

6. Customer trades and extreme events in foreign exchange

Mintao Fan and Richard K. Lyons*

1. INTRODUCTION

Professor Goodhart's impact has been substantial in many areas. One of those areas is a growing field that uses tools from microstructure finance to analyze exchange-rate behavior (the 'microstructure approach' to exchange rates). This chapter is an introduction to one of the frontiers of that field. It addresses the trading of FX customers: investors, importers, exporters, corporate Treasurers, etc. Past work within the microstructure approach has focused on FX trading between banks, or 'interdealer' trading. It is true that most of the trading in FX is interdealer. Nevertheless, interdealer trading is in a sense derivative: it is the demands of non-dealer customers that represent underlying demand for currencies in the real economy. Data necessary for this analysis have become available only recently, in part due to the wholesale shift to electronic trading (and the data capture it permits).

Our main results include the following. First, we find that aggregate customer order flow in our sample (10–15 per cent of the market-wide total) shows little evidence of mean reversion.¹ Indeed, customer order flow cumulated over time is approximately a random walk. Second, aggregate customer order flow tracks exchange rate movements at lower frequencies (e.g. annual) rather closely. Third, when aggregate customer order flow is disaggregated, we find that the parts behave rather differently. For example, our case study on the 10 per cent drop in the yen/\$ rate that occurred in a single day in 1998 (around the time of the Long Term Capital Management collapse) shows that hedge funds were not the trigger of the collapse, but instead were net providers of liquidity (i.e. net buyers of dollars). The trigger was the portfolio shift of financial institutions like mutual funds, pension funds, and insurance companies. Fourth, we find that extreme exchange-rate movements at high-frequency are generally associated with large net flows from financial institutions; in contrast, low frequency trends are associated with secular net flows from non-financial corporations.

The paper is organized in five sections. In section two we provide back-

ground on the available customer-trade data, including descriptive statistics. Section three introduces hypotheses from recent models of FX customer trading and puts them to the test. Section four is a case study of the behaviour of customer trades around the collapse of the yen/dollar rate in the fall of 1998. Section five examines other extreme movements in the dollar-yen and dollar-euro markets. Section six concludes.

2. BACKGROUND ON THE CUSTOMER TRADE SEGMENT

The role of order flow from non-marketmaker customers is central to microstructure theory. Indeed, customer flow is at the core of all the canonical models (e.g. Kyle 1985 and Glosten and Milgrom 1985, among others). In each of those models it is the customer orders that catalyze a market response. In this respect, it is not unreasonable to view different microstructure models as broadly similar: their basic implications for the relation between exchange rates and customer flow are the same (though the interim paths of price adjustment across models may differ).²

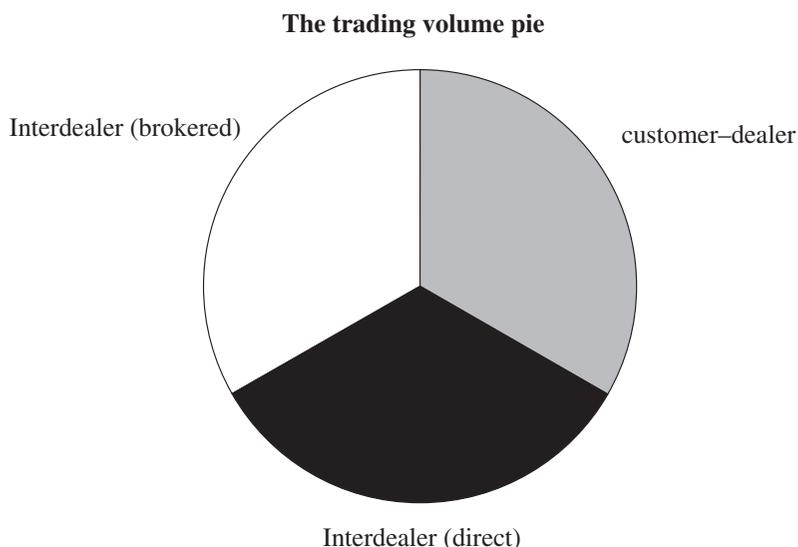
The importance of customer orders is obvious to practitioners as well. Any FX trader or trading-desk manager would agree. One trader we spoke with put it colourfully when he said that customer trades are the market's 'crack cocaine'. By this he meant that the customer orders are a catalyst, and a powerful one at that. In keeping with this notion of customer flow as the market's catalyst, proprietary information on those flows is a prime driver of proprietary trading at the largest banks. (Smaller banks see too little of the marketwide customer flow to make this information useful.) Embedded in this behaviour is the fact that banks find customer-flow information valuable for *predicting* exchange-rate movements. To date, work that adopts the microstructure approach to exchange rates has concerned itself with *explaining* movements, in the sense of accounting for movements using concurrent flow. That customer flow has predictive power as well (i.e. today's flow predicts future movements) adds a new dimension. It is this predictive dimension that most interests the practitioner audience.³

So why has previous literature focused so much on order flow between dealers? There are two reasons. The first is the simple fact that until the dataset described below became available, researchers had no alternative but to work with order flow between dealers. The second reason is that – despite the constraint on data availability – there is valid justification for focusing on flow between dealers. The justification relates to the differential transparency of customer-dealer versus interdealer flow. The reality of the FX market is that dealers do observe some order flow from interdealer

trades – including trades in which they are not involved. Customer-dealer trades, on the other hand, are not observable except by the bank that receives them. Dealers therefore learn about other dealers' customer orders as best they can by observing other dealers' interdealer trades, and they set market prices accordingly. Although this learning from interdealer orders is consistent with empirical models in the existing literature (e.g. Evans and Lyons 2002), the ultimate driver of that interdealer flow is customer flow.

Let us provide a bird's eye view of how the customer-flow data of this chapter relates to the order flow data analyzed in earlier papers. Volume in the major spot FX markets splits into three basic categories: customer-dealer trades, direct interdealer trades, and brokered interdealer trades. Figure 6.1 provides an illustration. The work of Lyons (1995), Evans (2001) and Evans and Lyons (2002), for example, uses data from the direct interdealer category. The work of Goodhart, Ito, and Payne (1996), Goodhart and Payne (1996), Payne (1999), and Killeen et al. (2001) uses data from the brokered interdealer category.

The customer data used in this study are qualitatively different than data used in other work on FX markets. In the past, data on customer orders in



Notes: In the 1990s, when empirical work on FX order flow began to develop, trades of the three basic types had roughly equal shares in total volume. By 2000, electronic interdealer brokers (principally EBS) had taken much market share away from direct interdealer trading. The BIS (2001) now estimates the brokered share of interdealer trading in major currencies at about 90 per cent.

Figure 6.1 Trading volume breakdown in the 1990s

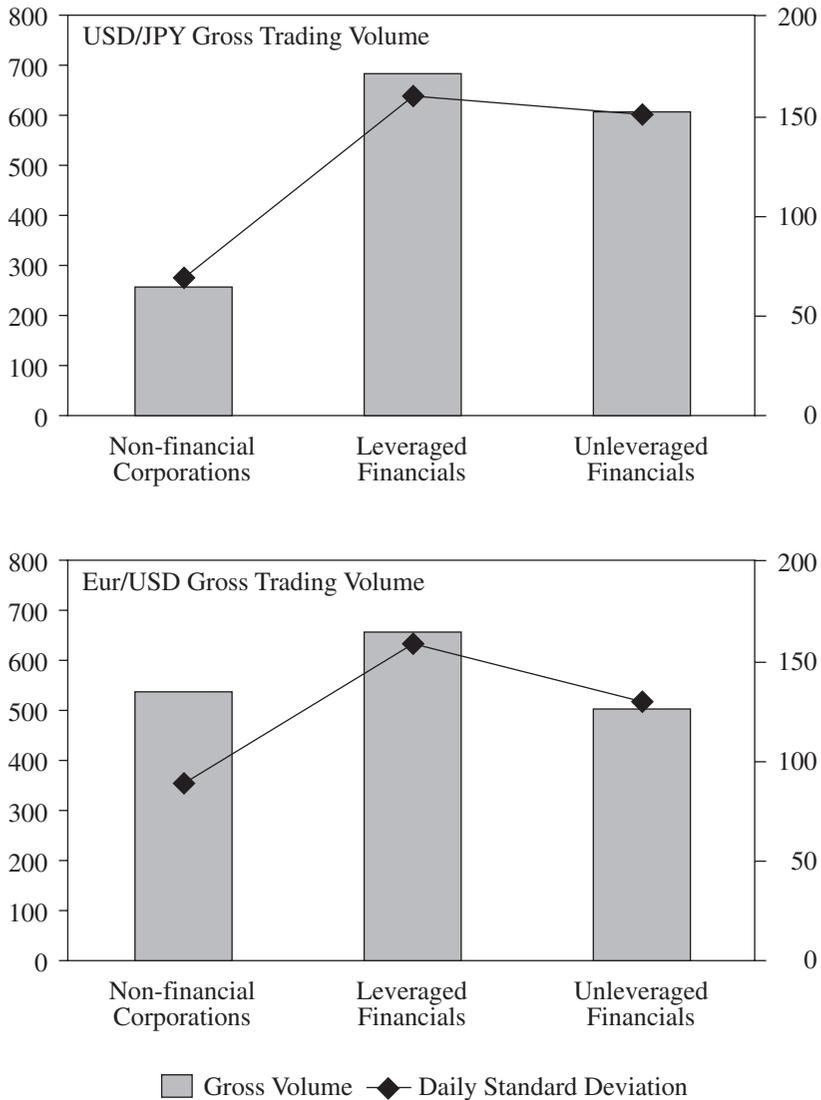
the major markets have been difficult to obtain. The only possible source – given the market’s current structure – is private banks themselves.⁴ But in general these banks consider these data highly proprietary. Recently, however, we obtained customer trade data from a leading FX trading bank. (The bank, Citibank, is among the top three worldwide, with a market share in major-currency customer business in the 10–15 per cent range.) The bank made these data available only on a timeaggregated basis. Specifically, all the customer orders received by this bank worldwide are aggregated into daily order flows (executed trades only). The data set therefore does not include individual trades. Consequently, transaction-level analysis along the lines of that in Lyons (1995) is not possible.

Against this drawback, the data set has many advantages:

- Length: The data span more than five years, so analysis at longer horizons (e.g., monthly) is possible.
- Markets: The data cover the two largest markets: \$/euro and \$/yen. (Before January 1999, order flow data for the ‘euro’ are constructed from flows in the constituent currencies against the dollar.)
- Transaction Types: The data include both spot and forward trades, but are netted of any trades in FX swaps (because FX swaps do not have net order flow implications).
- Components: The data are split into three customer-type categories, corresponding to three categories commonly applied by practitioners: nonfinancial corporations, unleveraged financial institutions (e.g. mutual funds), and leveraged financial institutions (e.g. hedge funds and banks’ proprietary trading desks).

The last advantage provides considerable statistical power for uncovering the underlying causes of order-flow’s impact on price. Are all orders equal in their price impact? Or might some order types – say the orders of hedge funds – convey more information than others? Our ability to disaggregate order flow to answer these questions brings us closer to a specification of the underlying information sources.

Figure 6.2 provides a visual description of the relative importance of these three customer categories.⁵ The sample is January 1993 to June 1999. For the euro, the total trading volume across the three customer categories is roughly balanced. For the yen this is not the case: non-financial corporate trading is less than half that for the other two categories. (These breakdowns may be specific to the data-source bank, however.) For both markets, the daily order flow from nonfinancial corporations is the least volatile. Though not included in the figure, cumulative order flow over the whole sample displays quite different characteristics across the three



Notes: The sample for both currencies is January 1993 to June 1999. (Before the launch of the euro in January 1999, volume and order flow are constructed from trading in the euro's constituent currencies.) Gross volume is the mean daily volume measured in millions of dollars for USD/JPY and millions of euros for EUR/USD. Daily standard deviation measures the standard deviation of daily order flow.

Figure 6.2 Customer trades: relative transaction volumes and transaction volatility

customer categories. For the dollar-euro market, nonfinancial corporations are the largest net sellers of euros and unleveraged financial institutions are the largest net buyers. (We return to this feature of the data in section five.) In the dollar-yen market, unleveraged financial institutions are the largest net sellers of dollars (though slight) and leveraged financial institutions are the largest net buyers.

Order Flows versus Capital Flows

A slight digression on the link between customer order flows and balance of payments flows may be useful. It is important to recognize that balance-of-payments flows and FX transactions are not one-to-one. To understand why, consider an import of \$100 million of Japanese goods into the US by a US multinational. (One could also use a capital-account transaction for this example.) Suppose the transaction is invoiced in yen, but the US multinational pays the invoice from yen it already holds at its Japanese subsidiary. An import is logged, but there is no corresponding order in the FX market; the link is not one-to-one. In time, one would expect some adjustment of the multinational's 'portfolio', but that need not occur immediately, and it need not involve order flow in the FX market (e.g. suppose the Japanese subsidiary responds by increasing its working capital borrowing in yen).⁶ The bottom line is that balance-of-payments flows do not necessarily generate corresponding order flow in the FX market. If gleaning information from order flow is how dealers determine price, then portfolio shifts in the form of balance-of-payments flows will not be counted unless and until they generate order flow.⁷

3. HYPOTHESES AND TESTS

Let us sketch the basics of the model introduced by Evans and Lyons (2002) and use it to formulate testable hypotheses for customer order flow. (The empirical analysis in the Evans-Lyons paper examines interdealer order flow only.) One of the conceptual contributions of the Evans-Lyons model is its explanation for why interdealer flow cumulated over time can follow a random walk, while at the same time individual dealers close out their positions each trading day. Empirically, both of these features are true (to a first approximation) in the data.

Evans and Lyons (2002) consider a pure-exchange economy with two assets, one riskless and one risky, the latter representing foreign exchange. Each day, foreign exchange earns a payoff R , publicly observed, which is composed of a series of random increments, one for each day t :

$$R_t = \sum_{i=1}^t \Delta R_i \quad (6.1)$$

The increments ΔR are i.i.d. Normal $(0, \sigma_R^2)$ and can be interpreted as the flow of public macroeconomic information (e.g., interest rate changes). The foreign exchange market has two participant types, dealers and customers. Within each day t there are three rounds of trading:

- Round 1: Dealers trade with customers (the public).
- Round 2: Dealers trade among themselves (to share inventory risk).
- Round 3: Dealers trade again with the public (to share risk more broadly).

The feature of their model that produces the result that dealer positions return to zero is their assumption that the public has a comparative advantage in holding overnight positions (i.e. the public is 'large' in a convergence sense relative to the risk-bearing capacity of the dealers). In equilibrium, the aggregate position of dealers is fully absorbed each day by the public. Put differently, for market-clearing in their model, any net trade by the public in round 1 must be reabsorbed by the public in round 3.⁸

This particular feature has strong implications for total customer flow. For example, it implies that marketwide, customer order flow each day should net to zero. Now, the data available to us on customer order flow represent the orders received by one bank, not the customer flow received marketwide, so this prediction is untestable. Suppose, however, that the single-bank data represent a random sample (10–15 per cent) of the marketwide customer order flow on any given day. In this case, the Evans–Lyons model predicts that:

- H1:** For a single bank, customer order flow each day should differ from zero due to random sampling error only.
- H2:** For a single bank, customer order flow each day should be uncorrelated with changes in the exchange rate.

Hypothesis 2 follows from the fact that the customer-order sample is assumed here to be random. (It should therefore contain on average as many realizations of the model's round 1 'shock' orders as it does end-of-day 'absorption' orders.)

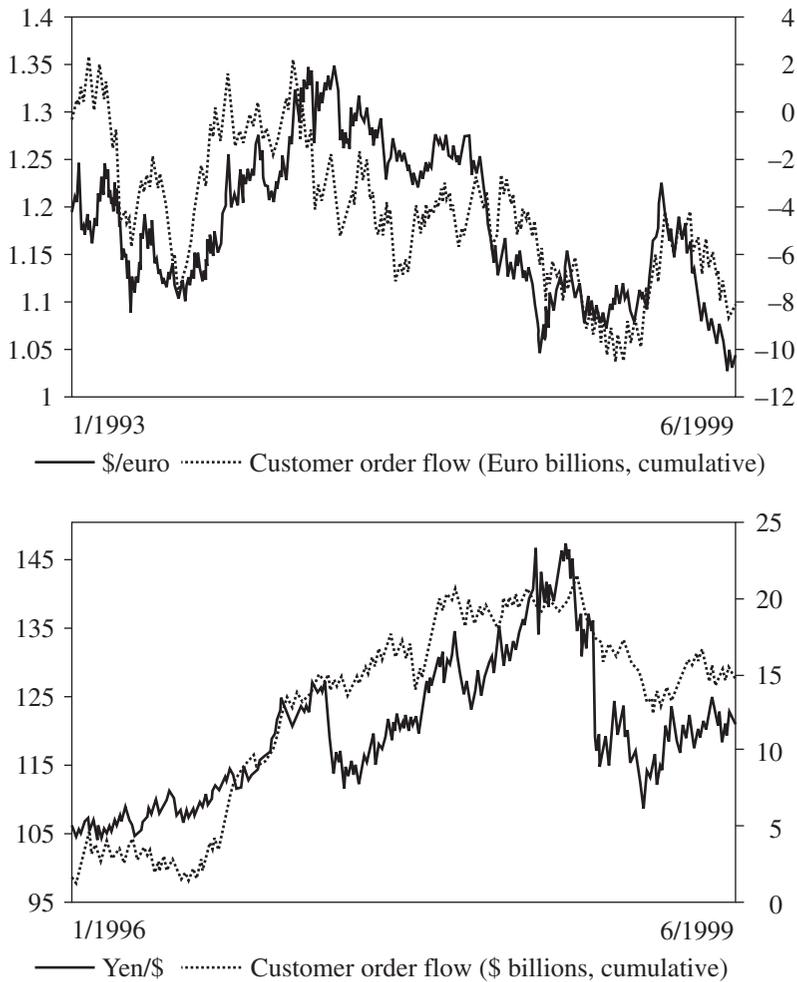
It is also possible, however, that all customer orders are not equally informative of subsequent market movements. Suppose for example that customer order flows are not alike in terms of their market impact. One might imagine two categories of customers: high-impact customers and low-

impact customers. If this were the right description of the world, then a bank's customer orders might not be representative of the customer-order population because the bank has a disproportionate amount of high-impact customers.

We turn now to plots of the customer flows, which provide a first glimpse of the possible link to exchange-rate movements. Figure 6.3 shows cumulative customer order flows and the level of the exchange rate in both the dollar-euro and dollar-yen markets.⁹ Positive correlation is evident. Comparing these plots to those for cumulative interdealer order flow in Evans and Lyons (2002), one sees that the correlation in Figure 6.3 is not as tight at higher frequencies. At lower frequencies, say monthly, the relation is manifested clearly. These plots also have implications for the two hypotheses introduced above. These hypotheses stated that this bank's daily customer flow should differ from zero due only to random sampling error, and should therefore be uncorrelated with exchange-rate movements. These hypotheses are clearly rejected: cumulative order flow received by this bank is correlated with exchange-rate movements.

What could explain this positive correlation? One possibility is that it is not really there – the correlation is not statistically significant. But the correlation between these series is statistically significant (p-value below 1 per cent). Another possibility is that marketwide customer flow each day does not net to zero. For example, collectively dealers may be maintaining non-zero positions. Though this would not be surprising from day to day, it would be quite surprising at weekly frequencies and lower (and these lower frequencies are more relevant for the correlation in Figure 6.3). Accordingly, we do not consider this a compelling reason to believe that rejection of hypothesis 1 is behind the positive correlation. Another reason why marketwide customer flow each day may not net to zero is that dealers *are* achieving their collective zero position, but do so by hedging with instruments that do not enter our sample (e.g. currency futures or options). The evidence suggests, however, that FX dealers use these methods of risk management rarely if at all (c.f. Naik and Yadav 2000 for dealer hedging in other financial markets).

There remain at least two other possibilities consistent with marketwide customer flow each day netting to zero that can explain the correlation in Figure 6.3. (We offer these as suggestions for future work; we cannot settle the issue based on analysis presented here.) First, it is possible that the customers of this bank are on average better informed. For example, the bank that is the source of the data is one of the very top FX trading banks in the world, which may attract a disproportionate share of the most informative customer business. (More concretely, suppose the orders of hedge funds are the most informative and this bank receives more than its share of



Notes: The plots show the spot exchange rate and cumulative customer order flow received by the source bank. The sample for the \$/euro plot is January 1993 to June 1999. The sample for the yen/\$ plot is January 1996 to June 1999 (the January 1993 to December 1995 period is not included due to the lack of Tokyo-office data). The spot exchange rate is expressed on the left-hand scale. The cumulative customer order flow is expressed on the right-hand scale (in billions of euros for \$/euro and in billions of dollars for yen/\$). Positive order flow in the case of the euro denotes net demand for euros (following the convention in that market of quoting prices in dollars per euro). Positive order flow in the case of the yen denotes net demand for dollars (following the convention in that market of quoting prices in yen per dollar).

Figure 6.3 Cumulative customer flow and exchange rates

hedge-fund orders.) A second possibility consistent with hypothesis 1 that can explain the correlation relies instead on this bank's sheer size. Suppose this bank's customers are the same as customers marketwide, but because the bank has such a large slice of total customer flow, its trades in the interdealer market generate disproportionate price impact. (A model along these lines could include a cost of 'monitoring' the trading activity of various banks; in this setting it may be cost-efficient to place disproportionate weight on the interdealer trades of large banks, despite their customers being no better informed than the average customer. See, e.g. Calvo 1999). Further theoretical work will undoubtedly produce additional explanations. As additional customer-flow data become available, empiricists will be able to distinguish between them.

4. CASE STUDY: THE COLLAPSE OF THE YEN/\$ RATE, OCTOBER 1998

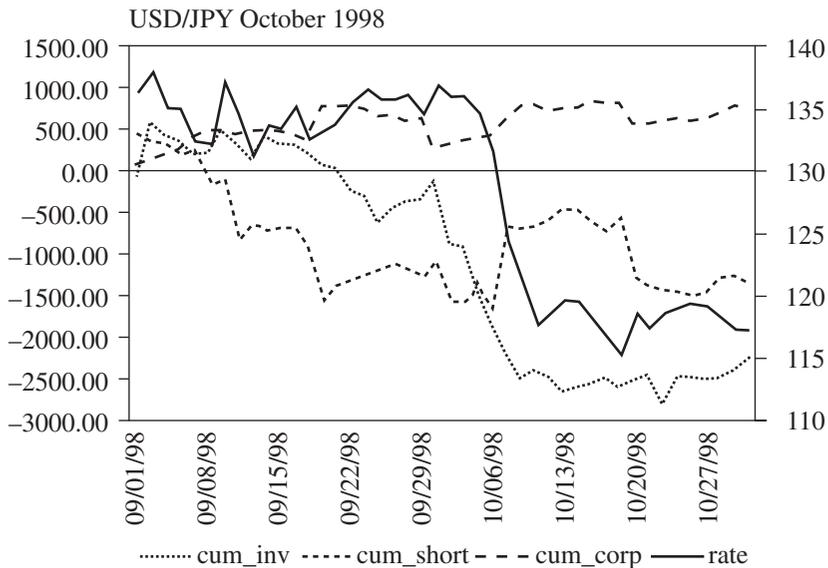
One of the most remarkable events in the post-Bretton-Woods era of floating exchange rates is undoubtedly the remarkable drop in the yen/\$ rate that occurred in October 1998. In a single day, the rate fell from about 130 to about 118, a change of roughly 10 per cent. On that day, bid-offer spreads were said to have topped one yen, i.e. one per cent or more in a market that usually trades with a spread of 1–2 basis points (and is arguably the second most liquid market in the world, behind the dollar-euro market). There was no identifiable macroeconomic news, at least not news that is usually associated with exchange-rate fundamentals. The financial news at the time was concentrated on the collapse of Long Term Capital Management (LTCM), a hedge fund whose positions around the world had become so illiquid that unwinding them became impossible without driving LTCM's capital below zero.

Major banks attribute the yen/\$ rate's drop to 'the unwinding of positions by hedge funds that had borrowed in cheap yen to finance purchases of highyielding dollar assets' – the so-called yen carry trade (*The Economist*, 10/10/98).¹⁰ This portfolio shift – and the selling of dollars that came with it – was forced by the scaling back of speculative leverage following the LTCM crisis. Though received wisdom suggests that this particular mechanism was at work, we still have little direct evidence. One paper, Cai et al. (2001), provides a first cut on the issue. They model volatility around the event using an aggregate order-flow measure and a comprehensive list of macro announcements. (Their aggregate order-flow measure is the same weekly data from the US Treasury used by Wei and Kim 1997.) They do find that there is an independent role for order flow, even after accounting

for an extensive list of public news. But their volatility model and aggregate flow measures cannot determine which players were pushing prices in which direction.

Here we adopt a case-study approach. We examine the behaviour of order flows by different players around the time of the event. Which institution types were doing the dollar selling? Identifying the sellers' types gives us insight into why the selling occurred (e.g. were they institutions that may have been 'distressed', in the sense of being compelled to sell due to institutional constraints like loss limits?).

Figure 6.4 plots the daily yen/\$ exchange rate and cumulative order flow from each of the three customer segments. The one customer segment that jumps out as having clear connection with the yen/\$ rate collapse is unleveraged financial institutions (denoted cum_inv). These institutions began strong selling before the exchange rate began to move. In total, they sold



Notes: Left axis shows cumulative flows in millions of dollars (positive is dollar purchases). Right axis is yen per dollar. Note that yen/dollar rate persisted at its new lower level following the sharp drop on 7 October. The series cum_inv denotes the cumulative order flow from unleveraged investors (e.g. mutual funds, pension funds, insurance companies, etc.). The series cum_short denotes cumulative order flow from leveraged investors (e.g., hedge funds). The series cum_corp denotes cumulative order flow from non-financial corporations.

Figure 6.4 October 1998 collapse of yen/\$ rate: cumulative flow of the three customer types

about 2.5 billion dollars over the days preceding and including the collapse. It will be interesting in the future to identify which of these unleveraged financial institutions were the most important. Were they US institutions or Japanese? Did they become distressed in some way precipitated by the LTCM collapse, or is the 'distressed players' view of the event misguided? Another common view is that this type of institution was heavily affected by common risk management systems (as opposed to actual 'distress'). There is much room for future work to address these possibilities.

It is also possible that the portfolio shift of unleveraged financial institutions' was, in the end, not uncommonly large, but it occurred at a time when leveraged financial institutions were distressed, and were therefore not willing to provide liquidity that they would normally provide. The cumulative order flow from leveraged investors in Figure 6.4 is consistent with this (denoted *cum_short*). It shows that leveraged investors did not account for any abrupt selling at the time of the collapse (contrary to received wisdom). Rather, they appeared to have *provided* liquidity, buying about 1 billion dollars at the time the price was collapsing. In the weeks prior to the big rate move, however, leveraged investors in aggregate were indeed selling dollar positions, and may have been reluctant to get back into them. Finally, the cumulative order flow from non-financial corporations shows they were also buying dollars at the time of the price collapse (denoted *cum_corp*); their total buying was small, however – only about half the size of the buying by leveraged financial institutions, or 0.4 billion dollars.

What are some of the larger implications of this yen/\$ case study? So abrupt a shift to a new exchange-rate level (without macro news) leads one to consider the possibility of path dependence. The type of path dependence we have in mind in this case is determined by the sequencing of trades by various customer types, and how that sequencing may have mattered for the ultimate exchange rate. For example, keeping the path of total customer flow the same, if it had been the unleveraged financial institutions that had gradually fled dollars in early September, rather than the leveraged institutions, might the new level of the yen/\$ rate in late October have been different? Though path dependence of this kind is not a property of existing FX microstructure models, it is an interesting possibility for future work to consider.

Another larger implication of the case study is that liquidity in FX markets varies over time, sometimes quite substantially. Existing empirical estimates of order flow's price impact imply that the price impact from the relatively small portfolio shifts in Figure 6.4 should have been much smaller, perhaps a few percent, rather than the roughly 10 per cent change that occurred.¹¹ Though timevarying liquidity is also not a property of the

early-vintage microstructure models, it is an issue that future work must consider. What triggers liquidity changes? Might what appears on the surface to be changing liquidity be due instead to changing order-flow composition?

5. OTHER EPISODES OF EXTREME PRICE MOVEMENT

Let us turn to some other examples of extreme exchange rate movements to determine whether the pattern shown above the yen/dollar rate in the fall of 1998 is representative. The rate move examined in the previous section is by far the largest daily move in our sample. Sticking with the yen/dollar market, we examine the remainder of the top five largest return days (either direction) in the yen market.

Figure 6.5, panels A-D, shows the flows in the days surrounding these other four large-return days for the two categories of financial institutions (leveraged and unleveraged). (Non-financial corporate flows are not presented in Figure 6.5 because they show no relationship to the exchange rate within these highfrequency windows.) The panels correspond to the following episodes in chronological order:

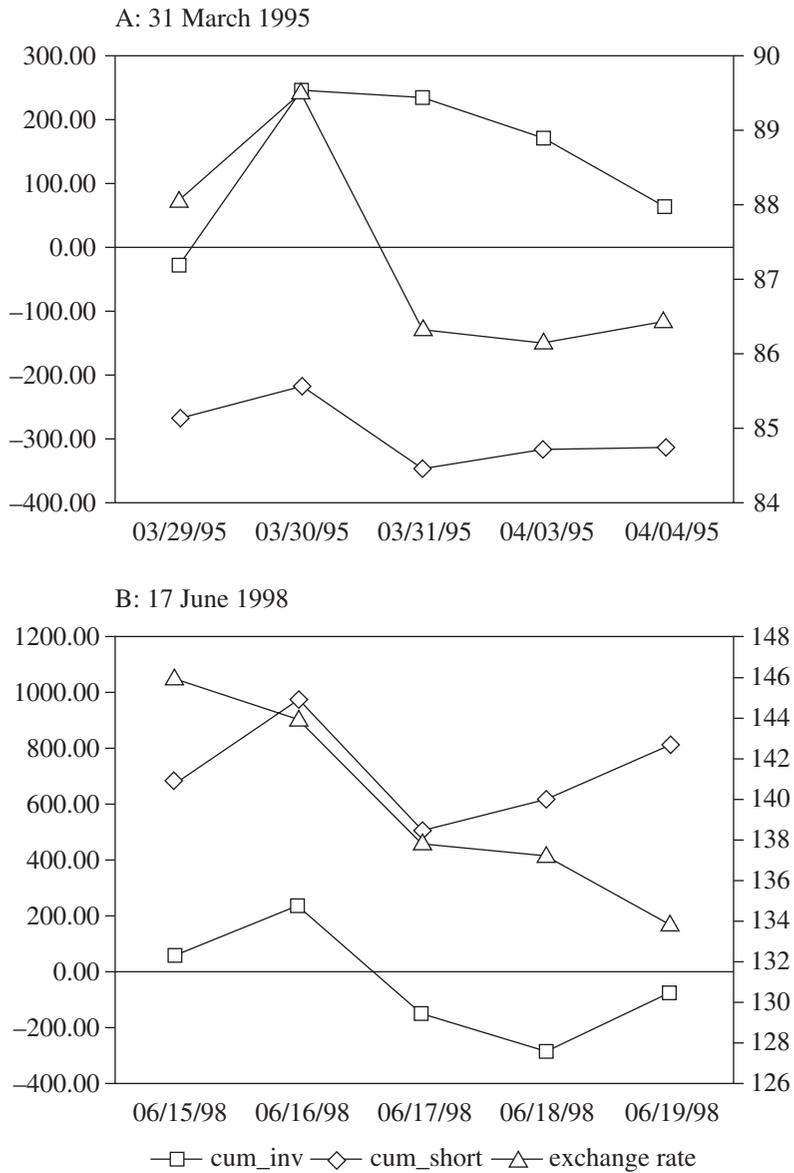
- Panel A: flows around the roughly 3 per cent drop in ¥/\$ on March 31, 1995.
- Panel B: flows around the roughly 4 per cent drop in ¥/\$ on June 17, 1998.
- Panel C: flows around the roughly 4 per cent drop in ¥/\$ on September 1, 1998.
- Panel D: flows around the roughly 4 per cent rise in ¥/\$ on September 9, 1998.

The figure shows that (high frequency) extreme events are typically associated with a large net flow from financial institutions. In three of the four cases, both of the financial institution types were in aggregate trading in the same direction as the exchange rate move on the day of the move. Though the net flow sizes are not as extraordinary as the returns on these days (given the distribution of daily net flows in the sample), they are substantial, and are likely to be positively correlated with the net flows received by other large banks (making the sizes on the left axis a downward biased measure of marketwide flow). The only one of the four days in which both categories did not go in the same direction as the exchange rate is September 1, 1998 (Panel C). In this case, the unleveraged financial institu-

tions were going in the direction of the rate change (albeit only slightly), but the leveraged financial institutions were buying dollars strongly as the price of dollars fell. This case, then, looks (qualitatively) more like the October 7, 1998, collapse of the yen/dollar rate: leveraged investors stepped in to provide liquidity, i.e. to oppose the price movement (at least this is true of the leveraged investors at this bank).

The word 'extreme' is typically associated with high frequency movements. But extreme events can also occur over time, in the form of secular trends that unexpectedly emerge. One such event is the secular depreciation of the euro since early 1995. This depreciation was particularly acute through 1999, the first official year of the euro's existence. (The impact on policy debates about EMU has been substantial.) Now, Figure 6.3 above suggested that customer flow data capture broad features of lower-frequency exchange rate trends. Is there a particular customer category whose qualitative behavior corresponds most closely to the trend depreciation of the euro? The answer is yes. Figure 6.6 shows that the trend depreciation of the euro is closely associated with net flows from non-financial corporations. (The flows of the financial institution categories are not shown because they do not evince any low-frequency association.) The slope of the cumulative non-financial corporate flow is not only negative, it is becoming increasingly negative throughout (concave down). Is this a causal relation? We have to leave that to future work to determine. The association is quite striking in any event. Further work with more disaggregated data on the trades of nonfinancial corporations would be quite valuable. Note, finally, that this focus on non-financial corporations is in stark contrast to the focus on portfolio investment flows in the broader literature as a possible explanation for the euro's weakness.

Are the customers of the bank that provided the data representative of FX customers more generally? There are two parts to this question. First, are the shares of the three customer types in this bank's flow 'portfolio' representative of marketwide average shares? Second, within a given customer category (e.g., leveraged investors), are this bank's customers representative of customers in that same category at other banks? Our conjecture is that the answer to the first question is no, while the answer to the second is a qualified yes. The reason we do not believe that the shares of this bank's customer types are representative of marketwide average is tied to our analysis of hypotheses 1 and 2. That is, one explanation for the positive correlations found in Figure 6.3 is that order flow from different customer categories may have different price impact, and this bank's customer shares may be tilted toward high price-impact categories. Regressions of returns on contemporaneous order flows from the separate categories confirm that the different types have quite different price impact (statistically and



Notes: Left axis shows cumulative flows in millions of dollars (positive is dollar purchases). Right axis is yen per dollar.

Figure 6.5 Top five big return days in US dollar–yen

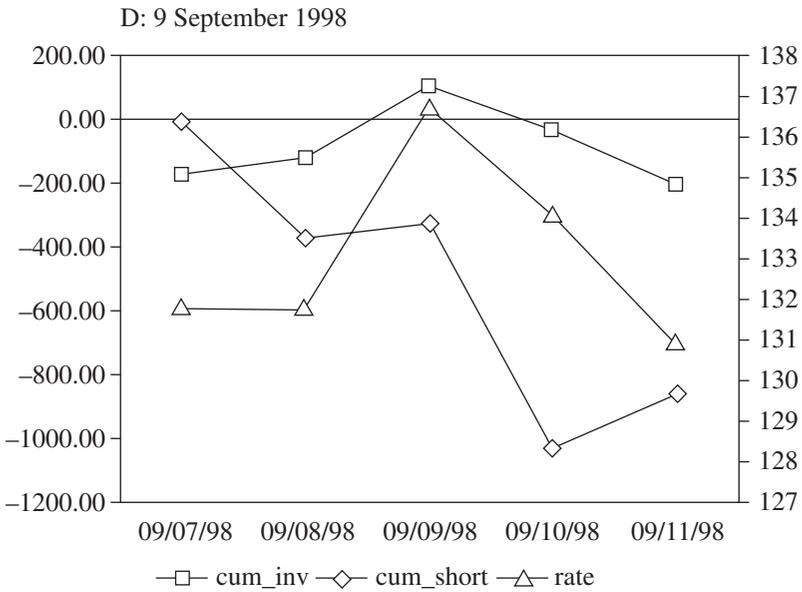
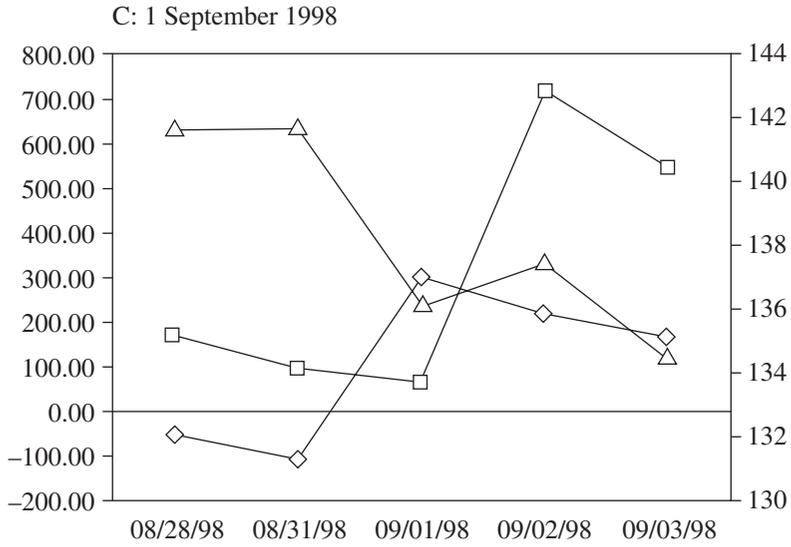
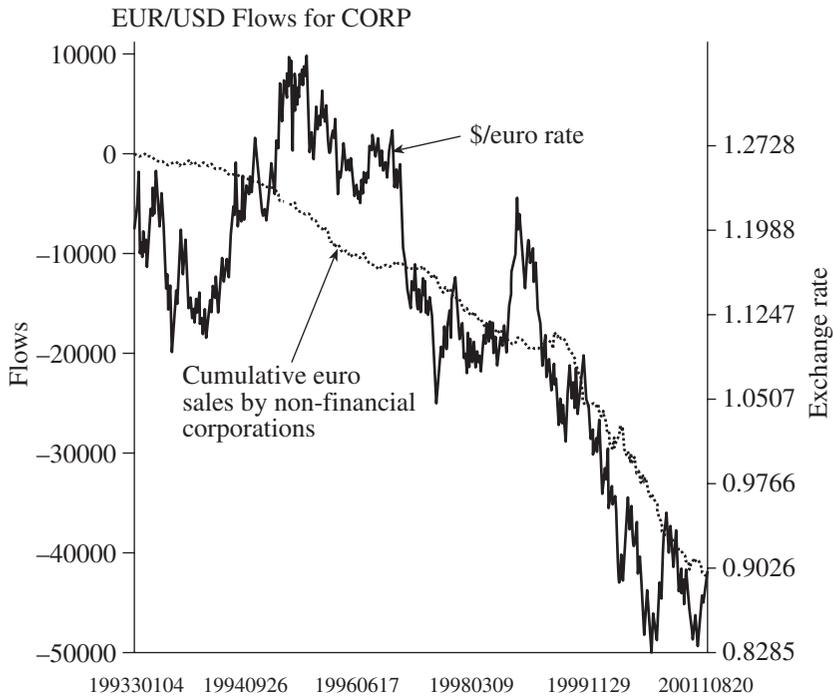


Figure 6.5 (continued)



Notes: Left axis shows cumulative flows in millions of euros (positive is euro purchases). Right axis is dollars per euro. Before 1 January, 1999 flow is constructed from flows in the constituent currencies.

Figure 6.6 Non-financial corporate flows out of the euro, 1993–2001

economically), with the lowest price impact being from the non-financial corporations. The source bank suggests that (relative to the rest of the banking system) their order flow is in fact tilted away from non-financial corporations, which is consistent with our conjecture. Finally, we do believe that this bank's customers are representative of customers in that same category at other banks, as long as 'other banks' is taken to mean other commercial banks with substantial market share in the FX market.

6. CONCLUSIONS

This paper addresses the trading of FX customers: investors, importers, exporters, corporate Treasurers, etc. Past work within the microstructure

approach has focused on FX trading between banks, or ‘interdealer’ trading. Though most trading in FX is interdealer, it is the demands of non-dealer customers that represent the underlying demand for currencies in the real economy. Data necessary for this analysis have become available only recently, in part due to the wholesale shift to electronic trading (and the data capture it permits).

Our main results include the following. First, we find that for customer order flow (10–15 per cent of the marketwide total) shows little evidence of mean reversion. Indeed, customer order flow cumulated over time is approximately a random walk. Second, customer flow tracks exchange rate movements at lower frequencies (e.g. annual) rather closely. Third, when customer order flow is disaggregated, we find that the parts behave rather differently. For example, our case study on the remarkable drop in the yen/\$ rate that occurred in October 1998 shows that hedge funds were not the trigger of the collapse, but instead were net providers of liquidity. The trigger was the portfolio shift of financial institutions like mutual funds, pension funds, and insurance companies. Fourth, we find that extreme exchange-rate movements at high-frequency are generally associated with large net flows from financial institutions; in contrast, low frequency trends are associated with secular net flows from non-financial corporations. We consider the largely graphical analysis of this paper as but a first step in a new arena for empirical work within FX microstructure, a research area that Professor Goodhart set in motion (see Goodhart 1988, Goodhart 1989, Goodhart and Figliuoli 1991).

NOTES

* We thank the following for helpful comments: Richard Adams and Andrew Rose. Lyons thanks the National Science Foundation for financial assistance.

1. Order flow is not synonymous with trading volume. Order flow – a concept from microstructure finance – refers to *signed* volume. Trades can be signed in microstructure models depending on whether the ‘aggressor’ is buying or selling. (The dealer posting the quote is the passive side of the trade.) For example, a sale of 10 units by a trader acting on a dealer’s quotes is order flow of -10 .

In rational-expectations (RE) models of trading, order flow is undefined because all transactions in that setting are symmetric. One might conclude from RE models that one could never usefully distinguish the ‘sign’ of a trade between two willing counterparties. A large empirical literature in microstructure finance suggests otherwise (Lyons 2001).

2. By analogy, it is not unreasonable to view firms that trade on the NYSE as fundamentally the same as firms that trade on the NASDAQ, other things equal (i.e. similar cost of capital, similar relative valuation, etc.).
3. For practitioner-oriented research on flow effects on exchange rates, see, e.g. Citibank’s *Citiflows Global Flow and Volume Analysis* (various issues), Deutschebank’s *Flowmetrics Monthly* (various issues), and Lehman Brothers’ *Global Economic Research Series*, particularly the issue on ‘FX Impact of Cross-Border M&A’. For evidence from practitioner

surveys, see Gehrig and Menkhoff (2000). It is noteworthy that – unlike fundamental and technical analysis – order-flow analysis is not available to everyone: one needs sufficient order-flow data.

4. Osler (2001) also obtains data on FX customer trades directly from a bank. Her focus is stop-loss and take-profit orders. She shows that clustering of these orders at particular prices helps to explain two familiar predictions from technical analysis, namely that (1) trends tend to be reversed at support and resistance levels and (2) trends tend to gain momentum if support and resistance levels are breached.
5. A natural question is where the trades of central banks appear. The source of these data is reluctant to disclose the specifics. Though not reported in the table, the source bank does maintain a fourth category of customer called ‘miscellaneous’. Though the trades within this category are quite small relative to the trades in the three main categories, the category is likely to include any central bank trades for which the source bank was the counterparty. (Central bank trades tend to be small relative to private trades.)
6. One could argue that in frictionless general equilibrium, starting from pareto-optimal allocations, it is not clear why firms’ ex-ante ‘portfolios’ are not instantaneously restored. As an empirical matter, this objection to the example is not so compelling: for many institution types, there are substantial (labour intensive) costs of adjusting their net positions in the market. This may produce path dependence in portfolio allocation, even if the relation between realized order flow and price is unique.
7. A direction for future research in this area is to isolate categories of trades that fit neatly into a particular balance-of-payments category. For example, one could isolate equity mutual funds. In this case, one could be confident that their FX trades fit neatly into the category called international portfolio investment.
8. By public we mean non-dealers, which includes the proprietary trading desks at dealing banks (a type of leveraged investor, as already noted).
9. The yen plot begins in January 1996 because Citibank did not include customer-flow data from its Tokyo office in its database until late 1995. (The Tokyo office is especially important for Citibank’s dollar-yen customer flow.) Note that this may account for the seemingly small share of non-financial corporate trading in total customer trading in dollar-yen shown in Table 1: if non-financial corporations tend to trade via their regional office, whereas financial institutions tend to trade on a 24-hour basis worldwide, then the customer trades in the database before 1996 would be tilted toward the financial institutions.
10. A comprehensive description of events is provided in BIS (1999).
11. For example, Evans and Lyons (2002) find that net order flow of \$1 billion produces a lasting price change of about 0.5 per cent. See also Furfine and Remolona (2001) for a transaction-level analysis of time-varying price impact in the US treasury bond market over the same period.

REFERENCES

- Bank for International Settlements (1999), *A Review of Financial Market Events in Autumn 1998*, Publication of the Committee on the Global Financial System, October.
- Bank for International Settlements (2001), *BIS 71st Annual Report*, June (at www.bis.org).
- Cai, J., Y. Cheung, R. Lee, and M. Melvin (2001), ‘“Once in a generation” yen volatility in 1998: fundamentals, intervention, and order flow’, *Journal of International Money and Finance*, 20: 327–47.
- Calvo, G. (1999), ‘Contagion in emerging markets: when Wall Street is a carrier’, University of Maryland working paper.

- Evans, M. (2001), 'FX trading and exchange rate dynamics', NBER Working Paper 8116, February, *Journal of Finance*, (DATA TO COME).
- Evans, M. and R. Lyons (2002), 'Order flow and exchange-rate dynamics', *Journal of Political Economy*, 110: 170–80.
- Evans, M. and R. Lyons (2001), 'Portfolio balance, price impact, and secret intervention', NBER Working Paper 8356, July.
- Fujiwara, I. (2000), 'Liquidity and leverage risk in the dollar/yen market', typescript, Oxford: Nuffield College, June.
- Gehrig, T. and L. Menkhoff (2000), 'The use of flow analysis in foreign exchange: exploratory evidence', typescript, University of Freiburg (Germany), Department of Economics, August, *Journal of International Money and Finance*, forthcoming.
- Glosten, L. and Milgrom, P. (1985), 'Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents', *Journal of Financial Economics*, 14: 71–100.
- Goodhart, C. (1988), 'The foreign exchange market: a random walk with a dragging anchor', *Economica*, 55: 437–60.
- Goodhart, C. (1989), 'News' and the Foreign Exchange Market, Manchester Statistical Society publication, October.
- Goodhart, C. and L. Figliuoli (1991), 'Every minute counts in financial markets', *Journal of International Money and Finance*, 10: 23–52.
- Goodhart, C., R. Ito, and R. Payne (1996), 'One day in June 1993: a study of the working of the Reuters 2000–2 electronic foreign exchange trading system', in *The Microstructure of Foreign Exchange Markets*, eds J. Frankel, G. Galli and A. Giovannini, Chicago: University of Chicago Press, IL, pp. 107–79.
- Goodhart, C. and M. O'Hara (1997), 'High frequency data in financial markets: issues and applications', *Journal of Empirical Finance*, 4: 73–114
- Goodhart, C., and R. Payne (1996), 'Micro-structural dynamics in a foreign exchange electronic broking system', *Journal of International Money and Finance*, December.
- Killeen, W., R. Lyons, and M. Moore (2001), 'Fixed versus flexible: lessons from EMS order flow', NBER Working Paper 8491, September.
- Kyle, A. (1985), 'Continuous auctions and insider trading', *Econometrica*, 53: 1315–35.
- Lyons, R. (1995), 'Tests of microstructural hypotheses in the foreign exchange market', *Journal of Financial Economics*, 39: 321–51.
- Lyons, R. (2001), *The Microstructure Approach to Exchange Rates*, Cambridge, Mass.: MIT Press (chapters at haas.berkeley.edu/~lyons).
- Morris, S. and H. Shin (1999), 'Risk management with interdependent choice', *Financial Stability Review*, November, 141–50.
- Naik, N. and P. Yadav (2000), 'Do market intermediaries hedge their risk exposure with derivatives?', typescript, London Business School, February.
- Osler, C. (2001), 'Currency orders and exchange rate dynamics: an explanation for the predictive success of technical analysis', typescript, Federal Reserve Bank of New York, August, *Journal of Finance*, forthcoming.
- Payne, R. (1999), 'Informed trade in spot foreign exchange markets: an empirical investigation', typescript, London School of Economics, January.
- Wei, S. and J. Kim (1997), 'The big players in the foreign exchange market: do they trade on information or noise?', NBER Working Paper 6256, November.