

Under the Microscope: The Structure of the Foreign Exchange Market*

Michael J. Sager ^a and Mark P. Taylor^{b†}

^aPutnam Investments and University of Warwick.

^bUniversity of Warwick and Centre for Economic Policy Research.

April 2005

Abstract

We provide a detailed, up-to-date description of the microstructure of the foreign exchange market and of the behaviour of participant groups. In the light of this, we highlight shortcomings in existing theoretical models of market interaction and present an outline alternative model that marries theoretical prediction and current market practice.

Keywords: Foreign exchange market; microstructure; order flow.

JEL Classification: F31

*While remaining solely responsible for any remaining errors in this paper, the authors are grateful to two anonymous referees and to Suzanne Brant, Jack Crawford, Rich Lyons, Michael Melvin, Dirk Morris, Carol Osler, as well as numerous market participants, for insightful and often extensive comments on earlier drafts.

†Corresponding author: Professor Mark P. Taylor, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. E-mail: mark.taylor@warwick.ac.uk

1 Introduction

Following the breakdown of the Bretton Woods system in the early 1970s, the foreign exchange market has been the subject of intense academic and practitioner research. In broad terms, two particularly important stylised facts have emerged from this body of evidence.

First, although the foreign exchange market is highly liquid,¹ it appears to depart significantly from the text-book efficient markets paradigm in the sense that there are significant deviations from uncovered interest rate parity that cannot be accounted for by stable models of risk premia (Taylor, 1995; Sarno and Taylor, 2002) and also in the sense that it is an opaque market, since the lack of a physical market place makes the process of price-information interaction difficult to understand (Dominguez, 1999; Lyons, 2002b). Both of these traits result from heterogeneities, including informational asymmetries, differing reaction speeds to information innovations, and diverse opportunity sets and risk-return expectations, that exist amongst the various participant groups. These heterogeneities imply the presence of persistent profit opportunities from informed trading relative to an underlying benchmark return.

Second, although empirical evidence over the post-Bretton Woods period suggests that fairly standard macroeconomic fundamentals—such as relative monetary velocity—may influence the long-run behaviour of real and nominal exchange rates (for surveys see Frankel and Rose, 1995; Froot and Rogoff, 1995; Taylor, 1995; Sarno and Taylor, 2002), the quality of shorter-term exchange rate models and forecasts continues to be an occupational hazard of the international financial economist, as fundamental variables are poorly correlated with high frequency exchange rate movements.

Largely motivated by these stylised facts, a growing literature on foreign exchange market microstructure has emerged in recent years to suggest that the quality of fundamental-based exchange rate forecasts can be improved by resort to measures of foreign exchange order flow (Froot and Ramadorai, 2001; Lyons, 2002a; Evans and Lyons, 2002a,b).² An important conclusion of this literature is that order flow allows the wider market to learn about the private information

¹Consistent with Kyle (1985), we define a liquid market as one that exhibits the following characteristics: tightness, so that bid-ask spreads for small transactions are tight; depth, so that bid-ask spreads for large transactions are small and; resilience, so that deviations of the spot rate from “fair value” should be corrected quickly. See Danielsson and Payne (2001) for a recent assessment of these criteria applied to the interdealer sector of the foreign exchange market.

²Order flow may be defined as transaction volume signed according to the initiator of the trade (Lyons, 2001); positive for a buy order, negative for a sell order. It is distinct from transaction volume, since the latter is a measure of trading activity between customers and dealers, or within the interdealer market, over a given period and in a particular exchange rate without indication of the *direction* of these transactions. Measured as *signed* transaction volume, order flow therefore provides an indication of the relative strength of buy and sell orders between, say, customers and dealers. In this way, order flow within particular investor groups will not necessarily sum to zero, but can instead exhibit persistent trends—for example if customers build a long position in a particular currency relative to an underlying strategic benchmark position, then order flow in that currency from this participant group will be positive and rising.

and trading strategies of better informed participants, and therefore represents the conduit through which informational asymmetries become embedded within market prices (Lyons, 1995; Bjønnes and Rime, 2001b).

The microstructure literature draws support and scepticism in equal measure (for more sceptical views, see Rogoff, 2002; Fisher and Hillman, 2003). Few disagree that the central hypothesis of this literature provides an intuitive explanation of the process of price discovery in the foreign exchange market. Disharmony surrounds the assessment of the practical value of this literature and, from the same quarter, discontent has also been expressed concerning the treatment given in the microstructure literature to assimilating the true structure of the foreign exchange market and the extent and motivation of interactions between the various participant groups. To date, theoretical predictions of this literature have been based upon highly stylised models of investor interaction (e.g. Evans and Lyons, 2002a,b). Ultimately, it is only through a detailed understanding of market structure and practice that one can fully appreciate the different impact on high frequency exchange rate returns of various categories of order flow.

Foreign exchange market structure and participant group interactions have altered substantially in recent years. Daily market turnover has increased in the last three years (BIS, 2004), with the majority of this increase concentrated upon the various customer groups in the market.³ The motivation for these flows has also altered, and investors increasingly view the foreign exchange market as a potential source of persistent returns that can provide important diversification benefits when combined in a portfolio with returns to more traditional assets, such as equities or bonds. The nature of interactions between customers and dealers has also undergone important changes in recent years, including rapid growth of electronic trading platforms—traditionally the preserve of the interdealer market—and important changes in the tactical interplay between customers and dealers.

The objective of this paper is to provide a comprehensive, up-to-date description of the structure of the foreign exchange market. In so doing, we marry the main tenets of existing theoretical models with the practical reality of the foreign exchange market.

The remainder of the paper is organised as follows. In the next section we pass a microscope over the foreign exchange market and, in addition to discussing trends in daily market turnover and geographic organisation, describe the various participant groups within the market and discuss the motivation and behaviour of each group, as well as group interaction. We then examine how closely both traditional asset-price exchange rate models and microstructural models reflect this structure. Drawing upon all of these aspects, we then present a thumbnail sketch of an alternative model that we believe would provide a better approximation to the foreign exchange market and yield testable implications. In a final section we draw some conclusions and offer suggestions for future research.

³The term customer encompasses asset management firms, hedge funds, commodity trading advisors (CTAs), central banks, corporations and high net worth private individuals. For expositional purposes, we exclude exchange rate transactions by tourists.

2 The Microstructure of the Foreign Exchange Market

2.1 Basic Features

The foreign exchange market is a decentralised market in the sense that market participants are generally separated from one another and transactions take place through electronic media such as by telephone or through computer networks. This is in contrast to, say, the New York Stock Exchange where traders physically interact with one another. Two implications of decentralisation are fragmentation and lack of transparency. The foreign exchange market is fragmented in the sense that transactions may (and do) occur simultaneously or near simultaneously in the market at different prices (Sarno and Taylor, 2001). It is opaque—or lacks transparency—in the sense that the absence of a physical market place makes the process of price-information interaction difficult to observe and understand (Dominguez, 1999; Lyons, 2002b).

The foreign exchange market is also the most liquid financial exchange in the world. Daily market turnover is estimated at approximately \$1.9 trillion, including spot, forward and swap transactions (BIS, 2004). This compares to daily turnover in the US government bond market of \$500 billion (Federal Reserve Bank of New York, 2005), and on the New York Stock Exchange of \$46 billion (NYSE, 2005). Although this estimate of daily turnover represents a substantial increase compared with 2001 (\$1.2bn; BIS, 2002), part of the improvement reverses the decline reported between 1998 (\$1.5 trillion) and 2001; nonetheless a longer-term upward trend in daily turnover is apparent.⁴ Including only spot transactions, daily turnover is approximately \$621 billion, emphasising the continued growth in the use of derivative contracts over spot transactions; the latter accounted for the majority of turnover in the late 1980s (BIS, 1999).⁵

Although the market is decentralised, it nevertheless has various physical locations or trading centres throughout the world where many of the most important market participants—in particular the market-making dealers—tend to concentrate. Thus, a foreign exchange transaction booked with the London office of a Japanese or American bank would be recorded as having taken place in the London market, even if it was conducted over a computer network and the customer was located elsewhere in the world. In this sense, London has consolidated its position as the major trading centre for foreign exchange, with 32% of daily market turnover transacted in this location (BIS, 2002). There are other important centres in New York and Tokyo, and smaller ones in Auckland, Sydney, Singapore, Hong Kong, Frankfurt and San Francisco. In terms of composition, the US dollar is on one side of approximately 90% of all daily transactions—this has changed little over the last fifteen years—followed by

⁴For instance, daily market turnover was estimated in 1989 at \$590 billion (BIS, 1996).

⁵Currency overlay specialists in asset management firms, as well as Hedge Funds and CTAs, implement currency hedging decisions primarily through the use of one or three month forward contracts. Forwards are preferred over spot or futures because they require no upfront payment and have no associated margin calls.

the euro (37%), yen and sterling (approximately 20% each; BIS, 2004). Euro-dollar is consequently the most liquid exchange rate, followed by yen-dollar and sterling-dollar. These rankings are also little changed over the last fifteen years, although the percentage of daily market turnover traded in sterling-dollar has risen by a third during this period, at the expense of yen-dollar.

2.2 Market Participants: Dealers and Customers

In broad generic terms, there are two types of participants in the foreign exchange market: dealers and customers.⁶ Dealers account for the majority of market liquidity, with approximately 59% of total daily market turnover occurring within the interdealer market. But the predominance of this group has been in decline for a number of years, and is down from over two-thirds in the early 1990s (BIS, 1996, 2002). Customers contribute the remaining 41% of daily market turnover. Of this figure, financial customer groups contribute the majority share, equivalent to 28% of total market turnover, and corporations 13% (BIS, 2002). Although a detailed breakdown of financial customer flows is not publicly available, our (albeit informal) discussions with a number of major investment banks suggest that asset management firms and hedge funds each account for approximately one quarter of customer flows, or together about 20% of total market flows, with commodity trading advisors (CTAs), central banks and individuals contributing the remaining 8%.⁷

2.2.1 Dealers

The interdealer market encompasses market-makers, leverage traders, designated proprietary or “prop” traders, and senior risk takers. Market-makers continue to perform their traditional core function of facilitating access for customers to interdealer liquidity and providing best execution for customer trades. But their wider role has evolved over recent years, in a number of ways. First, market-makers are typically allocated a book exchange rate upon which they focus their main attention. This compares with a few years ago when these individuals traded in a range of exchange rates. Second, Daniélsson and Payne (2001) suggest that market-makers often now focus upon one side of the market at a time, rather than posting genuine two-sided quotes. Put another way, market-makers are no longer typically the main source of price discovery in the foreign exchange market and do not attempt to generate excess profits from their market activity, but are instead largely facilitators of customer trades; in a similar vein, there has been a marked increase at many banks in the ratio of sales people to market-makers.

Prop traders have been characterised as intra-day, or even “nintendo”, traders (Bjønnes and Rime, 2003), whose investment time horizon extends from minutes

⁶Dealers are also known as broker-dealers; the two terms are synonymous.

⁷Market share data can vary substantially between investment banks who focus upon different segments of the customer market. Also, categorisation of flows between the various customer groups is somewhat arbitrary, with a number of investors spanning more than one segment. Consequently, these detailed market share data should be interpreted with caution.

to hours at most, and for whom there is no capacity to run overnight positions. In fact, this description is more consistent with the activities of leverage traders (also known as spot traders). These individuals trade primarily on the basis of order flow executed by the bank's trading desk, and typically have an investment horizon of a few hours, or at most days. The short-term nature of these traders' activity represents a key source of total market volatility, with tight stop-loss levels typically introduced around every position. As a result, price breaks of key technical levels typically extend further in the short term than they otherwise would in the absence of leverage interdealer trading (Osler, 2002, 2003).

In reality, prop traders are actually often discouraged by senior management from trading too actively and instead are encouraged to focus most of their allocated risk budget upon longer investment horizons than leverage traders—typically days but sometimes up to three months or even one year, depending upon the perceived opportunity set at any given time. Prop desk risk budgets are reduced overnight, reflecting the difficulty of monitoring positions outside of business hours, but remain positive.

An additional group of market participants—previously overlooked by academic analyses and surveys of the foreign exchange market—includes senior risk takers at large investment banks. These individuals perform a similar function to designated prop traders, but are allocated a much larger risk budget, for instance \$100-\$200 million compared with a typical \$25-\$40 million for designated prop traders, reflecting their relative seniority, experience and performance track record within the market.⁸ This risk budget will also be reduced overnight. Senior risk takers are similarly encouraged to focus upon relatively longer investment horizons, whenever the opportunity set allows.

A number of infrastructural trends are apparent within the interdealer market. First, industry consolidation has meant the number of banks that account for a majority of interdealer flows has fallen since 1998, with 17 banks in London and 13 banks in the US accounting for 75% of turnover transacted in these locations. This compares with 24 and 20 banks in 1998, respectively (BIS, 2002). Second, the general level of risk appetite within the interdealer market has declined in the wake of the 1998 Long-Term Capital Management crisis, leading to a reduction in the level of risk capital allocated to individual traders. Third, and related to the previous point, the rigour of risk management procedures and infrastructure related to dealer activity has improved substantially in recent years, to a generally high standard (Geithner, 2004). The predominant focus of risk control procedures remains the imposition of maximum intraday and overnight nominal—dollar—position limits for individual traders. Many

⁸Our estimate of the typical prop trader risk budget is smaller in absolute terms than Cheung and Chinn (1999), who studied the US sector of the market, and Cheung and Wong (2000), who examined trading activity and practices in Hong Kong and Singapore. But it is directionally consistent with both of these studies given broad market trends described below. Exact position limits are considered to be market-sensitive information, and are therefore confidential to individual banks. Our estimates result from informal conversations with market participants.

banks also make the amount of risk capital available to traders a function of past performance, rewarding good performance with an increase in risk capital and penalising bad results. Typically, traders assume responsibility for—or, in market jargon, “wear”—all losses, so that these reduce next month’s available risk capital on a commensurate basis. By contrast, profits are typically shared with the bank, such that risk capital available to a trader in the next month will increase only by some fraction of last month’s reported profits. Trading risk is also monitored using a variety of other metrics, including daily maximum drawdown, or capital loss, and daily Value at Risk (VaR) limits⁹; the use of VaR limits appears now to be more commonplace than reported by Cheung and Wong (2000). A loss that trips intraday limits is unlikely to signal the termination of a trader’s risk budget, but instead will trigger notification of senior management, require a detailed explanation of the circumstances surrounding the loss and of remedial steps to be taken, and closer scrutiny of the trader’s outstanding positions and activity in the immediate future. More generally, banks also conduct VaR sensitivity analysis on the activity of its trading desk in aggregate to ensure that senior management have a broad understanding of the probability of an extreme event that could threaten the solvency of the institution.

Fourth, most interbank trading occurs electronically, via either the Electronic Banking System (EBS) or Reuters D3000. Both systems were established in 1993 and were the primary facilitators of the subsequent marked increase in market liquidity. Their functionality is essentially equivalent, providing ex ante anonymous limit order bid-ask pricing to dealers. Combined, these systems account for approximately 85% of total interbank activity, with EBS dominating in all exchange rates other than sterling, Canadian and Australian dollar cross rates.¹⁰ Remaining interbank trading activity is shared by the remnants of voice broker (10%) and secure bank-to-bank chat lines provided by Reuters (5%). The dominance of EBS in most liquid exchange rates raises a question-mark over the robustness of existing empirical research which mainly uses data from Reuters to test the key predictions of the microstructure literature.

2.2.2 Customers

Customers interact with dealers in order to access the liquidity of the interbank market. Voice trading continues to account for the majority of customer trades. But three major electronic systems—FX Connect, FXAll and Currenex—have recently been introduced to the customer-dealer space.¹¹ All three systems are

⁹ Value at Risk (VaR) is defined as the maximum percentage value of an investment portfolio that could be lost during a fixed period (e.g. one day) within a certain confidence level (e.g. 95%).

¹⁰ Why this delineation by exchange rate originally arose is not clear. Unsuccessful attempts have periodically been made by both systems to gain market share in additional exchange rates.

¹¹ Volumes transacted on these systems have grown rapidly since their introduction. For instance, FXAll reported a rise in volumes of 104% in 2004, to an annual total of \$4.9 trillion. FX Connect transacts a similar annual volume to FXAll, and these two systems combined

multi-bank electronic trading portals linked directly to customer trading desks. They comprise an essential component of the push towards the introduction of Straight Through Processing (STP) that facilitates the complete automation of customer foreign exchange management from order creation through electronic trading portals to the settlement and confirmation of trades. A primary motivation behind the introduction of STP is a reduction in the risk of human error at various points in this process (Baird, 2002). The emergence of electronic portals in the customer-dealer space is unlikely to increase the availability of order flow data, and thereby will not improve the transparency of the foreign exchange market and the quality of microstructure-based exchange rate models as electronic order flow is treated as strictly confidential by system governing boards.¹²

Except for corporate hedging of the translation risk of international revenues and costs, all other customer activity in the foreign exchange market relates to the management of currency exposure in investment portfolios.¹³ By reputation, customer order flow is the most important source of private information in the foreign exchange market (Lyons, 1995; Ito, Lyons and Melvin, 1998; Bjønnes and Rime, 2003), with the sign of customer trades considered more informative than the associated nominal value of these trades (Bjønnes and Rime, 2001b). An important reason for this reputation is the heterogeneity that exists within this segment of the foreign exchange market: informational asymmetries; different reaction speeds to data innovations; diverse opportunity sets and risk/return expectations. We discuss each one of these heterogeneities in turn in the following sections. But in general terms, all of them suggest that a broad understanding of the behaviour of an exchange rate at a given time requires knowledge of the types of customers prevalent in the market at that time and of the ways in which they trade and interact with the wider market.

Passive versus active customers¹⁴ It is estimated that only a small fraction of daily customer flows are initiated by profit-maximising investors, typically in currency overlay firms, hedge funds and CTAs (Knott, 2002). Instead, the majority of flows reflect non-price level dependent hedging activity. Customer foreign exchange exposure occurs in a number of ways, reflecting the presence in the market of both *passive* and *active* customers. Passive customers inherit foreign exchange exposure from the sale and purchase of underlying assets, such

account for approximately 86% of all electronic customer trades conducted through multi-bank portals.

¹²This view appears contrary to that of Evans and Lyons (2002a).

¹³For instance, on behalf of pension funds of private sector corporations, government and supranational institutions, and wealthy individuals.

¹⁴Our use of the terms "active" and "passive" differs from the sense of Evans and Lyons (2002a,b), but is more consistent with market tradition. Whether a customer is passive or active is determined not by the group from which they are drawn—for example, corporations versus asset managers—but by the nature of their activity in the foreign exchange market. There are similarities in the categorisation of customers in this section with Kyle (1985). He defines a market in which three types of participants coexist: nondiscretionary liquidity, discretionary liquidity and informed traders.

as international equities or bonds, around a strategic portfolio benchmark, or from the accrual of international revenues and costs by corporations. This exposure is addressed by passive investors in one of two ways. First, and relating to financial customers, it is left unhedged. In this case, customers essentially adopt a policy of benign neglect towards currency exposure, presumably in the belief that returns to active currency management average zero over the long term. That this neutrality hypothesis has proven invalid by observation during the post-Bretton Woods period (Baldrige, Meath and Myers, 2000; Hersey and Minnick, 2000; Parker Global, 2004) suggests that the behaviour of these participants is inconsistent with the Rational Expectations Hypothesis (REH). Embarrassment risk associated with long-term underperformance of this policy of benign neglect should be a significant factor encouraging these investors to switch to the active customer camp; this switch may be encouraged by the trend towards more intense scrutiny of longer-term investment performance by pension fund trustees.

The second way that inherited currency exposure is addressed by passive customers—both financial and corporate customers—is by implementation of hedges that return exposure back to an underlying strategic benchmark position. These hedges are typically implemented without tactical consideration of the exchange rate at which they are conducted.¹⁵ As well as emphasizing the lack of a profit maximising motive on the part of many foreign exchange customers, this behaviour also seems to contradict the Efficient Market Hypothesis (EMH) assumption of investor risk neutrality; passive participant groups appear willing to pay a premium to other, active customer groups in order for these investors to assume their short-term foreign exchange risk exposure (Kearns and Manners, 2004).

Foreign exchange exposure is introduced by active customers into portfolios as part of an active currency programme. This behavior is the most consistent with the EMH, although one can reasonably question whether even these customers, or their associated investors, exhibit risk neutrality given the extent of evidence against this assumption generally (Lewis, 1994). Within active currency programmes, appointed currency managers are given discretion by investors to add value to underlying portfolios by implementation of appropriate active hedges around a strategic currency benchmark. These hedges are subject to explicit portfolio guidelines relating to position limits and admissible exchange rates that are designed to ensure that the contribution to total portfolio risk of these hedges is consistent with ex ante investor expectations. Strategic currency benchmarks have traditionally been defined in relation to a set of assets or liabilities to which the investor has a long-term exposure; this is the format of a traditional currency overlay programme. More recently, investor interest has grown in currency programmes implemented on the basis of an underlying notional capital value separate to this underlying asset exposure with a bench-

¹⁵This is the case for short-term hedging by corporate customers. This group also undertakes medium-term hedging that does have regard to exchange rate levels. However, market anecdote suggests evidence of associated systematic forecasting errors, undermining a key tenet of the Efficient Market Hypothesis.

mark equal either to zero or a cash–risk-free–interest rate; this investor interest has largely accounted for the recent rise in popularity of currency hedge funds, and leveraged funds offered by established asset management firms.¹⁶

Informed versus uninformed customers An alternative customer classification would divide participant groups into informed and uninformed investors. Which category a particular group belongs to reflects their ability to interpret the impact of data innovations upon future exchange rate returns. Most obvious amongst informed customers would be large central banks, who both observe innovations in fundamental data series relevant to their own currency in advance of other market participants, and who have access to the best streaming information on intraday order flow in their sector of the foreign exchange market. Reflecting both the sophistication of this customer group, as well as the magnitude of associated flows, central bank foreign exchange activity is among the most keenly sought source of private information. This information can relate either to prospective market intervention or compositional changes in the currency denomination of official foreign exchange reserves. For instance, on a number of occasions since the inception of the euro in January 1999, and particularly during 2003 and 2004, market anecdote has suggested important portfolio adjustments away from the US dollar and towards the euro amongst Middle Eastern, peripheral European and Asian central banks. As such activity occurs transacting dealers gain important information that can be used to inform their own foreign exchange trading strategies, at least in the short term.

An increasing number of central banks allocate a small part of official foreign exchange reserves to wealth creation, reflecting both an increase in the stock of reserves, particularly in Asia, and less need for these reserves to be held in liquid form due to wider use of derivative instruments. As a result, market anecdote suggests that central banks are now an active part of daily traded flows, and that investment activity on this portion of foreign exchange reserves has proven successful.¹⁷

Investment styles Customer heterogeneity, and contradictions of the EMH, can be observed in the range of investment styles that coexist within active–and informed–customer groups. These styles exploit different types of information–as well as the same information in different ways–leading to differences in exchange rate expectations.¹⁸

¹⁶The long-term performance of active currency managers compares favourably to that of managers in other asset classes, including equities and bonds (Morris, 2004). This performance provides a stark, and robust contradiction to the suggestion by Federal Reserve Chairman Alan Greenspan that the practice of active currency management is akin to a coin tossing contest (Greenspan, 2004).

¹⁷By contrast, evidence is mixed on whether foreign exchange intervention by central banks has been a profitable activity during the floating exchange rate era. For instance, see Becker and Sinclair (2004) and Dalio (2002) for contrasting conclusions.

¹⁸There has been some blurring of differences between investment styles in recent years as the product ranges of the various active, informed customers have converged.

Many CTAs are pure technical, or “black box”, managers. The associated investment process will typically comprise a set of optimised technical, or chartist, trading rules that have no intuitive underlying theoretical economic interpretation.¹⁹ News to this customer group is historical price innovation, over any period from minutes, hours, days, all the way out to years, combined with detailed trading rules related to key support and resistance levels, moving average cross-over levels, over-bought and over-sold calculations²⁰ and a range of other price patterns. Publicly announced macroeconomic news is only relevant indirectly, to the extent that it has some historic price impact. The importance of this type of technical trading within the foreign exchange market, particularly for high frequency exchange rate returns, is confirmed by Allen and Taylor (1990), Taylor and Allen (1992), Cheung and Wong (2000), and Euromoney (2002). This type of trading is often considered irrational, and therefore inconsistent with the EMH, as it represents an important source of systematic forecasting errors. In reality, the extent to which technical investors have demonstrated an ability to generate persistent excess returns suggests that they are in fact acting rationally and exploiting the proven unit root properties of spot exchange rates in the vicinity of equilibrium (Taylor, Peel and Sarno, 2001).

Hedge funds and currency overlay managers initiate order flow predominantly on the basis of publicly available macroeconomic information, with an expected pay-off schedule that is likely to stretch from around one to three months. The investment process of these customers is often highly quantified, with pre-defined trading rules based upon theoretical relationships between economic or financial variables and exchange rates. In this case, public data innovations will directly trigger customer order flow. Most hedge funds and currency overlay managers also employ technical analysis similar to CTAs²¹, and some also introduce tactical exchange rate hedges into portfolios on the basis of a purely qualitative interpretation of fundamental and order flow data, on the rationale that not all events-including central bank intervention-that determine exchange rate returns can be quantified in a consistent manner over time. The profitability of qualitative trading will rely upon the skill of portfolio managers to interpret the impact of order flows and market reaction to these events in a consistently accurate manner. Associated trades will-by design-typically be risk-reducing at the total portfolio level, thereby improving the Information

¹⁹Market anecdote suggests that some CTAs have begun to embrace fundamental-based trading strategies, in an effort to diversify returns. Others have adopted a more discretionary investment style. Models will be used to determine key technical levels and turning points but positions will be implemented on a discretionary basis by portfolio managers and traders. CTAs also appear to be increasing the breadth of portfolios by trading in a greater range of exchange rates; traditionally, trading activity concentrated on euro-dollar, yen-dollar and euro-yen.

²⁰Over-bought and over-sold calculations attempt to define when an exchange rate has moved too far and fast in either direction. They are typically calculated based on a moving average of the difference between the number of advancing and declining days over a certain period of time.

²¹A testable implication of the increasing use of technical analysis is that half-lives of exchange rate disequilibria should have risen over recent years, particularly for the most liquid exchange rates. We are grateful to Dirk Morris for this observation.

Ratio of the investment strategy.²²

Finally, risk-control currency managers introduce option replication strategies into client portfolios in order to minimise the downside risk attached to any level of foreign exposure over some fixed investment horizon, typically one year (Layard-Liesching, 2002). Consequently, this type of manager will react to price and data innovation indirectly to the extent that these affect the downside risk profile implicit within portfolios.

Reaction speeds Customer heterogeneity implies differences in the reaction speed of various customers to the arrival of new information. Many hedge funds are among the quickest to respond to significant data innovations by immediately trading in size to implement desired changes in portfolio construction. As hedge funds typically make more use of leverage than other active—or informed—customers, significant shifts in portfolio positioning will typically represent a larger proportion of total assets than is the case for other customers. While this implies that hedge funds can potentially benefit more than most other customers from significant data innovations, the downside of this approach is that sudden strategy shifts do not allow hedge fund traders to exploit intra-day market liquidity—by feeding orders into the market—to ensure that trades are implemented at best price;²³ another, more systemic consequence of this tendency to effect sudden shifts in positioning is that this behaviour of hedge funds can amplify the impact of shocks to the foreign exchange market during periods of significant market stress (Geithner, 2004).

By contrast, and at odds with the REH that argues public information is instantaneously reflected in exchange rates, currency overlay firms typically exploit investment strategies that build or reduce active hedges gradually over the course of several trading sessions in response to significant data innovations. This contradiction is consistent with practical observation that the contemporaneous relationship between order flow and exchange rates is relatively persistent (Bjønnes and Rime, 2001a; Froot and Ramadorai, 2001; Fan and Lyons, 2001; Evans and Lyons, 2002a,b). It may be explained by a desire to minimise the impact of portfolio shifts on exchange rate returns, and therefore expected profits, given the size of assets under management; this explanation is only valid as long as one believes that exchange rate returns exhibit significant positive serial correlation.

Differences in customer reaction speeds can provide dealers with important information about the size and sign of customer trades, and the likely persistence of this trading. To the extent that a dealer sees order flow from large active,

²²Information Ratio is defined as a measure of excess return per unit of risk taken, where excess return is defined as the difference between the currency portfolio's return and the return to a benchmark index.

²³It may also be the case that sudden shifts in hedge fund positioning are due to position liquidation as a result of significant losses; leverage implies that the probability of significant losses will be higher within the hedge fund community than amongst asset management firms. The recent trend towards contingent credit agreements mitigates this probability, to some degree (Geithner, 2004).

informed currency customers it is generally reasonable to assume that this is only a small part of the total trade being executed, and that the remaining orders are likely to be fed into the market throughout the current trading session (in the case of hedge funds) or several trading sessions (currency overlay managers, for instance). This knowledge will allow the dealer either to “piggy-back” on the trade, committing some of his own risk capital to the same trade (Bjønnes and Rime, 2001b), or to net off trades from other customers.

Information about large customer order flow is made available by dealers to other large customers, on a quasi-anonymous basis via direct voice links, once a dealer has executed a customer order. An indication of the exchange rate, type of initiating customer, and size and sign of the transaction will be provided, as well as the ease with which the market absorbed the order.²⁴ This information helps other customers gauge the extent of technical support and resistance levels around current spot exchange rates, including any option-based trading structures, and the probability and likely extent of any break-outs from prevailing price trend channels.

In a similar manner, ad hoc information concerning the level and volume of limit orders placed by customers on dealer order books will be provided by dealers to preferred clients on a daily basis. Osler (2002) finds that information regarding limit orders has historically had an important explanatory role for yen-dollar. But the quality of this information is likely to have deteriorated in recent years as customers have become more reticent to place limit orders with dealers, reflecting greater cogniscence of their own strategic market impact. For informed customers, this cogniscence has increasingly focused efforts upon ensuring that internal risk management systems can incorporate appropriate, continuous position gain-loss monitoring procedures to allow creation and execution of trades as stop levels are approached.

3 Theoretical Models of the Foreign Exchange Rate

Given the foregoing discussion of the microstructure of the foreign exchange market, we are now in a position to analyse alternative approaches to modelling the exchange rate, beginning with traditional asset-price exchange rate models and then moving to more recent microstructural models.

3.1 Traditional Asset-Price Models

Traditional asset-price models of exchange rate determination take the form:

$$S_t = \beta' F_t + \alpha S_{t+1}^e, \tag{1}$$

²⁴For instance, "Asian names" is a market pseudonym for the major central banks in Asia. But dealers will never explicitly provide the name of customers initiating trades to other customers, in order to ensure client confidentiality.

where S_t is the spot exchange rate at time t , and S_{t+1}^e is the one-period-ahead expected spot rate given information at time t (both expressed as domestic currency per unit of foreign currency). F_t represents a vector of fundamental variables that exhibit some persistent explanatory power for the determination of exchange rates, and β represents a vector of factor loadings. This formulation is sufficiently general to encompass simple monetary models of the exchange rate, sticky-price overshooting models and portfolio balance models (see Taylor, 1995; Sarno and Taylor, 2002; Lyons, 2001, chapter 7).

The first thing to note about this class of models is that they are explicitly macroeconomic and apparently independent of either the way information is processed in the market or the institutional structure of the market (Lyons's "two i 's" —Lyons, 2001). Thus, asset-price models are essentially equilibrium models that marginalise the importance given to the means by which that equilibrium is reached or the institutional setting in which currency prices are determined. Moreover, whose expectation is S_{t+1}^e ? Many of these models were developed (or have been interpreted) under the strict assumption that agents are endowed with rational expectations, so that S_{t+1}^e is the true conditional mathematical expectation. If that is the case, then news on fundamentals is instantly imparted into the market as soon as it is released and the exchange rate will jump to its new equilibrium level. This cannot, however, be the whole story, since this would imply a complete consensus of opinion among all market participants as to the precise change in "fair value" due to any data innovation; furthermore, our discussion above of the behaviour of the various participant groups in the foreign exchange market makes clear that this is not the practical reality of the market. A less restrictive assumption, which has its roots in Friedman's (1953) apologia of floating exchange rates, would be that the market in some average sense conforms to the rational expectations hypothesis and reaches this equilibrium because uninformed participants are driven out of business (or to change their behaviour) by informed participants. In this case, equation (1) represents a snapshot of equilibrium in the market as dependent on macroeconomic fundamentals, but says nothing at all about how the market gets to that equilibrium. It is perhaps for that reason that asset-price models have performed spectacularly poorly in explaining short-run exchange rate movements (Meese and Rogoff, 1983; Taylor, 1995; Sarno and Taylor, 2002), although there is some evidence that they may be able to explain longer-run movements. Indeed, given that the profession seems to have moved towards support for some form of long-run purchasing power parity equilibrium for exchange rates (see e.g. Taylor and Taylor, 2004), then grafting this onto an existing consensus of long-run monetary neutrality implies that the long-run behaviour of the exchange rate must reflect relative movements in monetary velocity between countries (see e.g. Frankel and Rose, 1995; MacDonald and Taylor, 1993).

Viewed as a long-run equilibrium relationship, therefore, asset-price models may be innocuous but they are also anodyne: they tell us nothing about short-run exchange rate behaviour and, in particular, nothing at all about how information on macroeconomic fundamentals gets compounded into the current exchange rate; they draw a veil over the most interesting part of the show.

As Lyons (2001) notes, microstructural approaches to exchange rate determination are not necessarily at odds with asset-price models. For example, if—as predicted by most asset-price models—relative monetary velocity affects the exchange rate, then agents will not wait for this month’s money supply figures to be released before adjusting the exchange rate accordingly. Rather, an excess increase in the money supply (decline in velocity) will itself have created downward pressure on the external value of the currency as interest rates decline (perhaps as a result of the central bank’s enthusiastic purchase of government bonds in their open market operations) and net capital inflows diminish. While the release of the money supply figures at the end of the month will explain the downward movement in the foreign price of domestic currency (i.e. the exchange rate), the effect of the excess money supply growth will have been reflected throughout the month in a rise in negative order flow associated with the domestic currency; this is one explanation why order flow appears to be highly correlated with contemporaneous exchange rate movements (Evans and Lyons, 2002a).

But while customer order flow in particular may represent the missing piece of the exchange rate puzzle, this does not mean that these data can necessarily be exploited for the purposes of predicting exchange rate movements even at very short horizons, since they are only contemporaneously observable on a real-time basis by a limited number of well-informed market participants. These participants include the custodian or investment bank or electronic trading platform that collects and collates the data, and on an *indirect*, verbal basis, a select group of preferred customers of these organisations—for instance large asset management firms, hedge funds and CTAs.²⁵ Limited access to order flow data reflects the confidentiality concerns of major market participants whose trading activity is captured by these data and which could therefore be identified by competitors, and a wish on the part of the collecting institution to sustain and profit from this potential information advantage for as long as possible. For this reason, *direct* access to proprietary order flow data by preferred customers is provided only after the raw data have been filtered and collated into an index to ensure customer anonymity, and with at least some publication lag.²⁶ Dealers in smaller banks see most customer order flow data only indirectly once they have been embedded in prices. Furthermore, most customers do not gain direct access to even filtered flow data, and instead have to content themselves with interpretative analysis provided by dealers that includes qualitative trading strategies based upon developments in their own collated flow data (HSBC,

²⁵Preferred clients are those identified by investment banks to be of strategic importance to their position in the foreign exchange market. As well as best execution, these clients will be provided with better streaming information than other clients on intraday and daily order flow and market positioning, and will be allocated resources for collaboration on confidential, customised research. Importantly, information concerning order flow data made available to preferred customers never includes details of individual customer transactions, so that customer anonymity is assured.

²⁶As we report in a companion paper, filtering and publication lags appear to significantly reduce the information content of available order flow data, from a customer perspective (Sager and Taylor, 2005).

2003; Citibank, 2003). (We address the issue of the predictive content of order flow for exchange rate movements in Sager and Taylor (2005).)

In summary, while traditional asset-price exchange rate models may possibly tell us about equilibrium conditions or “fair value” in the foreign exchange market, they ignore extremely important aspects of the foreign exchange market which are evident in our preceding discussion of its microstructure. They can tell us nothing at all about the way in which heterogenous groups interact with one another in the market and transmit information to one another. In other words, they ignore issues of information transmission and institutional structure.

3.1.1 The Evans-Lyons Model

The academic literature presents a number of stylised models of the interaction between the various customer groups and dealers described above. One of the most recent and influential is due to Evans and Lyons (2002a,b), which in turn borrows aspects of Kyle’s (1985) sequential equilibrium model. The model is consistent with (but does not rely on) equilibrium or “fair value” in the market being determined in accordance with asset-price models of the kind discussed above, but also attempts fill in missing gaps in these models in terms of information transmission and institutional features. Its broad characteristics may be set out as follows.

There are N dealers in the market, a continuum of non-dealers (or customers), and an infinite number of trading days. Dealers observe a periodic payoff on foreign exchange, denoted R_t , representing the flow of macroeconomic information. This payoff is proxied by Evans and Lyons by interest differentials. There are three stages to the foreign exchange trading day. First, and having observed R_t at the start of each day, dealers independently and simultaneously set bid-ask spreads for, and trade with, customers. Second, dealers trade amongst themselves, independently and simultaneously posting a bid-ask spread for other traders. These spreads lead to trades, as dealers spread risk generated by earlier customer trades through the interbank market. Once this second round of trading is complete, all dealers are able to observe the order flow, X_t , that has occurred from interdealer trading during that day. These data are assumed by Evans and Lyons (2002a,b) to convey important information about the size and sign of customer trading during Round One. Finally, in the third round of trading dealers once more trade with customers, in order to share overnight risk more widely across the market. A crucial assumption here is that dealers set prices in Round Three such that customers willingly absorb all dealer inventory imbalances, so that dealers run no overnight open positions. Consequently, the closing price at the end of each day within the Evans and Lyons model will be

$$P_t = \beta_1 \sum_{\tau=1}^t \Delta R_t + \beta_2 \sum_{\tau=1}^t X_t \quad (2)$$

and the change in prices from the end of day $t - 1$ to the end of day t can be written as

$$\Delta P_t = \beta_1 \Delta R_t + \beta_2 X_t \quad (3)$$

where $X_t = \sum_{i=1}^N T_{it}$, is the total order flow generated by interdealer trading during the day and $T_{it} = \alpha C_{it}^1$ is the sum of customer orders received by individual trader i during the first round of trading; α is a constant coefficient. A key distinction in the model exists between “active” and “passive” customers. Round One customers are active in the sense that they generate initial imbalances in dealers’ positions, while Round Three customers are passive in that they passively absorb dealers’ positions at the end of the day.²⁷

Although the Evans and Lyons (2002a,b) characterisation of the foreign exchange market is a useful taxonomic device for clarifying one’s thoughts about market microstructure and customer-dealer interactions, we believe that it should not be interpreted too literally as a description of the foreign exchange market (as it appears to have been in a number of applied microstructural studies of the foreign exchange market). In particular, the notion that customers willingly absorb the daily inventory imbalance of dealers seems at odds with market practice. Indeed, customers have traditionally paid dealers to assume this type of price and credit risk on their behalf, thereby preserving customer anonymity and allowing access to the liquidity of the interbank market through a relatively small number of contact points.²⁸ This payment is made in soft terms, with dealers traditionally able to divide customer trades into smaller tranches as they spread risk through the interbank market. As the width of bid-ask spreads is a positive function of transaction size for transactions beyond a minimal value (Bjønnes and Rime, 2001b), tranching saves the dealer transaction costs and, along with knowledge of, and trading around customer limit order books, has traditionally represented a major source of trading profit for banks (Rime, 2001).

This profit source has recently been undermined by the behaviour of those customers increasingly aware of the strategic price impact of their trading activ-

²⁷Evans and Lyons’ use of the terms “active” and “passive” differs from ours above, or market norms.

²⁸It is standard practice for a currency overlay manager (the agent) to request prior approval from a client (the principal) to trade with a maximum of ten dealers on a continuous basis on the client’s behalf. This ensures that the price and credit risk of the client with respect to any one of the dealers is minimised and that the overlay manager maximises its opportunity to achieve best execution for the client. More recently, interest has grown in prime brokerage agreements, particularly amongst hedge funds. Under these agreements, all of a customer’s foreign exchange orders are transacted through a single dealer. The prime broker is usually a large, highly rated institution that allows the customer to initiate trades, subject to credit limits and collateralisation, with a pre-determined limited group of high quality third party banks (so-called “Spoke” banks) in the prime broker’s name, rather than the customer’s. Consequently, customer credit risk is limited to the prime broker. These agreements also allow for a reduction in operational risk for customers, as all positions are with the prime broker and only one dealer back office will be involved in trade reconciliation and netting (E-Forex, 2003).

ity. For example, hedge fund traders will often trade against their own positions in small size in order to encourage the interdealer community to adopt similar positions, subsequently reversing this tactic with substantially greater volumes to effect a short squeeze on the market. In a similar vein, these traders will often execute many trades simultaneously across the interdealer community, rather than feeding them into the market, in order to maximise the noise surrounding their activity and its subsequent price impact. Moreover, currency overlay managers will now regularly engage in the tranching activity traditionally employed by dealers, dividing trades of \$100 million-\$1 billion into smaller amounts of, say, \$25-\$50 million. These tranches are then spread amongst dealers intermittently during the course of one or more trading sessions. Although consistent with customer profit-maximisation, this practice undermines the trading profits of dealers and implies that banks are now accepting risk on behalf of customers without adequate compensation in return. How dealers can regain payment for this risk in the future is not immediately apparent, given also the rise of electronic trading that implies both a compression of dealer bid-ask spreads on this segment of trading and an increase in ex ante customer anonymity (Portes, 2002).²⁹ Consequently, risk sharing in the foreign exchange market would seem to be an important area of future research.

3.2 A Thumbnail Sketch of a Push-Pull Model

Given our discussion of the structure of the foreign exchange market, we can sketch the features of a theoretical framework that would correspond more closely to the real world than existing theoretical models while paying close attention to Lyons’s “two *i*’s” (information and institutions) and having, in principle at least, have testable implications.

First, there would be only one round to the trading day and customers would be free to trade at any time. A key distinction between customers would hinge on whether they are active or passive, and informed or uninformed—in the sense that we have used these terms in this paper—and therefore whether they push prices or are pulled into the market by a price movement. Net customer orders from either large active customers, or those considered by other market participants to be particularly well informed, would be expected to initiate price movements—these would be the so-called “push customers”. Other customers, perhaps passive or less well informed, or whose investment process is predominantly driven by technical analysis, are attracted or pulled into the market by price movements; effectively they exercise an option to trade once the price crosses their implicit “strike price”. These would be the “pull customers”. There would be no requirement for dealers to square positions overnight, and no assumption that customers absorb dealers’ daily inventory. There would also be evidence of rational errors, no assumption of risk neutrality for any participant group and, depending upon the particular customer group, no profit motive from exchange rate hedging.

²⁹ Manipulation of bid-ask spreads in response to larger size, ad hoc customer order flow remains a potential source of profits for dealers.

Among the various customer groups participating in the market, it would be possible to test hypotheses concerning whether a certain category could be described as “push” rather than “pull” by examining the correlation of their order flow with subsequent price movements: a positive and significant correlation would indicate that the customer is pushing the market; a negative and significant correlation would be consistent with customers being pulled into the market to buy as prices drop and to sell as prices rise.³⁰

4 Conclusions

The foreign exchange market incorporates a heterogeneous set of participants. Heterogeneities include informational asymmetries, different reaction speeds to significant data innovations, diverse opportunity sets and risk-return expectations. The behaviour and interaction of market participants, which both generates persistent inefficiencies and ensures that the process of price discovery is opaque, is poorly understood. Existing theoretical models of market microstructure make a number of simplifying assumptions that bear little resemblance to actual market practice. In particular, although most models assume that customer participant groups are the main source of private information in the foreign exchange market, theoretical work has concentrated upon the behaviour of the interdealer market, overlooking the subtleties of the customer segment. On the other hand, traditional asset-price models of exchange rate determination ignore crucial issues of information dissemination and institutional structure.

In this paper we set out to improve understanding of foreign exchange market structure, and participant group behaviour and interaction. We reviewed recent structural developments in the market, including trends in daily market turnover, geographic concentration, liquidity of individual exchange rates and trading and settlement of customer-dealer transactions. We also presented a detailed description of the behaviour of individual participant groups within both the interdealer and customer segments of the market, as well as the ways in which these various groups interact. This description allowed us to sketch a model of market behaviour that differs in a number of ways from existing theoretical models. A model conforming to our “push-pull” framework would represent a necessary practical foundation to the market microstructure literature that emphasises the role of order flow in the determination of high frequency exchange rate returns.

³⁰In a recent paper, Carlson and Osler (2005) develop a theoretical and empirical model that conforms to some of the main features of our suggested framework.

References

- [1] Allen, H., and M.P. Taylor (1990), Charts, Noise and Fundamentals in the Foreign Exchange Market, *Economic Journal*, 100, pp. 49-59.
- [2] Baird, D. H. (2002), Integrating eFX with Institutional Investor Needs. JPMorgan Fleming Asset Management Mimeo.
- [3] Baldridge, J., B. Meath, and H. Myers (2000), Capturing Alpha Through Active Currency Overlay, Frank Russell Company.
- [4] Bank for International Settlements (1996), Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1995 - Final Results. www.bis.org.
- [5] Bank for International Settlements (1999), Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1998 - Final Results. www.bis.org.
- [6] Bank for International Settlements (2002), Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity 2001 - Final Results. www.bis.org.
- [7] Bank for International Settlements (2004), Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity 2003 - Preliminary Results. www.bis.org.
- [8] Becker, C. and M. Sinclair (2004), Profitability of Reserve Bank Foreign Exchange Operations : Twenty Years After the Float. Reserve Bank of Australia Research Discussion Paper # 2004-06.
- [9] Bjønnes, G.H. and D. Rime (2001a), US Exchange Rates and Currency Flows. Unpublished mimeo Stockholm Institute for Financial Research. www.sifr.org.
- [10] Bjønnes, G.H. and D. Rime (2001b), Customer Trading and Information in Foreign Exchange Markets. Working paper # 4. Stockholm Institute for Financial Research.
- [11] Bjønnes, G.H. and D. Rime (2003), Dealer Behavior and Trading Systems in Foreign Exchange Markets. Norges Bank working paper 2003/10. <http://www.norges-bank.no/publikasjoner/arbeidsnotater/pdf/arb-2003-10.pdf>.
- [12] Brown, B.W. and S. Maital (1981), What Do Economists Know? An Empirical Study of Experts' Expectations.. *Econometrica*, 49, 2, 491-504.
- [13] Carlson, J.A. and C.L. Osler (2005), Short-Run Exchange Rate Dynamics: Theories and Evidence. Mimeo, Brandeis International Business School.

- [14] Cheung, Y.-W. and M. D. Chinn (1999), Currency Traders and Exchange Rate Dynamics: A Survey of the U.S. Market". *Journal of International Money and Finance* 20, 439-471.
- [15] Cheung, Y.-W. and C. Y.-P. Wong (2000), A Survey of Market Practitioners' Views on Exchange Rate Dynamics. *Journal of International Economics*, 51, 401-419.
- [16] Citibank (2003), Order Book and Option Vols. Report. Citibank.
- [17] Consensus Economics (2004), Foreign Exchange Consensus Forecasts. <http://www.consensuseconomics.com/>.
- [18] Dalio, R. (2002), Inefficiencies and Opportunities in the Currency Markets, in *A Guide to Currency Overlay Management*, D. Knott (ed.) Deutsche Bank. Global Market Research.
- [19] Daniélsson, J. and R. Payne (2001), Measuring and Explaining Liquidity on an Electronic Limit Order Book: Evidence from Reuters D2000-2. LSE Financial Markets Group Mimeo.
- [20] Dominguez, K. M. (1999), The Market Microstructure of Central Bank Intervention. NBER Working Paper # 7337.
- [21] E-Forex (2003). Foreign Exchange Joint Standing Committee E-Commerce Subgroup Report.
- [22] Euromoney (2002), 2nd International Forex Forum. www.euromoney.com.
- [23] European Central Bank (2003), Review of the Foreign Exchange Market Structure. www.ecb.int.
- [24] Evans, M. and R.K. Lyons (2001), Order Flow and Exchange Rate Dynamics. *BIS Papers* #2. <http://www.bis.org/publ/bispap02.htm#pgtop>.
- [25] Evans, M. and R.K. Lyons (2002a), Order Flow and Exchange Rate Dynamics. *Journal of Political Economy*, February, 170-180.
- [26] Evans, M. and R.K. Lyons (2002b), Informational Integration and FX Trading. *Journal of International Money and Finance*, 21, 807-31.
- [27] Evans, M. and R.K. Lyons (2003), *New Micro Exchange Rate Economics*, typescript, U.C. Berkeley.
- [28] Fan, M. and R.K. Lyons (2001), Customer Trades and Extreme Events in Foreign Exchange. U.C. Berkeley Mimeo.
- [29] Federal Reserve Bank of New York (2005). <http://www.ny.frb.org/>.
- [30] Fisher, P. and R. Hillman (2003), Explaining and Forecasting Exchange Rates with Order Flows. Bank of France Economic Policy Forum. <http://www.banque-france.fr/fr/fondatio/telechar/paulfisher.pdf>.

- [31] Frankel, J. A. and K. Froot (1985), Using Survey Data to Test Some Standard Propositions Regarding Exchange Rate Expectations. NBER Working Paper # W1672.
- [32] Frankel, J.A. and A.K. Rose (1995), Empirical Research on Nominal Exchange Rates, in G.M. Grossman and K. Rogoff (eds.) Handbook of International Economics, III, pp.1689-729. Amsterdam: North-Holland.
- [33] Friedman, M. (1953), Essays in Positive Economics. University of Chicago Press.
- [34] Froot, K. A. and J.A. Frankel (1989), Interpreting Tests of Forward Discount Bias Using Survey Data on Exchange Rate Expectations. NBER Working Paper # W1963.
- [35] Froot, K.A. and K. Rogoff (1995), Perspectives on PPP and Long-Run Real Exchange Rates, in G.M. Grossman and K. Rogoff (eds.) Handbook of International Economics, III, pp.1647-88. Amsterdam: North-Holland.
- [36] Froot, K.A. and T. Ramadorai (2001), Currency Returns, Institutional Investor Flows and Exchange-Rate Fundamentals. State Street Associates Working Paper # 5.
- [37] FXAll (2005), FXAll Press Release. <http://www.fxall.com/about/pressrelease58.html>.
- [38] Geithner, T. J. (2004), Hedge Funds and their Implications for the Financial System. Federal Reserve Bank of New York. <http://www.newyorkfed.org/newsevents/speeches/2004/gei041117.html>.
- [39] Greenspan, A. (2004), Euro in Wider Circles. Remarks to the European Banking Congress, Frankfurt. <http://www.federalreserve.gov/boarddocs/speeches/2004/200041119/default.htm>.
- [40] Hersey, B. and J. Minnick (2000), Currency Overlay. Global Pensions. February.
- [41] HSBC (2003), FX Flow Dynamics. April. HSBC.
- [42] Ito, T., R.K. Lyons and M.T. Melvin (1998), Is There Private Information in the FX Market? The Tokyo Experiment. The Journal of Finance, June, 1111-1130.
- [43] Jansen, D.J. and J. de Haan (2003), Statements of ECB Officials and their Effect on the Level and Volatility of the Euro-Dollar Exchange Rate. CESifo Working Paper # 927.
- [44] Kearns, J. and P. Manners (2004), The Profitability of Speculators in Currency Futures Markets. Research Discussion Paper # 2004-07. Reserve Bank of Australia.

- [45] Knott, D. (2002), A Guide to Currency Overlay Management. Deutsche Bank. Global Market Research.
- [46] Kyle, A.S. (1985), Continuous Auctions and Insider Trading. *Econometrica*, Vol. 53, 6, 1315-1336.
- [47] Layard-Liesching, R. (2002), Adding Value Through Currency Management. In A Guide to Currency Overlay Management, D. Knott (ed.) Deutsche Bank. Global Market Research.
- [48] Lewis, K.K. (1994), Puzzles in International Finance Markets. NBER Working Paper #4951.
- [49] Lyons, R. K. (1995), Tests of Microstructural Hypotheses in the Foreign Exchange Market. *Journal of Financial Economics*, 39, 321-351.
- [50] Lyons, R.K. (2001), The Microstructure Approach to Exchange Rates. MIT Press.
- [51] Lyons, R.K. (2002a), Foreign Exchange: Macro Puzzles, Micro Tools. Federal Reserve Bank of San Francisco Annual Review.
- [52] Lyons, R.K. (2002b), Theoretical Perspective on Euro Liquidity. *Economic Policy*. Vol. 17, 35, 572-597.
- [53] Lyons, R.K. (2003), Explaining and Forecasting Exchange Rates with Order Flows. 3RD Bank of France Economic Forum. http://www.banque-france.fr/gb/fondatio/telechar/R_k_lyons.pdf.
- [54] Morris, D.M. (2004), Presentation to Currency Roundtable. Putnam Investments mimeo.
- [55] NYSE (2005), New York Stock Exchange. <http://www.nyse.com/home.html>.
- [56] Osler, C.L. (2002), Stop-Loss Orders and Price Cascades in Currency Markets. Federal Reserve Bank of New York Staff Report # 150.
- [57] Osler, C.L. (2003), Currency Orders and Exchange Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis. *Journal of Finance*, 58, 5, pp.1791-1820.
- [58] Parker Global (2004), Parker FX Index. www.parkerglobal.com.
- [59] Portes, R. (2002), Comment on Hau, H., W. Killeen and M. Moore, How has the Euro Changed the Foreign Exchange Market? *Economic Policy*, vol. 17, 34, 149-191.
- [60] Rime, D. (2001), Private or Public Information in Foreign Exchange Markets? An Empirical Analysis. www.sifr.org.

- [61] Rogoff, K. (2002), Discussion, in Theoretical Perspective on Euro Liquidity, R. K. Lyons, Economic Policy, Vol. 17, 35, 572-597.
- [62] Sager, M.J. and M.P. Taylor (2005), Order Flow and Exchange Rate Forecasting. University of Warwick mimeo.
- [63] Sarno, L. and M. P. Taylor (2002), The Economics of Exchange Rates. Cambridge: Cambridge University Press.
- [64] Taylor, M.P. (1995), The Economics of Exchange Rates. Journal of Economic Literature, 33, pp.13-47.
- [65] Taylor, M.P. (1988), What do Investment Managers Know? An Empirical Study of Practitioner's Predictions. Economica, 55, 218, 185-202.
- [66] Taylor, M.P. and H. Allen (1992), The Use of Technical Analysis in the Foreign Exchange Market. Journal of International Money and Finance, 11, 304-314.
- [67] Taylor, M.P., D.A. Peel and L. Sarno (2001), Non-Linear Mean Reversion in Real Exchange Rates: Towards a Solution to the Purchasing Power Parity Puzzles. International Economic Review, 42, 1015-42.