US Stocks (CFTC-SEC 2010). More broadly, price-contingent trading in the form of technical trading, algorithmic trading, and dynamic option hedging are common in equity and futures markets, and some of the four key properties of price-contingent order flow are observed there. For example, equity order executions cluster at the market's open and closing times, as indicated by the familiar U shape of trading volume (Biais et al. , 1995), and equity limit orders are known to cluster at certain price levels (Niederhoffer and Osborne, 1966; Kavajecz and Odders-White, 2004).

The rest of the paper has four sections and a conclusion. Section 2 presents our data and some important background information. Section 3 focuses on three of our four properties of price-contingent orders: kurtosis in the distribution of order sizes, intraday clustering in order execution, and clustering in the rates that trigger order executions. Section 4 analyzes the fourth factor, feedback between returns and price-contingent order-flow and evaluates the relative and overall contributions of all four factors. Section 5 extends our analysis by considering the influence of market growth, news, and trading that is not price-contingent. Section 6 concludes.

2. Background

This section discusses the institutional treatment of stop-loss and take-profit orders in currency markets, our data, and kurtosis in currency returns.

2.1. Institutional context

To implement their intended price-contingent trades, many currency market participants place stop-loss or take-profit orders with dealers. Stop-loss orders can be formalized this way:

Stop-loss buy = Buy Q at market price IFF $(S_{t+\tau} \geq S > S_t \text{ for } \tau > 0)$

Stop-loss sell = Sell Q at market price IFF $(S_{t+\tau} \leq S < S_t \text{ for } \tau > 0)$,

where Q is the quantity to trade measured in the base (denominator) currency of the exchange rate as conventionally quoted, S_t is the exchange rate when the order is placed, and S is the order's trigger price.¹ Take-profit orders can be formalized this way:

Take-profit buy = Buy Q at market price IFF $(S_{t+\tau} \leq S < S_t \text{ for } \tau > 0)$

Take-profit sell = Sell Q at market price IFF $(S_{t+\tau} \geq S > S_t \text{ for } \tau > 0)$.

Currency dealers, who work for banks, manage these orders for customers, for other banks, and for other traders within the bank. A hedge fund might instruct a dealer to buy €10 million once euro-dollar rises to \$1.35/€ (a stop-loss); similarly, an exotic options trader at the Royal Bank of Scotland might instruct a dealer at his own bank to buy \$75 million after dollar-yen falls to ¥115/\$ (a take-profit). Each bank manages its own private book of orders. (Osler (2009) provides a detailed description of the institutional structure of foreign exchange markets.)

The institutional treatment of stop-loss orders differs between the foreign exchange and equity markets, in part because currency traders face no short-sales restrictions. Indeed, short selling is not even defined in currency markets, because the sale of one currency can always be

reabeled as the purchase of another.² The lack of short-sale restrictions presumably explains why foreign-exchange dealers do not typically ask whether the customer placing a given order has a pre-existing position, and why stop-loss buy and sell orders are placed with roughly equal frequency.³

The market convention is that stop-loss and take-profit orders are only executed once the currency actually trades at the trigger price. Further, the prices that trigger an order's execution must be observed in a setting that can be audited, to minimize concerns about dealer manipulation. Many orders are "good 'til cancelled," where cancellation requires a proactive input from the order placer. Other orders have an explicit sunset provision, such as "good 'til close," meaning the order is cancelled if not executed by the end of the (local) trading day.

Take-profit orders are similar to limit orders insofar as they both involve price-contingent negative-feedback trading that can interrupt trends (Osler, 2005; Neiderhoffer and Osborne, 1966). The two differ fundamentally, nonetheless, since take-profits – like stop-losses – are conditional market orders and should in theory absorb liquidity rather than supply it.

2.2. *Data*

Our primary data set comprises all stop-loss and take-profit orders processed by the currency trading desk at the Royal Bank of Scotland (RBS), one of the world's top dealing banks (Euromoney, 2007), over two time periods: September 8, 1999 through April 11, 2000 and June

¹ In the over-the-counter foreign exchange market all currency pairs are quoted according to convention. For example, euro-dollar is always quoted as dollars per euro; dollar-yen is always quoted as yen per dollar.

² The use of currency "stops" is aptly illustrated by a recent recommendation by Fernandez in *Currency Trader Magazine* (2010) of the "adaptive breakout technique." Traders place stop-loss buys and sells a given distance beyond the price range of the first fourteen trading hours on Monday, opening positions if the orders are reached.

³ In online retail trading, some ECNs do require stop-loss and take-profit orders to match existing positions. Other ECNs are more flexible, and label orders intended to open a position "stop entry orders" or "entry limit orders." Dealers in the still-dominant wholesale market have not historically used these distinctions.

12, 2001 through September 20, 2002. We focus on the three most active currency pairs: euro-dollar, dollar-yen, and sterling-dollar. The data include the time and date the order was placed, the non-dollar currency, the amount, the trigger rate, the prevailing market exchange rate, the order's 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.

Basic descriptive statistics are shown in Table 1. The 47,312 orders in our dataset are worth in aggregate over \$253 billion. About 57 percent of the orders are take-profits. The mean (absolute) order size is \$5.4 million; the standard deviation of (signed) order sizes is \$7.9 million. About 27 percent of orders were actually executed; most of the rest were cancelled though a few remained open at the end of the sample period. Trigger rates are generally within one percent of the market mid-rate at the time the orders were placed, though the two can differ by as much as eighty percent (a gap this large could represent, among other things, the hedge for a barrier option). The majority of orders are executed or deleted within one day, though some remain open for months. Throughout the succeeding analysis we focus on those orders actually executed, since the others are unlikely to have much influence on returns.

By placing a stop-loss or take-profit order, a firm interested in price-contingent trading effectively enlists the dealer's services in monitoring the market, thereby ensuring prompt trade execution upon fulfillment of the contingency. Because dealers do not charge for this service, placing such orders is rational for many market participants. For corporate treasury staffers and portfolio managers, among others, it is not efficient to monitor this fast-moving market personally. Only a subset of price-contingent trades are recorded as formal orders, however. Large firms occasionally hire their own traders to monitor the market and some market

participants avoid placing large orders – and monitor as best they can – based on the concern that dealers might share information about their orders with others.

It is natural to wonder how large price-contingent trading is relative to all foreign exchange order flow. Unfortunately, there are as yet no data with which to evaluate these questions rigorously. When asked, dealers consistently suggest that the share of stop-loss and take-profit orders in all order flow is five percent or less, but of course much price-contingent trading is never formalized as orders. It is also worth noting that the influence of price-contingent trading on returns could exceed its share of trading, as suggested by the analysis below.

We complement the orders data with Reuters indicative mid-quotes sampled at five-minute intervals from January 2, 2000 through November 9, 2002, exclusive of September 13, 2001, through October 8, 2001.⁴ Since foreign exchange is a 24-hour market, we include all trading hours outside of weekends, which we define as Friday at 4 p.m. EST (when trading trails off in New York) to Sunday at 4 p.m. EST (when trading opens in Sydney).

2.3. *Extreme returns and kurtosis*

Kurtosis is essentially the ratio of a distribution's fourth central moment to its squared second central moment. This number is always three for normal distributions, and distributions with kurtosis above three are said to have "excess" kurtosis. Distributions with a high frequency of extreme observations are often described as having "fat tails" and high kurtosis. Technically, however, kurtosis involves not only fat tails but also a relatively high frequency of observations near the mean. Most of the sources of kurtosis identified in this paper increase the frequency of observations near the sample mean as well as the frequency of observations in the tails.

⁴ These data were kindly provided by the Federal Reserve Bank of New York. Daniélsson and Payne (2002) show that indicative quotes and tradable quotes move closely together at high frequencies.

Exchange-rate returns have long been known to exhibit excess kurtosis (Westerfield, 1977). As shown in Table 2, Panel A, the sample kurtosis of exchange-rate returns is as high as 24 at frequencies below an hour. Sample kurtosis generally declines as returns are measured over longer horizons but significant excess kurtosis remains at the one- and two-day horizons.

For many statistical distributions kurtosis and other higher moments do not exist, which raises the question of whether kurtosis exists for exchange-rate returns. Extreme-value theory reveals that a moment of order n only exists for distributions with a tail index, α , less than n . The inverse of the tail index, $\gamma = 1/\alpha$, known as the shape parameter, can be estimated using the OLS Hill estimator on the most extreme k percent of the distribution (LeBaron, 2008).⁵ Table 2, Panel B, shows OLS Hill estimators of the tail index for the right and left tails of the distribution of half-hour and 12-hour returns for our three exchange rates. These suggest that kurtosis exists for 12-hour returns – additional results show that it also exists for 1-day returns – but does not exist for half-hour returns. At intermediate time horizons the results are mixed.⁶

Fortunately, the true distribution of returns is not important for this paper. In our effort to investigate the economic origins of extreme returns, kurtosis usefully summarizes the extent to which a given distribution differs from the normal near the mean and in the tails. We rely on sample kurtosis purely as a descriptive device. Table 2, Panel C, shows that returns near and far

⁵ Suppose the conditional probability of an observation beyond u in some distribution has the following shape: $Pr(x < y | x < u) = g|y|^{-\alpha}$, where g is a scaling factor and $\alpha > 0$ is the tail index. Any moments beyond α will not exist. The

Hill estimator of γ , γ_k , is calculated as follows: $\gamma_k = \frac{1}{N_k} \sum_{i=1}^{n_k} (\log x_{n-i+1} - \log x_{n-n_k})$, where k is the fraction of

observations to be considered, N_k is number of observations to be considered, and the x_i are the order statistics. To find the OLS Hill estimator, one calculates γ_k for a variety of k 's and then runs the following regression: $\gamma_k = a + b k + e_k$. The estimated constant term provides an unbiased estimate of γ .

⁶ Unfortunately there is still no formal test for the estimator's statistical significance.

from the sample mean do occur relatively frequently for all three currencies, as suggested by their high kurtosis.

3. Distribution of price-contingent order flow

Empirical evidence shows that order flow – defined as buy-initiated trades minus sell-initiated trades – affects returns. A currency's value tends to rise (fall) when buy (sell) orders dominate aggregate order flow (Evans and Lyons, 2002). If order flow were constant, the only source of extreme returns would be news; for order flow to trigger extreme returns it must itself be extreme. This section identifies three reasons why price-contingent order flow can sometimes be extreme.

3.1. Size distribution of price-contingent orders

Figure 1 uses the size distribution of executed euro-dollar stop-loss and take-profit orders to illustrate that distribution for all three currencies. The distributions are close to symmetric with respect to direction (buy, sell) and they do not differ noticeably across order types (stop-loss, take-profit). However, it takes only a quick glance to see that this distribution differs strikingly from the normal, because it has spikes at ± 1 , ± 2 , ± 3 , ± 5 , ± 10 , and ± 20 , where the units are € millions. This raises the possibility that order sizes are influenced by a strong human preference for round numbers (Yule, 1927). The origin of these spikes in preferences is underscored by the fact that the order-size distributions for dollar-yen and sterling-dollar have spikes at the same numbers: for example, there are peaks at ± 1 , ± 2 , etc. in sterling-dollar, where the units are £ millions. Anderson-Darling statistics confirm that the underlying distribution of order sizes is highly unlikely to be the normal.⁷ The test statistic for euro-dollar

⁷ The Anderson-Darling test is a particularly powerful test based on a weighted average gap between the cumulative distributions for the normal and for the sample (D'Agostino and Stephens, 1986).

orders, 37.2, easily exceeds the critical value for 0.01 percent significance, which is 8.0. The test statistics for dollar-yen and sterling-dollar orders are 35.6 and 45.1, respectively.

These order-size distributions also differ from the Gaussian because of their high or "lepto-" kurtosis. For euro-dollar, sample kurtosis of executed order sizes is 725; for sterling-dollar it is 21, and for dollar-yen it is 26. The high sample kurtosis of order sizes reflects a relatively high proportion of both small and large orders. Using euro-dollar orders once again to illustrate the point, at the center of the distribution 79 percent of executed order sizes are within one-half standard deviation of the mean, roughly twice the 38 percent share under the normal distribution; at the middle of the distribution, 15 percent of executed orders are within 1.5 and 2.5 standard deviations of the mean, less than one-third the 48 percent share under the normal; finally, in the tails of the distribution, 0.59 percent of observed order sizes are within 3.5 to 4.5 standard deviations of the mean, far more than the 0.05 percent share under the normal.

Kurtosis for euro-dollar orders may seem unrealistically high, especially since it declines from 725 to 31 if the single largest order is excluded.⁸ This does not mean, however, that an unrepresentative large order has distorted the euro-dollar sample. To the contrary, the sample of executed euro-dollar orders seems to be fairly representative of the overall population of such orders (as a reminder, the full sample includes cancelled 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 in euro-dollar was €36.3 million, quite close to the corresponding value of €39.9 million for the full sample. Likewise, the fraction of order sizes greater than ten times the median (using absolute values) was 1.4 percent for executed orders and 1.5 percent for all placed orders.

⁸ The extremely-large orders in our sample were often placed by the bank's exotic options dealers, presumably as dynamic hedges for barrier options.

The dollar-yen and sterling-dollar orders placed with the bank included some extremely large orders that were never executed. Thus the absence of orders of this magnitude from the sample of executed dollar-yen and sterling-dollar orders, which accounts for their relatively low sample kurtosis, appears to be a matter of chance rather than a reflection of differences in the underlying order-size distributions. If the data had been collected over a longer time period, more large orders would presumably have been executed in dollar-yen and sterling-dollar and kurtosis for those orders would have been correspondingly large.

To evaluate the contribution of order-size kurtosis to order-flow kurtosis, a convenient conceptual baseline is the outcome when just one order is executed per period. Specifically, suppose one signed order size is picked at random each period, where a positive (negative) order represents a customer buy (sell), and assume that order flow uniquely determines returns. In this case the distribution of order flow is isomorphic to the distribution of order sizes. Now suppose $N > 1$ orders are executed each period, so aggregate order flow becomes the sum of randomly-selected signed order values. By the central limit theorem, sample order-flow kurtosis declines toward a limit of 3.0 with rising N .

We evaluate the speed with which kurtosis declines using Monte Carlo simulations. For each currency we create thirty separate series of orders, sampling order sizes at random from that currency's order-size distribution. We take one "period" to represent a half hour, since this appears to represent the time until a transaction has its maximum exchange-rate impact (Payne and Vitale, 2003). To simulate five years of 24-hour trading days, each order series lasts 62,400 periods. The resulting kurtosis values, reported in Table 3, are calculated as the average of sample kurtosis across the thirty simulations. Order-flow kurtosis in euro-dollar falls from over 500 with one order per period to 105 with five orders per period, 55 at 10 orders per period, and

8 at 100 orders per period (all values are significantly above three). Other currencies behave similarly.

For our sample period, what would be a realistic estimate of price-contingent orders executed per half hour? Since no data on price-contingent currency orders are available beyond those examined here, we construct the following educated guess. On average, during the more recent sample period, 32 new euro-dollar orders were placed with the bank each day, of which an average of 8.4 were executed. The bank informally estimates that it captured about four percent of the world's currency business at that time. This suggests that roughly 210 euro-dollar price-contingent orders were executed per day across the entire market. 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 have been about four in euro-dollar. Similar back-of-the-envelope calculations yield five and three executed orders per half hour for dollar-yen and sterling-dollar respectively. These values are assumed in our baseline analysis, though we later consider higher order numbers. As shown in Table 3, at these values order flow kurtosis is 129.8 for euro-dollar, 9.5 for dollar-yen and 10.3 for sterling-dollar.

3.2. Order clustering by time of day

In addition to variation in order sizes, unevenness in order flow is due in part to intraday variation in the number of orders executed per period. As shown in Figure 2A, the share of a given day's euro-dollar orders executed per hour ranges from about zero during the market's short "overnight" period (20 to 23 hours GMT) to ten percent around 15 hours GMT. Intraday order execution patterns are qualitatively similar for dollar-yen and sterling-dollar.

One factor that determines order execution is the pattern of order origination. As shown in Figure 2B, this depends noticeably on the biological imperatives to eat and sleep: order

origination is extremely low during the overnight period, for example, when traders everywhere are at home. Order execution also depends on intraday variation in price volatility. Since each order is triggered by the exchange rate's arrival at a specific price, a bigger exchange-rate move triggers more orders, *ceteris paribus*. As shown in Figure 2C, the number of exchange rates crossed per half-hour for euro-dollar varies from a low of three during the market's overnight period to a high of twelve at mid-afternoon in London. Corresponding patterns for dollar-yen and sterling-dollar are again similar. Differences between the intraday patterns of order execution and exchange-rate volatility correspond well to features of the pattern of order origination. During the overnight period, for example, volatility never disappears entirely but the share of executed orders – like the share of orders originated – effectively reaches zero.

This unevenness in order flow should raise order-flow kurtosis. Since order flow represents the sum of many order sizes, periods of high order flow should have high variance and periods of low order flow should have low variance. Even if the distribution of order sizes were normal – in which case the conditional distribution of order flow would be normal each period – the unconditional distribution of order flow would have fat tails, since mixing normal conditional distributions with the same mean but different variances creates an unconditional distribution with high kurtosis.

To evaluate the contribution of intraday variation in order execution to order-flow kurtosis, we turn again to Monte Carlo simulations calibrated to the underlying data. For each currency we draw N orders at random each period. To isolate the contribution of order-flow variation we draw simulated order sizes from a normal distribution with the same mean and standard deviation as the original distribution for that currency, rather than the observed distribution. The share of orders in a given half-hour of each simulated trading day is calibrated

to the frequency with which orders were executed during that period in the original data. As before, we create 30 series of 62,400 periods.

In these simulations, sample order-flow kurtosis at the half-hour horizon averages 4.0, 3.8, and 4.4 for euro-dollar, dollar-yen, and sterling-dollar, respectively, with all three values significantly different from three. Since these values are far below those associated with kurtosis in the order-size distributions, one would tend to infer that intraday variation in executed order flow contributes less to exchange-rate kurtosis than the distribution of order sizes. Our later analysis indicates that when other factors are operative intraday variability in order flow makes a substantial indirect contribution to kurtosis.

3.3. Order clustering by exchange-rate level

Unevenness in the flow of orders will also be driven by the exact level of the exchange rate, because the preference for round numbers evident in the distribution of order sizes is also evident in the distribution of trigger rates. Consider the final two digits of the trigger rate (for the exchange rates \$1.2345/€ and ¥123.45/\$ the last two digits would be 45). If traders had no special preference for certain numbers, the frequency of each two-digit combination would be one percent. As shown in Figure 3, however, there are big spikes in the frequency distribution of trigger rates. Almost six percent of executed orders end in 00 and on average about three percent of orders ending in each of the ten combinations ending in 0. Trigger rates also tend to cluster at levels ending in 5, for which the average frequency is about two percent. Finally, trigger rates end more frequently in 2, 3, 7, and 8 than in 1, 4, 6, and 9 (Osler, 2003).

This order clustering implies that the number of orders executed per period depends on the particular exchange rate levels just crossed, where each level represents one “pip” of the exchange rate as conventionally quoted. (In euro-dollar, for example, 1.2224 and 1.2225 are

separated by one pip.) If the rate crosses a level ending in 0, many orders will be executed; if it crosses a level ending in 1, few orders will be executed. Variation in order numbers clearly makes order flow uneven, raising the possibility of extreme returns. Technically, as with time-of-day clustering, the conditional standard deviation of order flow will be large (small) when many (few) orders are executed. Even if the order-size distribution were Gaussian, the unconditional distribution of order flow would be a mixture of normal distributions with shared mean but differing conditional variances. The unconditional distribution would therefore have fat tails and high kurtosis.

The dependence of order execution on exchange rates also works through order type. Any exchange-rate move triggers both stop-loss and take-profit orders, and one type will require purchases while the other will requires sales. A rising rate, for example, triggers take-profit sell orders and stop-loss buy orders. If stop-loss sell orders cluster at the same rates as take-profit buy orders, mean order flow will likely be small because the orders will tend to offset each other. If these two order types cluster differently, however, order flow will sometimes include many of one order type but few of the other, generating an extreme return. As indicated in Table 4A, stop-loss and take-profit orders cluster differently. Above levels ending in 00, for example, there are relatively few stop-loss sell orders and relatively more take-profit buy orders (Osler, 2003).

To evaluate the contribution of exchange-rate clustering to order-flow kurtosis we turn once again to calibrated Monte Carlo simulations in which we assume that order sizes are normally distributed with a mean and standard deviation that match the underlying dataset. For each simulation we choose a number of levels crossed each period, L . For each period within a simulation we choose a random initial two-digit exchange-rate level, Z , and a random direction

of change (up or down with equal probability). If L is four, Z is 27, and the direction of change is down, the exchange rates crossed that period will end in 27, 26, 25, and 24.

To generate order flow for a given period we randomly select L sets of one hundred order sizes from the appropriate normal distribution. To match the original data, forty-three of every 100 are assigned to be stop-loss orders and the others are assigned to be take-profit orders. The orders are signed to fit market conditions; if the rate is falling, for example, stop-loss (take-profit) orders are assigned a negative (positive) sign. For each order k associated with level Z we draw a random number X_k from $U[0,1]$. We select the subset of these orders included in order flow by comparing each order's X_k to the share of that order type and direction with trigger rate ending in Z . Let Y_{Z}^{AB} represent the share of order-type A , direction B , with trigger rate ending in Z . An order of type AB is included in order flow if $X_k < Y_{Z}^{AB}$. We repeat this process for the remaining $L-1$ levels crossed in the period. Aggregate order flow for the period is the sum of (signed) sizes of all included orders.

We examine $L = 1, 2, 4, 6$, and 10. As before, each simulation has 62,400 periods and we run 30 simulations for each case. With $L = 1$ average sample kurtosis for euro-dollar is 10.7 with standard error 0.08 (Table 4B). Mean order-flow kurtosis for dollar-yen and sterling-dollar are 10.4 and 12.8, respectively.

In theory, rising L should moderate the kurtosis associated with exchange-rate clustering. If the rate rises through levels ending in 49 and 50, for example, few orders will be triggered by the rate ending in 49 but many orders will be triggered by the rate ending in 50. The number of orders executed will therefore be closer to the mean than it would be if the rate crossed only one of those levels. If the exchange-rate levels 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, L , rises towards infinity. If $L = 100$, for example, the rate would cross every 2-digit level in every period and there would be no exchange-rate clustering. Together, the facts that (i) $L \ll 100$ and (ii) the levels must be adjacent to each other slow the convergence of the distribution towards the normal. As this analysis predicts, simulated sample kurtosis is inversely related to the number of exchange rates crossed per period and approaches three as L rises (Table 4B).

Euro-dollar exchange rates cross an average of 6.4 levels per half-hour, corresponding values for dollar-yen and sterling-dollar are 7.5 and 7.7, respectively. At those levels, average order-flow kurtosis in our simulations remains significantly above three, though not by much. When six exchange rates are crossed per period, for example, order-flow kurtosis is 3.9, 3.8, and 4.0 for euro-dollar, dollar-yen, and sterling-dollar, respectively (with standard errors of 0.01). We conclude that traders' exchange-rate preferences have a smaller direct contribution to kurtosis than fat tails in the distribution of order sizes, and that their direct contribution is comparable to that of intraday variation in order execution.

3.4. *Interactions*

We next examine order-flow kurtosis when all three of the properties of price-contingent orders analyzed so far are active simultaneously. We use Monte Carlo simulations because first principles are not sufficient to anticipate the qualitative results. The three properties could conceivably offset each other's effects, bringing smoother order flow akin to the way combining uncorrelated variables brings lower variance, or they could intensify each other's effects.

According to the simulations, the three properties intensify each other. As shown in Table 5, average order-flow kurtosis at the half-hour horizon when the factors interact is 305.3 for euro-dollar, 20.3 for dollar-yen, and 23.8 for sterling-dollar. These values substantially exceed

the sum of the three factors' direct contributions to kurtosis, which are 138.2, 17.6, and 19.1, respectively. We conclude that time-of-day clustering and exchange-rate clustering make sizeable indirect contributions to excess kurtosis, since without these two factors kurtosis would be the lower value associated with the size distribution of orders – 129.8, 9.5, and 10.3, respectively for the three currencies. Equivalently, we conclude that interactions among the properties of order flow jointly constitute an important source of extreme returns.

Intuitively, this seems to make sense. If order sizes were uneven but exactly L orders were executed per period, the most extreme order flow would involve L of the biggest (absolute) order size, all in one direction. When time-of-day and exchange-rate clustering also influence order flow, its extremum will involve more orders and will thus be more extreme. This could occur, for example, if the rate crosses a level ending in 00 at mid-afternoon London time, when Europe and North America are both active.

All three order-flow factors imply that the distribution of order flow converges towards the normal and order-flow kurtosis declines as time horizon lengthens. When the time horizon increases from a half-hour to an hour, for example, kurtosis should be reduced by a doubling in the average number of orders executed per period; by moderation in the intraday seasonality of volatility; and by a rise in the number of exchange rates crossed per period. Consistent with this analysis, simulated order-flow kurtosis decreases monotonically as time horizon lengthens (Table 5). The rate of decline is initially quite rapid: as the horizon moves from a half to a full hour, kurtosis drops by roughly half. Nonetheless, order-flow kurtosis remains statistically significantly above three even at the two-day horizon. Feedback effects between order flow and returns, to which we turn next, introduce the potential for non-monotonicity in this relationship.

4. Feedback effects

This section first examines how feedback between order flow and exchange-rate returns increases the kurtosis of each. It then shows that our four properties of price-contingent orders, when they operate simultaneously, can account for a large portion of excess return kurtosis.

4.1. *Price-contingent orders and exchange-rate trends*

Price-contingent order flow and returns influence each other in three ways that generate unevenness in order flow: price cascades, price halts, and induced randomness in volatility.

Price Cascades: Price cascades are self-reinforcing price moves (Morris and Shin, 2004; Osler, 2005). These are often ascribed to stop-loss orders and other sources of positive-feedback trading, such as the portfolio insurance that contributed to the 1987 stock market crash (Genotte and Leland, 1990). To illustrate, an exchange-rate decline might trigger stop-loss sell orders, which would cause the rate to fall further, which would trigger more stop-loss sell orders, etc. According to currency-market participants, price cascades are extremely fast, even “gappy” (meaning the exchange rate skips some levels), and they happen perhaps a few times per week. Osler (2005) presents formal evidence for the existence of price cascades in currency markets.⁹ Price cascades increase the fat tails of the return distribution at horizons of two periods or longer.

Price Halts: Price halts occur when a rate change triggers the execution of take-profit orders which, since they involve negative-feedback trading, impede further trending. For example, a rate decline in period 1 could trigger a cluster of take-profit buy orders in period 2 that stops or reverses any further decline in that period. With little change in the exchange rate in period 2, few price contingent orders would be executed in period 3, so the price could continue

⁹ If 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 closely by banks that hold the orders, to protect their clients.

to change little. Osler (2005) presents evidence for this effect. Price halts raise kurtosis at horizons of two periods or longer by increasing the frequency of small returns.

Random exchange-rate volatility: In the simulations of Section 3.2., euro-dollar always crossed exactly twelve levels per half hour at 4 p.m. London time and four levels per half hour during overnight trading. In reality, however, the number of levels crossed per half hour is random: sometimes it will cross fewer than four levels during overnight trading and sometimes it will cross more than twelve levels at 4 p.m. London time. This widens the range of intraday return volatility and, by the mixture-of-distributions effect, there should also be a rise in sample kurtosis of both order flow and returns.

4.2. *Dynamic simulations*

To measure the influence of feedback effects on kurtosis we turn again to Monte Carlo simulations. We first analyze how feedback from order flow to returns influences kurtosis in isolation from the three order-flow factors considered previously, and then analyze the joint effects of all four factors. In these simulations, the distribution of order flow in period t is determined by the actual exchange rates crossed in period $t-1$, consistent with the earlier simulations. In addition, order flow in period t influences contemporaneous returns:

$$s_{t+1} - s_t = \lambda(\text{OrderFlow}_t), \quad (1)$$

where s_t is the log of the exchange rate and $\lambda(\cdot)$ is a function.¹⁰

These dynamic simulations are calibrated to match five critical features of trading:

1. *The average number of orders executed per period and its intraday distribution.*
2. *The average number of exchange-rate levels crossed per period and its intraday distribution.*

¹⁰ In reality, quoted rates have only five significant digits. Therefore we set period $t+1$'s quoted exchange rate by taking the preliminary log value calculated in Equation (1), inverting that value, and rounding it off.

3. *The frequency distribution of order execution with respect to two-digit exchange-rate levels across the four order types (slb, sls, tpb, tps).*
4. *The share of take-profits within all executed orders.*
5. *The share of half-hour exchange-rate changes that equal zero*, which was 8.2 percent for euro-dollar, 6.3 percent for dollar-yen, and 8.0 percent for sterling-dollar.

We initially assume a linear relationship between order flow and returns, consistent with most of the empirical literature (e.g., Evans and Lyons, 2002). To focus on the influence of feedback effects, we assume a normal distribution for order sizes and exclude time-of-day and exchange-rate clustering. Sample return kurtosis at the half-hour horizon averages 12.5 for euro-dollar, 11.0 for dollar-yen and 14.0 for sterling-dollar, respectively (Table 6, Panel A). These values only reflect volatility randomness, since price cascades and price halts are multi-period phenomena. Return kurtosis at the 1-hour horizon, which includes price cascades and price halts, is 10.0 for euro-dollar, 10.1 for dollar-yen, and 13.7 for sterling-dollar. Clearly the effect of order-flow feedback on returns is substantial.

We next re-introduce the three order-flow factors. The simulated and original intraday seasonals in the number of executed orders and volatility are juxtaposed in Figures 4A and 4B. The simulated seasonals for euro-dollar match the true seasonals reasonably closely, given that we do not impose these distributions directly. The calibration accuracy is similar for dollar-yen and sterling-dollar. Figure 5 plots the exchange-rate path generated by one of these simulations. Price cascades are apparent as occasional vertical lines while price halts are visible as stretches where the rate changes little.

In these comprehensive linear simulations, average kurtosis is extremely high (Table 6, Panel B). For euro-dollar, kurtosis of half-hour returns is 945.8 with standard error 241.1. This

exceeds kurtosis due to the dynamic factors alone (12.5), it exceeds the sum of this value and kurtosis due to the three order-flow factors (317.5), and it certainly exceeds kurtosis in the original exchange-rate returns (18.7). A similar conclusion emerges for dollar-yen and sterling-dollar, for which simulated half-hour return kurtosis in these comprehensive simulations is 98.9 and 156.9, respectively. The dramatic increases in kurtosis relative to all previous simulations, which is mirrored in similarly high levels of order-flow kurtosis, indicate once again that interactions among our four factors – or equivalently, the factors’ indirect effects – are important in their own right.

The distribution of returns in these simulations could diverge from the normal not just in having high kurtosis but also in being lumpy. The distributions of both order sizes and trigger rates are spiky due to the common preference for round numbers. If the number of price-contingent orders is small, the order-flow and return distributions from combining these factors – especially in the tails – may be difficult to capture accurately by the smooth distributions proposed in the literature, such as the mixture of normals or the stretched exponential.

The level of return kurtosis in these simulations would be disconcerting if the simulations’ linear structure were realistic. But order flow probably has a concave relation with returns because of order splitting (Bertsimas and Lo, 1998). This practice, common in equity markets, is also the standard approach to handling large orders in currency markets.¹¹ Hasbrouck (1991) provides direct evidence for a concave relationship in his VAR treatment of the U.S. equity market. Berger et al. (2008) provide indirect evidence of a concave relationship in currency markets when they show that the proportionate price impact of interdealer order flow declines with the amount. Direct evidence for currency markets using individual orders is

¹¹ Note that this is concavity of price with respect to order flow – which has a time dimension – not concavity of price with respect to individual trade sizes.

provided in Osler and Vandrovych (2009), who find that the square root function best captures the relation between order size and price impact. In the absence of other evidence, we adopt that same functional form.

Table 6, Panel C shows that when order flow has a concave impact on returns, order-flow kurtosis remains high but return kurtosis is dramatically lower than it was in the linear simulations though still significantly above three. At the 1-hour horizon, which includes price cascades and price halts, return kurtosis is 10.5 for euro-dollar, 7.7 for dollar-yen, and 6.8 for sterling-dollar.¹² For all three currency pairs, simulated 1-hour return kurtosis is now meaningfully lower than observed kurtosis, leaving room for contributions from other factors. Nonetheless, simulated excess kurtosis of 1-hour returns represents 69 percent of observed excess return kurtosis for euro-dollar, 52 percent for dollar-yen, and 65 percent for sterling-dollar. This suggests that price-contingent trading is an important source of the high frequency of extreme exchange-rate returns.

In the non-linear simulations return kurtosis is not monotonically related to time horizon. Instead, it rises from the half-hour to the hour horizon before beginning its convergence to the limit of three. These increases can be attributed to price cascades and price halts, whose influence becomes more apparent when the concave functional form compresses the influence of order flow per se.

Table 7 presents the frequency with which 1-hour returns of varying sizes (relative to the standard deviation) would be observed in these simulations. As predicted by the high kurtosis of returns, the frequency of extreme returns is substantially higher than would occur under the normal distribution. One-hour returns of 3.5 to 4.5 standard deviations can be expected to occur

¹² The calibration of the number of orders per period and the number of exchange-rate levels crossed per period is about as accurate in the non-linear simulations as the linear simulations. Details available upon request.

between 4.8 and 7.6 times more frequently than under the normal while those larger than 5.5 standard deviations would occur between 80 and 90 times more frequently than under the normal distribution. To be more concrete, a 5.5-standard-deviation 1-hour return in euro-dollar would be 0.6 percent. Returns of this or larger magnitude could be expected 26 times per year, or almost twice each month. Similarly, a 5.5-standard-deviation hourly return in dollar-yen would be 0.7 percent; returns of this or larger magnitude could be expected 115 times per year, or roughly twice per week.

5. Extensions

This section examines three additional influences on return kurtosis: higher order numbers, news, and non-price-contingent trading.

5.1. Higher trading volume

Considering the influence of higher order numbers is important for two reasons. First, a substantial fraction of price-contingent trading is never formalized as stop-loss or take-profit orders. Second, foreign exchange trading volume has grown rapidly in recent years. According to the B.I.S. Triennial Survey (2007), spot plus forward trading grew by more than 200 percent between 2001 and 2007. Price-contingent trading probably increased even faster, since volume expanded particularly fast among active currency managers and others who rely heavily on algorithmic trading (Chaboud et al., 2009).

We re-run our comprehensive non-linear simulations of euro-dollar with an average of 11 executed orders per period, rather than the previous average of four per period. Figure 6 compares return kurtosis with high and low order numbers. Kurtosis at the half-hour horizon declines roughly 20 percent from the previous simulations because, as discussed in Section 3, the additional trading shifts the distribution of order flow towards the normal. Kurtosis

at the longest horizons actually increases because additional trading intensifies the feedback between order flow and returns. Though these changes are visually small, they represent substantial proportionate increases in excess kurtosis.

5.2. *News*

To evaluate the contribution of news to return kurtosis, we re-run the comprehensive non-linear simulations assuming each simulation begins on January 2, 2004 and including actual news surprises over the ensuing five years. Our sample of news includes the eight U.S. macroeconomic variables found to have a significant effect on dollar exchange rates in Savaser (2010) and Andersen et al. (2007).¹³ This is, of course, only a subset of all the relevant news: much potentially relevant news does not arrive as a scheduled macroeconomic statistical release (Dominguez and Panthaki, 2006). Some excluded news categories may not have huge effects on exchange rates. Fed Funds target changes, for example, have typically caused less trading volume and volatility in recent years, presumably because in this era of central bank transparency the Fed's own "guidance" usually succeeds in helping traders anticipate these changes (Berger et al. 2008). Surprisingly, most foreign macro statistical releases also have relatively modest effects on dollar exchange rates (e.g., Ito and Roley 1987, Almeida et al. 1998).

For each announcement, we take the standardized surprise (surprise divided by the standard deviation of surprises for that variable) and multiply it by the associated estimated response coefficient. The estimates are based on five-minute responses, but they should capture permanent effects since they reflect the arrival of fundamental information. For euro-dollar we

¹³ The eight announcements are: Durable goods orders, the trade balance, non-farm payrolls, retail sales, consumer confidence, the National Association of Purchasing Managers Index, initial claims for jobless benefits, and the advanced estimate of real GDP growth.

take the coefficients from Savaser (2010), whose sample covers the same period as ours. For dollar-yen and sterling-dollar we take coefficients from Andersen et al. (2007). We add the estimated news responses on the dates and times corresponding to each announcement.

As one might expect, news tends to raise return kurtosis at horizons of six hours or less but otherwise has little effect (Table 8): the rise is economically small and usually insignificant. The simulated effects of news would presumably be more substantial if we could examine the universe of news rather than a subset.

5.3. Non-price-contingent trading

Price-contingent trading is only a portion of total foreign-exchange trading, so we finish by considering how trading that is not price-contingent influences return kurtosis. We can rule out any influence through exchange-rate clustering and feedback effects, which by definition do not apply to non-price-contingent trading. Analyzing the other two factors – the size distribution of trades and the intraday distribution of order volume – is challenging because relevant data are scarce. We rely on data from EBS, one of the two dominant electronic interdealer brokerages. Interdealer trading accounts for roughly half of all foreign-exchange trading (Bank for International Settlements, 2007).

Kurtosis in the size distribution of EBS trades in 1999, a period that overlaps our first sample, is 98.8. This factor could therefore be strongly influential, given our earlier conclusion that the size distribution of price-contingent orders makes the strongest direct contribution to excess kurtosis. Intraday EBS trading volume, which is depicted in a number of existing studies (e.g., Berger et al., 2008), follows roughly the same oscillating path as intraday return volatility (Figure 2C). As noted in Section 3, heightened variation in order flow across the trading day

directly increases the frequency of extreme returns. Our analysis indicates that this property has a small direct effect but a considerable indirect effect on the frequency of extreme returns.

Overall, it appears likely that non-price-contingent foreign-exchange trading raises return kurtosis through two of the four dimensions highlighted earlier, and that its contribution could be substantial.

6. Conclusion

This paper examines the forces behind exchange-rate crashes, focusing on patterns in the placement and execution of price-contingent orders. We highlight four properties of such trading that contribute to a high frequency of extreme returns and, more broadly, to high return kurtosis. These are: (1) the distribution of order sizes, which itself has high kurtosis; (2) time-of-day clustering in order execution; (3) exchange-rate clustering of order execution; and (4) feedback from order execution to returns, which generates price cascades and price halts.

We evaluate the relative contributions of these factors using simulations calibrated to the properties of price-contingent orders at the Royal Bank of Scotland, one of the world's largest dealing banks. When the factors are taken in isolation from each other, the single most important factor is leptokurtosis in the distribution of order sizes. However, when the factors interact with each other, their interactions generate as much or more kurtosis as any single factor in isolation.

The amount of kurtosis attributable to price-contingent orders is substantial. According to our simulations, price-contingent trading may account for one-half to two-thirds of observed kurtosis in high-frequency returns. Overall, our analysis indicates that price-contingent trading may be a major source of fat tails in returns, and that extreme exchange-rate returns are statistically inevitable even in the absence of news.

This research may be relevant to other markets with active price-contingent trading. Indeed, stop-loss orders have been implicated in the flash crash of May, 2010. Price-contingent

trading is associated with technical analysis, options hedging, and algorithmic trading, all of which are known to be important in equity and futures markets. In addition, price-contingent trading in these markets may have some of the same properties it has in the foreign exchange market. For example, trading clusters strongly at the opening and closing of these markets. Further, equity limit orders tend to cluster at certain price levels (Neiderhoffer and Osborne, 1966; Kavajecz and Odders-White, 2004).

By enhancing our understanding of the sources of extreme currency returns, this analysis may illuminate the forces behind the ambiguous jump processes frequently invoked in the analysis of option pricing. Our analysis may also help risk managers anticipate the likelihood and magnitude of adverse tail events. Finally, our analysis may help explain the crash risk that may help justify high average carry-trade returns in currency markets (Brunnermeier et al., 2008). It also suggests that these equilibria may be partially self fulfilling: crash risk prompts traders to place stop-loss orders on one side of the market and those orders then contribute to crash risk.

REFERENCES

- Abreu, D., Brunnermeier, M., 2003. Bubbles and crashes. *Econometrica* 71, 173-204.
- Akgiray, V., Booth, G. G., 1988. Mixed diffusion-jump process modeling of exchange rate movements. *Review of Economics and Statistics* 70, 631-637.
- Almeida, A., C. Goodhart, R. Payne, 1998. The effects of macroeconomics news on high frequency exchange rate behavior. *Journal of Financial and Quantitative Analysis* 33, 383-408.
- Andersen, T., Bollerslev, T., Das, A., 2001. Variance-ratio statistics and high-frequency data: Testing for changes and in intraday volatility patterns. *Journal of Finance* 56, 305-328.
- Andersen, T., Bollerslev, T., Diebold, F. X., Vega, C., 2007. Real-time price discovery in global stock, bond, and foreign exchange markets. *Journal of International Economics* 73, 251-277.
- Bank for International Settlements, 2007. Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007 – Preliminary global results. Bank for International Settlements: Basle.
- Berger, D., Chaboud, A., Chernenko, S., Howorka, E., Wright, J., 2008. Order flow and exchange rate dynamics in electronic brokerage system data. *Journal of International Economics* 75, 93-109.
- Bertsimas, D., Lo, A. W., 1998. Optimal control of execution costs. *Journal of Financial Markets* 1, 1-50.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655-1689.
- Brunnermeier, M., Nagel, S., Pedersen, L., 2008. Carry trades and currency crashes. NBER Working Paper 14,473.
- Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2010a. Do Peso Problems Explain the Returns to the Carry Trade? NBER Working Paper 14,054.
- Burnside, C., B. Han, D. Hirshleifer, and T. Y. Wang. 2010b. Investor Overconfidence and the Forward Discount Puzzle. *Review of Economic Studies*, Forthcoming.

- CFTC-SEC (Commodities Futures Trading Commission-Securities and Exchange Commission), 2010. Preliminary findings regarding the market events of May 6, 2010.
<http://www.sec.gov/sec-cftc-prelimreport.pdf>
- Chaboud, A., Chiquoine, B., Hjalmarsson, E., Vega, C., 2009. Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. FRB International Finance Discussion Paper 980.
- Covrig, V., and Melvin, M., 2005. Tokyo Insiders and the informational efficiency of the Yen/Dollar exchange rate. *International Journal of Finance and Economics* 10, 185-193.
- Cutler, D. M., Poterba, J. M., Summers, L. H., 1991. Speculative Dynamics. *Review of Economic Studies* 58, 529-546.
- D'Agostino, R. B., Stephens, M. A., 1986. Goodness-of-fit techniques. Marcel Dekker: New York.
- Daniélsson, J., de Vries, C., 1998. Value-at-risk and extreme returns. Tinbergen Institute, Tinbergen Institute discussion Papers 98-017/2.
- Daniélsson, J., Payne, R., 2002. Real trading patterns and prices in spot foreign exchange markets. *Journal of International Money and Finance* 21, 203-222.
- Dominguez, K., Panthaki, F., 2006. What Defines 'News' in Foreign Exchange Markets? *Journal of International Money and Finance* 25, 168-198.
- Euromoney Survey, 2007. FX Poll 2007: Overall market share. <http://www.Euromoney.com/Article/1330390/Article.html>.
- Evans, M., and Lyons, R., 2002. Order flow and exchange-rate dynamics. *Journal of Political Economy* 110, 170–180.
- Fama, E. F., 1965. The behavior of stock market prices. *Journal of Business* 38, 34-105.
- Fernandez, D., 2010. Weekly trend prediction: An adaptive breakout technique analysis. *Currency Trader Magazine* (February), 24-27.
- Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H. E., 2003. A Theory of power-law distributions in financial market fluctuations. *Nature* 423, 267-270.
- Genotte, G., Leland, H., 1990. Market liquidity, hedging, and crashes. *American Economic Review* 80, 999-1021.

- Hasbrouck, J., 1991. Measuring the information content of stock trades. *Journal of Finance* 46, 179–207.
- IIROC (Investment Industry Regulatory Organization of Canada), (2010). Review of the market events of May 6, 2010. <http://docs.iiroc.ca/DisplayDocument.aspx?DocumentID=C6CB526F93254DED9F4F98CDCB935FF3&Language=en>
- Ito, R., and V. V. Roley, 1987. News effects from the U.S. and Japan: Which moves the yen/dollar exchange rate? *Journal of Monetary Economics* 19, 255-277.
- Kavajecz, K., Odders-White, E., 2004. Technical analysis and liquidity provision. *Review of Financial Studies* 17, 1043-1071.
- Laherrere, J., Sornette, D., 1998. Stretched exponential distributions in nature and economy: Fat tails with characteristic scales. *European Physical Journal B* 2, 525-539.
- LeBaron, B., 2008. Robust Properties of Stock Return Tails. Mimeo.
- Levich, R., Pojarliev, M., 2010. Detecting Crowded Trades in Currency Funds. NBER Working Paper 15,698.
- Morris, Stephen, and Hyun Song Shin, 2004, Liquidity Black Holes, *Review of Finance* 8:1-18.
- Niederhoffer, V., Osborne, M.F.M., 1966. Market making and reversals on the stock exchange. *Journal of the American Statistical Association* 61, 897-916.
- Osler, C. L., 2003. Currency orders and exchange-rate dynamics: Explaining the predictive success of technical analysis. *Journal of Finance* 58, 1791-1819.
- _____, 2005. Stop-loss orders and price cascades in currency markets. *Journal of International Money and Finance* 24, 219-241.
- _____, 2009. Foreign Exchange Microstructure: A Survey. In: Robert A. Meyers (Ed), *Encyclopedia of Complexity and System Science*, Springer: New York.
- Osler, C. L., Vandrovych, V., 2009. Which customers bring information to the in foreign exchange market? Presented at the Third Annual Microstructure Workshop, Emerging Markets Group, Cass Business School, London, May 1, 2009.
- Payne, R., Vitale, P., 2003. A transaction level study of the effects of central bank intervention on exchange rates. *Journal of International Economics* 61, 331-52.
- Plantin, G., Shinn, H., 2009. Carry trades and speculative dynamics. Mimeo.

- Roll, R., 1970. The Behavior of Interest Rates: An Application of the Efficient Market Model to U.S. Treasury Bills. Basic Books: New York.
- Savaser, T., 2010. Exchange rate response to macro news: Through the lens of microstructure. Journal of International Financial Markets, Institutions and Money. Forthcoming.
- Shiller, R., 1989. Investor behavior in the October 1987 stock market crash: Survey evidence. In Shiller, Robert, Market Volatility. MIT Press: Cambridge, MA.
- Tucker, A. L., Pond, L., 1988. The Probability distribution of foreign exchange price changes: Tests of candidate processes, Review of Economics and Statistics, 638-647.
- Westerfield, J. M., 1977. An examination of foreign exchange risk under fixed and floating rate regimes. Journal of International Economics 7, 181-200.
- Yule, G. U., 1927. On reading a scale. Journal of the Royal Statistical Society 90, 570-579.

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, dollar-yen, and sterling-dollar currency pairs processed by the Royal Bank of Scotland over the period September 8, 1999, through April 11, 2000, and June 12, 2001, through September 20, 2002. “Distance to market” refers to the distance between an order’s trigger rate and the market mid-quote when the order is placed.

	All orders	Stop-loss	Take-profit
Number of orders	47,312	20,213	27,099
Total value (\$ billions)	253.9	114.6	139.5
Share of orders (%)	100.0	42.7	57.3
Share executed (%)	27	26	29
Mean size (\$ million)	5.4	5.7	5.2
Mean distance to market (%)	0.5	0.5	0.5
Median days open	0.6	0.4	0.7

Table 2: Exchange-rate kurtosis

Data are Reuters mid-quotes sampled at five-minute intervals from January 2, 2000, through November 9, 2002, exclusive of September 13, 2001, through October 8, 2001.

2A: Sample kurtosis

Absolute kurtosis of returns (bold) and standard error (in parentheses).

	15 Min.	30 Min.	1 Hour	2 Hours	6 Hours	12 Hours	24 Hours	48 Hours
EUR	23.8 (0.02)	18.7 (0.03)	13.8 (0.04)	12.3 (0.05)	6.7 (0.09)	5.2 (0.13)	4.1 (0.18)	4.7 (0.26)
JPY	18.1 (0.02)	14.4 (0.03)	11.9 (0.04)	8.8 (0.05)	7.3 (0.09)	6.6 (0.13)	3.9 (0.18)	3.5 (0.26)
GBP	13.3 (0.02)	10.5 (0.03)	8.8 (0.04)	8.1 (0.05)	5.4 (0.09)	5.1 (0.3)	3.7 (0.18)	4.4 (0.26)

2B: Estimated tail indexes for exchange-rate returns

OLS Hill estimates, where k is the fraction of observations on which the estimates are based.

		EUR		JPY		GBP	
		Left tail	Right tail	Left tail	Right tail	Left tail	Right tail
1/2-hour:	k = 0.10	3.58	3.38	3.60	3.47	3.68	3.76
	k = 0.20	3.33	3.25	3.55	3.51	3.54	3.61
12 hours:	k = 0.10	5.34	5.07	5.20	4.39	4.16	5.38
	k = 0.20	4.21	4.41	5.03	4.01	4.05	5.09

2C: Fraction of half-hour euro-dollar returns at varying distances from the mean.

The fraction of returns differing from the mean by a given multiple of the standard deviation is shown as a ratio to the corresponding fraction under the normal distribution.

Standard deviation multiples					
	< 0.5	0.5 to 1.5	1.5 to 2.5	2.5 to 3.5	> 3.5
<i>Half-hour returns relative to normal distribution</i>					
EUR	1.4	0.7	0.6	1.5	19.4
JPY	1.4	0.8	0.6	1.4	17.7
GBP	1.5	0.7	0.6	1.4	13.3
<i>12-hour returns relative to normal distribution</i>					
EUR	1.3	0.8	0.7	2.0	13.6
JPY	1.3	0.8	0.7	1.2	18.1
GBP	1.3	0.8	0.8	2.1	10.6
<i>Share of observations</i>					
Gaussian	38.3	48.4	12.1	1.2	4.7E-2

Table 3: Average order-flow kurtosis with calibrated order-size distribution

The table shows kurtosis of order flow in Monte Carlo simulations with varying numbers of orders executed per period. In each period, order sizes are drawn at random from an underlying sample comprising all euro-dollar, dollar-yen, and sterling-dollar orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. Sample kurtosis (in bold) is calculated for 30 separate simulations per currency, each with 62,400 periods. Standard errors in parentheses.

	Orders per period								
	1	2	3	4	5	10	20	50	100
EUR	513.0 (68.04)	252.4 (41.95)	173.1 (31.52)	129.8 (18.04)	105.3 (15.59)	54.6 (8.31)	28.6 (3.86)	13.2 (1.54)	8.2 (0.70)
JPY	35.9 (3.00)	19.4 (2.01)	14.1 (1.12)	11.3 (0.80)	9.5 (0.67)	6.4 (0.34)	4.7 (0.16)	3.7 (0.07)	3.3 (0.04)
GBP	24.9 (0.62)	14.0 (0.46)	10.3 (0.25)	8.5 (0.23)	7.4 (0.18)	5.2 (0.10)	4.1 (0.06)	3.4 (0.03)	3.2 (0.02)

Table 4A: Asymmetries in the frequency of trigger rates near round numbers

The table shows the percent of executed orders of each order type with requested execution rates ending in the indicated set of two-digit numbers (weighted by number of orders). The data comprise all euro-dollar, dollar-yen, and sterling-dollar stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002.

	Stop-loss orders		Take-profit orders	
	Buy	Sell	Buy	Sell
Just above 00: 01-10	10.9	5.9	11.8	9.3
At 00	5.0	6.1	5.2	6.1
Just below 00: 90-99	6.7	11.3	11.0	8.8
Just above 50: 51-60	11.9	7.8	9.9	9.2
At 50	4.7	6.0	4.4	4.1
Just below 50: 40-49	7.3	14.1	8.6	10.0

Table 4B: Order-flow kurtosis with calibrated frequencies of trigger rates

The table shows kurtosis from Monte Carlo simulations of aggregate order flow when the frequency of trigger rates is calibrated to its original distribution. For each period, the initial exchange rate and the direction of movement are chosen randomly. Sample kurtosis (in bold) is calculated for 30 separate simulations per currency, each with 62,400 periods. Standard errors in parentheses.

	Number of exchange-rate levels crossed per period				
	1	2	4	6	10
EUR	10.7	6.1	4.4	3.9	3.5
	(0.08)	(0.02)	(0.01)	(0.01)	(0.01)
JPY	10.4	5.8	4.3	3.8	3.5
	(0.06)	(0.03)	(0.01)	(0.01)	(0.01)
GBP	12.8	7.0	4.5	4.0	3.5
	(0.09)	(0.03)	(0.01)	(0.01)	(0.01)

Table 5: Average order-flow kurtosis with calibrated order flow.

The table shows kurtosis of aggregate order flow from Monte Carlo simulations calibrated to match the order-size distribution, the intraday variation in executed orders, and the frequency distribution of trigger rates in our original data. Original data comprise all euro-dollar, dollar-yen, and sterling-dollar stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. Sample kurtosis (in bold) is calculated for 30 separate simulations per currency, each with 62,400 periods. Standard errors in parentheses.

	Time horizon						
	1/2 Hour	1 Hour	2 Hours	6 Hours	12 Hours	24 Hours	48 Hours
EUR	305.3 (14.63)	155.9 (7.58)	80.0 (3.88)	28.7 (1.40)	16.3 (0.87)	9.3 (0.49)	5.8 (0.23)
JPY	20.3 (0.49)	12.0 (0.16)	7.8 (0.13)	4.9 (0.05)	4.1 (0.06)	3.3 (0.04)	3.2 (0.04)
GBP	23.8 (0.40)	13.8 (0.16)	8.9 (0.09)	5.7 (0.07)	4.9 (0.07)	3.4 (0.04)	3.2 (0.03)

Table 6: Return kurtosis with dynamic interactions between order flow and returns

The table shows kurtosis of exchange-rate returns from Monte Carlo simulations in which the number of executed orders and exchange-rate returns influence each other. In Panel A the simulations assume that order flow has a linear effect on returns, that order sizes are normally distributed, and that order flow is not influenced by time of day or by the exchange-rate's current level. Panel B modifies the simulations of Panel A to incorporate the true size distribution of orders and the influence of both time-of-day and exchange-rate levels. Panel C modifies the simulations of Panel B by assuming that the square root of order flow influences returns. Simulations are calibrated to a dataset comprising all euro-dollar, dollar-yen, and sterling-dollar stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. Simulations are also calibrated to the intraday properties of half-hour returns over the period January 2, 2000 through November 9, 2002. Sample kurtosis (in bold) is calculated for 30 separate simulations per currency, each with 62,400 periods. Standard errors in parentheses.

	Time horizon						
	1/2 Hour	1 Hour	2 Hours	6 Hours	12 Hours	24 Hours	48 Hours
<i>A: Linear, no order-flow factors</i>							
EUR	12.5 (0.16)	10.0 (0.09)	9.9 (0.16)	6.1 (0.09)	4.7 (0.06)	3.9 (0.08)	3.5 (0.06)
JPY	11.0 (0.18)	10.1 (0.09)	7.9 (0.10)	5.2 (0.06)	4.2 (0.06)	3.6 (0.06)	3.3 (0.06)
GBP	14.0 (0.29)	13.7 (0.26)	9.5 (0.17)	5.7 (0.13)	4.4 (0.08)	3.9 (0.08)	3.4 (0.06)
<i>B: Linear, comprehensive</i>							
EUR	945.8 (241.14)	497.1 (82.06)	271.2 (39.78)	107.1 (16.41)	58.1 (8.84)	30.8 (4.55)	16.4 (2.15)
JPY	98.9 (13.09)	77.8 (6.77)	48.9 (4.94)	18.6 (1.13)	11.7 (0.72)	7.2 (0.34)	5.3 (0.21)
GBP	156.9 (16.72)	125.8 (15.20)	76.1 (8.64)	33.0 (4.04)	17.6 (1.49)	10.5 (0.83)	6.8 (0.51)
<i>C: Concave, comprehensive</i>							
<i>Return kurtosis</i>							
EUR	7.8 (0.11)	10.5 (0.19)	8.5 (0.09)	5.8 (0.09)	4.8 (0.07)	3.7 (0.05)	3.4 (0.05)
JPY	4.9 (0.03)	7.7 (0.06)	6.6 (0.05)	5.3 (0.07)	4.7 (0.08)	4.1 (0.08)	3.8 (0.07)
GBP	5.1 (0.03)	6.8 (0.04)	6.1 (0.05)	4.9 (0.05)	4.3 (0.05)	3.5 (0.04)	3.3 (0.05)
<i>Order-flow kurtosis</i>							
EUR	506.4	408.2	245.5	92.5	47.7	25.0	13.7
JPY	33.0	26.8	17.9	9.6	6.5	4.7	3.8
GBP	35.8	27.7	18.3	9.6	6.8	4.6	3.8

Table 7: Fraction of returns at varying distances from the mean.

The fraction of 1-hour returns differing from the mean by a given multiple of the standard deviation is shown as a ratio to the corresponding fraction under the normal distribution. Returns generated in Monte Carlo simulations in which the square root of order flow influences returns, orders are calibrated to their true size distribution, and order flow is influenced by both time-of-day and the exchange-rate level.

	Standard deviation multiples						
	< 0.5	0.5 to 1.5	1.5 to 2.5	2.5 to 3.5	3.5 to 4.5	4.5 to 5.5	5.5 to 6.5
EUR	1.2	0.9	0.9	1.7	7.6	75.9	843.5
JPY	1.1	0.9	0.9	1.4	4.8	18.7	1889.2
GBP	1.1	0.9	0.9	1.6	6.3	61.6	5060.8

Table 8: Average return kurtosis with news

The table shows kurtosis of exchange-rate returns from Monte Carlo simulations in which returns respond to realized U.S. macroeconomic statistical releases for the five years from 2004 through 2008. The square root of order flow influences simulated returns, simulated orders are calibrated to the true size distribution, and simulated order flow is influenced by both time of day and exchange-rate level. Simulations are also calibrated to the intraday properties of half-hour returns over the period January 2, 2000 through November 9, 2002. Sample kurtosis (in bold) is calculated for 30 simulations per currency, each with 62,400 periods. Average kurtosis with news minus average kurtosis without news in italics. † indicates 5-percent significance.

	Time horizon						
	1/2	1	2	6	12	24	48
	Hour	Hour	Hours	Hours	Hours	Hours	Hours
EUR: With news	8.02	10.53	8.26	6.24	5.07	3.75	3.33
<i>News - no news</i>	<i>0.18</i>	<i>0.07</i>	<i>-0.25</i>	<i>0.42†</i>	<i>0.23</i>	<i>0.04</i>	<i>-0.07</i>
JPY: With News	5.12	6.76	6.00	4.81	4.36	3.54	3.31
<i>News - no news</i>	<i>0.01</i>	<i>0.00</i>	<i>-0.09</i>	<i>-0.07</i>	<i>0.11</i>	<i>0.02</i>	<i>0.01</i>
GBP: With news	5.01	7.82	6.57	5.36	4.54	4.04	3.78
<i>News - no news</i>	<i>0.08</i>	<i>0.17†</i>	<i>0.01</i>	<i>0.04</i>	<i>-0.17</i>	<i>-0.06</i>	<i>-0.03</i>

Figure 1: Distribution of executed order sizes

The figure shows the frequency distribution of order sizes for executed euro-dollar stop-loss and take-profit 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 24 months of the period September, 1999, through September, 2002. Horizontal axis truncated at €50 million for clarity.






Figure 2: Intraday patterns of order execution, order origination, and exchange-rate volatility

Underlying Figures 2A and 2B are the complete record of executed stop-loss and take-profit orders processed by the Royal Bank of Scotland during 24 months of the period September, 1999 through September, 2002. Figure 2A shows intraday variation in the number of executed orders. Figure 2B shows intraday variation in the number of placed orders. Figure 2C shows intraday variation in the volatility of returns, measured as the number of exchange rates crossed per half hour. Underlying data for Figure 2C are five-minute Reuters exchange rate mid-quotes during the period January 2, 2000 through November 9, 2002.



2C: Exchange-rate volatility




Figure 3: Frequency distribution of trigger rates of price-contingent orders by final two digits

Underlying data are the complete record of stop-loss and take-profit orders in euro-dollar, dollar-yen, and sterling-dollar executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. The final two digits are 45 for orders with trigger rates \$1.2345/€, ¥123.45/\$ or, \$1.2345/£. Orders weighted by number.





Figure 4: Intraday seasonals: Realized versus dynamic simulations

The figures compare realized and simulated intraday seasonals in the frequency of euro-dollar order execution and in the number of exchange-rate levels crossed per half hour. In the simulations, order flow influences returns linearly, order sizes are calibrated to the true distribution, and simulated order flow is influenced by both time of day and exchange-rate level. Simulated properties are the averages over 30 simulations, each with 62,400 periods.

4A: Actual vs. simulated orders executed



4B: Actual vs. simulated volatility




Figure 5: Simulated exchange rate path

The figure shows a simulated euro-dollar path from a Monte Carlo simulation in which order flow has a linear influence on returns, orders are calibrated to the true size distribution, and order flow is influenced by both time-of-day and the exchange-rate level. Original order-flow properties taken from a dataset comprising all euro-dollar stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. Simulations are also calibrated to the intraday properties of half-hour euro-dollar returns over the period January 2, 2000 through November 9, 2002. The simulation's 64,200 observations correspond to roughly five years of half-hour time periods.




Figure 6: Return kurtosis under low and high order numbers

The figures compare kurtosis of euro-dollar returns from Monte Carlo simulations in which the average number of orders per period is four (“low”) and eleven (“high”). In Figure 6A the simulations assume that order flow has a linear effect on returns; in Figure 6B the simulations assume that the square root of order flow influences returns. Otherwise, all simulations are calibrated to a dataset comprising all euro-dollar stop-loss and take-profit orders executed by the Royal Bank of Scotland during 24 months of the period September, 1999, through September, 2002. Simulations are also calibrated to the intraday properties of half-hour euro-dollar returns over the period January 2, 2000 through November 9, 2002. Sample kurtosis is calculated for 30 simulations with 62,400 periods each.

