

Foreign Exchange Market Microstructure

Martin D. D. Evans¹

Georgetown University and NBER

Abstract

This paper provides an overview of the recent literature on Foreign Exchange Market Microstructure. Its aim is not to survey the literature, but rather to provide an introductory tour to the main theoretical ideas and empirical results. The central theoretical idea is that trading is an integral part of the process through which information relevant to the pricing of foreign currency becomes embedded in spot rates. Micro-based models study this information aggregation process and produce a rich set of empirical predictions that find strong support in the data. In particular, micro-based models can account for a large proportion of the daily variation in spot rates. They also supply a rationale for the apparent disconnect between spot rates and fundamentals. In terms of forecasting, micro-based models provide out-of-sample forecasting power for spot rates that is an order of magnitude above that usually found in exchange-rate models.

Keywords: Exchange Rates, Microstructure, Information Aggregation, FX Trading.
JEL No. F3, F4, G1

¹ Department of Economics, Georgetown University, Washington DC 20057, Tel: (202) 687-1570, Email: evansm1@georgetown.edu. This paper was prepared for the New Palgrave Dictionary of Economics. I thank Richard Lyons for valuable discussions and gratefully acknowledge the financial support of the National Science Foundation.

Introduction

Models of foreign exchange (FX) market microstructure examine the determination and behavior of spot exchange rates in an environment that replicates the key features of trading in the FX market. Traditional macro exchange rate models pay little attention to how trading in the FX market actually takes place. The implicit assumption is that the details of trading (i.e., who quotes currency prices and how trade takes place) are unimportant for the behavior of exchange rates over months, quarters or longer. Micro-based models, by contrast, examine how information relevant to the pricing of foreign currency becomes reflected in the spot exchange rate via the trading process. According to this view, trading is not an ancillary market activity that can be ignored when considering exchange rate behavior. Rather, trading is an integral part of the process through which spot rates are determined and evolve. Recent micro-based FX models also differ from other areas of microstructure research in their focus on the links between trading, asset price dynamics, and the macroeconomy.

Recent research on exchange rates stresses the role of heterogeneity (e.g., Bacchetta and van Wincoop 2003, and Hau and Rey 2002). Micro-based exchange-rate models start from the premise that much of the information about the current and future state of the economy is dispersed across agents (i.e., individuals, firms, and financial institutions). Agents use this information in making their every-day decisions, including decisions to trade in the FX market at the prices quoted by dealers. Dealers quote prices (e.g. dollars per unit of foreign currency) at which they stand ready to buy or sell foreign currency; they will purchase foreign currency at their bid quote, and sell foreign currency at their ask quote. Agents that choose to trade with an individual dealer are termed the dealer's customers. The difference between the value of purchase and sale orders *initiated* by customers during any trading period is termed customer order flow. Importantly, order flow is different from trading volume because it conveys information. Positive (negative) order flow indicates to a dealer that, on balance, their customers value foreign currency more (less) than his asking (bid) price. By tracking who initiates each trade, order flow

provides a measure of the information exchanged between counterparties in a series of financial transactions.

Trading in the FX market also takes place between dealers. In direct interdealer trading, one dealer asks another for a bid and ask quote, and then decides whether he wishes to trade. When the dealer initiating the trade purchases (sells) foreign currency, the trade generates a positive (negative) interdealer order flow equal to the value of the purchase (sale). Interdealer trading can also take place indirectly via brokerages that act as intermediaries between two or more dealers. In recent years electronic brokerages have come to dominate interdealer trading, but the interdealer order flow generated by brokered trades plays the same informational role as the order flow associated with direct interdealer trading.

Micro-Based Exchange Rate Determination

At first sight, the pattern of FX trading activity seems far too complex to provide any useful insight into the behavior of exchange rates. However, on closer examination, two key features emerge: First, the equilibrium spot exchange rate does not come out of a “black box”. Instead, it is solely a function of the foreign currency prices quoted by dealers at a point in time. This is a distinguishing feature of micro-based exchange rate models and has far-reaching implications. Second, information about the current and future state of the economy will only impact on exchange rates when, and if, it affects dealer quotes. Dealers may revise their quotes in response to new public information that arrives via macroeconomic announcements. They may also revise their quotes based on orders they receive from customers and other dealers. This order flow channel is the means through which dispersed information concerning the economy affects dealer quotes and hence the spot exchange rate. The role played by order flow in transmitting information to dealers, and hence to their quotes, is another distinguishing feature of micro-based exchange rate models.

Micro-based models incorporate these two features of FX trading into a simplified setting. Canonical multi-dealer models, such as Lyons (1995) and Evans

and Lyons (2002a), posit a simple sequence of quoting and trading. At the start of each period, dealers quote FX prices to customers. These prices are assumed to be good for any amount and are publicly observed. Each dealer then receives orders from a subset of agents, his customers. Dealers next quote prices in the interdealer market. These prices, too, are good for any quantity and are publicly observed. Dealers then have the opportunity to trade among themselves. Interdealer trading is simultaneous and trading with multiple partners is feasible.

In this trading environment, optimal quote decisions take a simple form; all dealers quote the same FX price to both customers and other dealers. We can represent the period- t quote as

$$s_t = (1-b) \sum_{i=0}^{\infty} b^i E[f_{t+i} | \Omega_t^D], \quad (1)$$

where $0 < b < 1$. s_t is the log price of foreign currency quoted by all dealers, and f_t denotes exchange rate fundamentals. The form for fundamentals differs according to the macroeconomic structure of the model. For example, in Evans and Lyons (2004b), f_t includes home and foreign money supplies and household consumption. In models where central banks conduct monetary policy via the control of short-term interest rates (i.e., follow Taylor-rules), f_t will include variables used to set policy. More generally, f_t will include a term that identifies the foreign exchange risk premium.

While equation (1) takes the present value form familiar from standard international macro models, here it represents how dealers quote the price for foreign currency in equilibrium. All dealers choose to quote the same price in this trading environment because doing otherwise opens them up to arbitrage, a costly proposition. (Recall that quotes are publicly observed and good for any amount, so any discrepancy between quotes would represent an opportunity for a riskless trading profit.) Consequently, the month- t quote must be a function of information known to all dealers. Equation (1) incorporates this requirement with the use of the expectations operator, $E[. | \Omega_t^D]$, that denotes expectations conditioned on

information common to all dealers at the start of month t , Ω_t^D . This is not to say that all dealers have the same information. On the contrary, the customer order flows received by individual dealers represent an important source of private information so there may be a good deal information heterogeneity across dealers at any one time. The important point to note from equation (1) is that due to the “fear of arbitrage”, individual dealers choose not to quote prices based on their own private information. In this trading environment, dealers use their private information in initiating trade with other dealers, and, in so doing, contribute to the process through which all dealers acquire information.

The implications of micro-based models for the dynamics of spot rates are most easily seen by rewriting (1) as

$$\Delta s_{t+1} = \frac{1-b}{b}(s_t - E[f_t | \Omega_t^D]) + \varepsilon_{t+1}, \quad (2)$$

where $\Delta s_{t+1} = s_{t+1} - s_t$, and

$$\varepsilon_{t+1} = \frac{1-b}{b} \sum_{i=1}^{\infty} b^i (E[f_{t+i} | \Omega_{t+1}^D] - E[f_{t+i} | \Omega_t^D]). \quad (3)$$

Equation (2) decomposes the change in the log spot rate (i.e., the depreciation rate for the home currency) into two components: the expected change $E[\Delta s_{t+1} | \Omega_t^D]$ identified by the first term, and the unexpected change, $\varepsilon_{t+1} = s_{t+1} - E[s_{t+1} | \Omega_t^D]$, shown in equation (3). Both terms contribute to exchange rate dynamics in micro-based models. In equilibrium, dealers’ period- t quote must be based on expectations, $E[\Delta s_{t+1} | \Omega_t^D]$, that match the risk-adjusted returns on different assets. This means that variations in the interest differential between home and foreign bonds can contribute to the volatility of the depreciation rate via the first term in (2). The second term, ε_{t+1} , identifies the impact of new information received by all dealers between the start of periods t and $t+1$. Equation (3) shows that new information impacts on the FX price quoted in period $t+1$ to the extent it revises

forecasts of the present value of fundamentals based on dealers' common information.

As an empirical matter, depreciation rates are very hard to forecast, so the dynamics of spot rates are largely attributable to the effects of news. Here micro-based models have a big advantage over their traditional counterparts because their trade-based foundations provide detail on how news affects spot rates. In particular, as equation (3) indicates, micro-based models focus on how new information about the fundamentals reaches dealers and induces them to revise their FX quotes.

News concerning fundamentals can reach dealers either directly or indirectly. Common knowledge (CK) news operates via the direct channel. CK news contains unambiguous information about current and/or future fundamentals that is simultaneously observed by all dealers and immediately incorporated into the FX price they quote. In principle, macroeconomic announcements (e.g. on GDP, industrial production or unemployment) could be a source for CK news, but in practice they rarely contain much unambiguous new information. In fact, CK news events appear rather rare. The indirect channel operates via order flow and conveys dispersed information about fundamentals to dealers. Dispersed information comprises micro-level information on economic activity that is correlated with fundamentals. Examples include the sales and orders for the products of individual firms, market research on consumer spending, and private research on the economy conducted by financial institutions. Dispersed information first reaches the FX market via the customer order flows received by individual dealers. These order flows have no immediate impact on dealer quotes because they represent private information to the recipient dealer. The information in each customer flow will only impact on quotes once it is known to all dealers. Interdealer order flow is central to this process. Individual dealers use their private information to trade in the interdealer market. In so doing, information on their customer orders is aggregated and spread across the market. This process is known as information aggregation. Dispersed information is incorporated into dealer quotes once this process is complete.

Empirical Evidence

The appeal of micro-based models is not solely based on their theoretical foundations. In marked contrast with traditional exchange-rate models, micro-based models have enjoyed a good deal of empirical success. Evans and Lyons (2002a) first demonstrated their empirical power when studying the relation between depreciation rates and interdealer order flow at the daily frequency. In particular, they show that aggregate interdealer order flow from trading in the spot dollar/dmark market on day d accounts for 64 percent of the variation in the depreciation rate, Δs_{d+1} , between the start of days d and $d+1$. This is a striking result because macro models can account for less than 1 percent of daily depreciation rates. It is also readily explained in terms of equations (2) and (3). Aggregate interdealer order flow during day d trading provides a measure of the market-wide information flow that dealers use to revise their quotes between the start of days d and $d+1$. This contemporaneous relationship between depreciation rates and interdealer order flows appears robust. It holds for many different currencies, and for different currency-order flow combinations (e.g., Evans and Lyons 2002b, Payne 2003 and Froot and Ramadorai 2005). It is also worth emphasizing that order flow's impact on spot rates is very persistent. There is very little serial correlation in the daily depreciation rates for major currencies, so the order flow impact on current FX quotes persists far into the future.

While consistent with the idea that dispersed information is impounded into spot exchange rates via interdealer order flow, these results do not provide direct evidence on the ultimate source of exchange rate dynamics. According to micro-based models, the analysis of customer order flows should provide the evidence. In particular, if interdealer order flows measure the market-wide information flow that carries the information concerning fundamentals originally motivating customer orders, customer orders should also have explanatory power for depreciation rates. This is indeed the case. Evans and Lyons (2004b) show that a significant contemporaneous relationship exists between depreciation rates and the customer order flows of a single large bank. Moreover, the strength of this relationship

increases as we move from a one day to a one month horizon. This, too, is consistent with micro-based models: At longer horizons, customer flows from a single bank should be a better proxy for the market-wide flow of information driving spot rates.

Micro-based models also make strong empirical predictions about the relationship between order flows and fundamentals. According to equation (1), dealers are forward-looking when quoting FX prices, so spot rates embody their forecasts for fundamentals based on common information, Ω_t^D . One empirical implication of this observation is that spot exchange rates should have forecasting power for fundamentals. While there is some evidence that this is true for variables that comprise fundamentals in many models (Engel and West 2005), the forecasting power is rather limited. Micro-based models also have implications for the forecasting power of order flows: If order flows convey information about fundamentals that is not yet common knowledge to all dealers (i.e., not in Ω_t^D), then they should have incremental forecasting power for fundamentals, beyond the forecasting ability any variable in Ω_t^D . This is a strong prediction: it says that order flow should add to the forecasting power of all other variables in Ω_t^D , including the history of spot rates and the fundamental variable itself. Nevertheless, Evans and Lyons (2004b) find ample support for this prediction using customer order flows and candidate fundamental variables such as output, inflation and money supplies. These findings provide direct evidence on the information content of customer order flows, and provide a new perspective on the link between exchange rates and fundamentals.

Dispersed information concerning fundamentals need not only come from the activities of individuals, firms and financial institutions. Scheduled announcements on macroeconomic variables (e.g. GDP, inflation, or unemployment) can also be a source of dispersed information. If agents have different views about the mapping from the announced variable to fundamentals, then the news contained in any announcement, while simultaneously observed, will not be common knowledge. For example, two firms may interpret the same announcement on last quarter's GDP as having different implications for future GDP growth. Differing interpretations about

the implications of commonly observed news will be a source of customer order flows because they imply heterogeneous views about future returns, which in turn, induces portfolio adjustment. Thus, micro-based models raise the possibility that the exchange rate effects of macro announcements operate via both a direct channel (i.e., when the announcement contains CK news) and an indirect channel. Love and Payne (2002) and Evans and Lyons (2003, 2005b) find evidence that both channels are operable. Evans and Lyons estimate that roughly two-thirds of the effect of a macro announcement is transmitted indirectly to the dollar/mark spot rate via order flow, and one-third directly into quotes. With both channels operating, macro news is estimated to account for more than one-third of the variance in daily depreciation rates. This level of explanatory power far surpasses that found in earlier research analyzing the impact of macro news on exchange rates (e.g., Andersen et al. 2003). It also further cements the link between spot rates and the macro variables comprising fundamentals.

Order Flows, Returns and the Pace of Information Aggregation

The process by which the information contained in the customer flows becomes known across the market, and hence embedded into FX quotes, is complex. The individual customer and interdealer orders received by each dealer contain some dispersed information about the economy, but extracting the information from each order constitutes a difficult inference problem. Under some circumstances, the inference problems are sufficiently simple for every dealer to learn all there is to know about fundamentals in a few rounds of interdealer trading. In this case, the pace of information aggregation is very fast, so that new information concerning fundamentals is quickly reflected in dealer quotes whether the news is initially dispersed or common knowledge. The resulting dynamics for exchange rates over weeks, months or quarters will be indistinguishable from the predictions of macro models. Under other circumstances, the inference problem facing individual dealers is sufficiently complex to slow down the pace of information aggregation. Here it

takes many rounds of interdealer trading before the dispersed information concerning fundamentals becomes known across the market. This scenario is much more likely from a theoretical perspective. Evans and Lyons (2004a) show that the conditions needed for fast information aggregation are quite stringent. Of course, because interdealer trading takes places continuously, dispersed information could be completely embedded in FX quotes in a short period of calendar time (e.g., a day), even if the pace of information aggregation is slow. In principle, dealers might be able to learn a good deal from the multitude of orders they receive in a typical day, even if individual orders are relatively uninformative. The question of whether it takes significant amounts of calendar time before dispersed information is embedded in FX quotes can only be answered empirically.

If the pace of information aggregation is slow, customer order flows across the market contain information that will only become known to all dealers at a later date. So, if the customer orders received by an individual bank are representative of the market-wide flows, they should have forecasting power for the future market-wide flow of information that drives quote revision. Recent empirical findings support this possibility. Evans and Lyons (2004b, and 2005) show that customer order flows have significant forecasting power for future depreciation rates both in and out of sample. These results are qualitatively different from the contemporaneous empirical link between order flows and depreciations rates discussed above. In the context of equations (2) and (3), the market-wide flow of information from period- t trading impacts on the depreciation rate, Δs_{t+1} , via ε_{t+1} . The contemporaneous link arises because period- t interdealer order flows measure the market-wide information flow, ε_{t+1} . In contrast, the forecasting power of customer flows for the depreciation rate arises because ε_{t+1} contains information that was originally in the customer orders received by individual banks *before* period- t trading.

These forecasting results are surprising both in terms of their horizon and strength. In particular, out-of-sample forecasts based on customer flows from month $t-1$ can account for roughly 16 percent of the variation in next month's depreciation

rate, Δs_{t+1} . This finding suggests that the pace of information aggregation is far, far slower than was previously thought; it seems to take weeks, not minutes, for dispersed information to be fully assimilated across the market. The level of forecasting power is also an order of magnitude above that usually found in exchange rate models. For example, the in-sample forecasting power of interest differentials for monthly depreciation rates is only in the 2 – 4 percent range.

The slow pace of information aggregation may shed light on one of the long-standing puzzles in exchange rate economics; the disconnect between spot exchange rates and fundamentals over short and medium horizons (Meese and Rogoff 1983). The idea is quite simple. If changes in fundamentals are only reflected in spot rates once information concerning the change is recognized by dealers across the market, the slow pace of information aggregation will mask the link between the depreciation rate and the change in fundamentals over short horizons, because the latter is a poor proxy for the market-wide flow of information. Simulations in Evans and Lyons (2004a) show that this masking effect can be quite substantial. Fundamentals account for only 50 percent of variation in spot rates at the two-year horizon even though information aggregation takes at most 4 months.

One factor that might contribute to the slow pace of information aggregation is the presence of price-contingent order flow generated by feedback trading. Stop-loss orders, for example, represent a form of positive feedback trading, in which a fall in the FX price triggers negative order flow from customers wishing to insure their portfolios against further losses. Feedback trading of a known form does not complicate the inference problem facing dealers because the orders it generates are simply a function of old market-wide information. However, when the exact form of the feedback is unknown, it makes inferences less precise and so slows down the pace of information aggregation. Osler (2005) argues that feedback trading will be an important component of order flow when quotes approach the points at which stop-loss orders cluster. A fall in FX quotes at these points can trigger a self-reinforcing price-cascade where causation runs from quotes to order flow.

Some economists argued that the early empirical findings linking order flow and the depreciation rate reflected the presence of positive feedback trading rather than the transmission of dispersed information. Indeed, there is no way to tell whether intraday causation runs from order flows to quotes or vice versa from just the contemporaneous correlation between order flow and the depreciation rate measured in daily data. However, the new evidence on the forecasting power of order flow for both depreciation rates and fundamentals firmly points to order flow as the conveyor of dispersed information. This is not to say that feedback trading is absent. Portfolio insurance and other price-contingent trading strategies (e.g., liquidity provision) undoubtedly contribute to order flows and their presence may actually explain why the pace of information aggregation is so slow.

Future Research

Exchange rate research using micro-based models is still in its infancy. The past few years have seen a rapid advance in theoretical modeling and some surprising empirical results. Advances on the empirical side will be spurred by the greater availability of trading data. On the theoretical side, micro-based modeling may provide new insights into the determinants of the foreign exchange risk premium, the efficacy of foreign exchange intervention, and the anatomy of financial contagion.

References

- Andersen, T., T. Bollerslev, F. Diebold, and C. Vega (2003), Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review*, 93: 38-62.
- Bacchetta, P., and E. van Wincoop (2003), Can information dispersion explain the exchange rate disconnect puzzle? NBER Working Paper 9498, February, forthcoming, *American Economic Review*.
- Engel, C., and K. West (2005), Exchange rates and fundamentals, *Journal of Political Economy*, 113: 485-517.
- Evans, M., and R. Lyons (2002a), Order flow and exchange rate dynamics, *Journal of Political Economy*, 110: 170-180.
- Evans, M., and R. Lyons (2002b), Informational integration and FX trading, *Journal of International Money and Finance*, 21: 807-831.
- Evans, M., and R. Lyons (2003), How is Macro News Transmitted to Exchange Rates? NBER Working Paper 9433, January.
- Evans, M., and R. Lyons (2004a), A New Micro Model of Exchange Rates, NBER Working Paper 10379, March.
- Evans, M., and R. Lyons (2004b), Exchange Rate Fundamentals and Order Flow, typescript, Georgetown University, available at <http://www.georgetown.edu/faculty/evansm1/>.
- Evans, M., and R. Lyons (2005a), Do Currency Markets Absorb News Quickly? *Journal of International Money and Finance*, 24: 197-217.
- Evans, M., and R. Lyons (2005b), Meese-Rogoff Redux: Micro-Based Exchange Rate Forecasting, *American Economic Review P&P*, May.
- Froot, K., and T. Ramadorai (2005), Currency Returns, Intrinsic Value, and Institutional-Investor Flows, *Journal of Finance*, LX: 1535-1565

- Hau, H., and H. Rey (2002), Exchange rates, equity prices, and capital flows, NBER Working Paper 9398, December, *Review of Financial Studies*, forthcoming.
- Love, R., and R. Payne (2002), Macroeconomic news, order flows, and exchange rates, typescript, London School of Economics, December.
- Lyons, R. (1997), A Simultaneous Trade Model of the Foreign Exchange Hot Potato, *Journal of International Economics*, 42: 275-298.
- Meese, R., and K. Rogoff (1983), Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.
- Osler, C. (2005), Stop-loss orders and price cascades in currency markets, *Journal of international Money and Finance*, 24: 219-241.
- Payne, R. (2003), Informed trade in spot foreign exchange markets: an empirical investigation, *Journal of International Economics*, 61: 307-329.