Économie internationale 96 (2003), p. 107-127.

ÉCONOMIE INTERNATIONALE

ECONOMIC POLICY FORUM

EXPLAINING AND FORECASTING EXCHANGE RATES WITH ORDER FLOWS

Richard K. Lyons[1]

ABSTRACT. This paper summarizes key lessons learned from using models from microstructure finance to explain and forecast exchange rates. The first section is an executive summary, which outlines seven lessons that pertain to how different transaction-flow measures (e.g., interbank flows versus end-user flows) perform in explaining concurrent returns and forecasting future returns. Section 2 addresses three overarching topics, including: (1) how various transaction-flow measures differ, (2) causality between transaction flows and returns and how to think about it, and (3) strategies for pinning down underlying flow drivers. Section 3 addresses empirical results underlying the seven lessons in section 1.

JEL Classification: F31; G14; G15.

Keywords: Exchange rates; Order flow; Price determination; Forecasting; Microstructure.

RÉSUMÉ. Cet article présente les principaux enseignements tirés de l'application des modèles de la microstructure financière à l'analyse de la détermination et de la prévision des taux de change. Le résumé introductif met en évidence sept leçons qui montrent que des mesures différentes des flux de transaction (c'est-à-dire les flux inter-bancaires *versus* les flux entre utilisateurs finaux) permettent d'expliquer les rendements courants et de prévoir des rendements futurs. La deuxième partie traite de trois sujets fondamentaux : la manière dont les mesures des flux de transaction diffèrent ; les relations de causalité entre les flux de transaction et les rendements, et quelles conclusions en tirer ; enfin, les stratégies permettant d'identifier les déterminants sous-jacents des flux. L'article conclut par une présentation de résultats empiriques qui étayent les sept leçons présentées dans le résumé introductif.

Classification *JEL* : F31 ; G14 ; G15.

Mots-clefs : taux de change ; flux d'ordre ; fixation des prix ; prévision ; microstructure.

1. Richard K. LYONS, Professor, UC Berkeley, Haas Shool of Business (lyons@haas.berkeley.edu).

A new approach has emerged within exchange rate economics. Its foundations lay in microeconomics, in particular, the economics of information. Unlike macro models of exchange rates, where relevant information is symmetric economy-wide (or, sometimes, asymmetrically assigned to a central bank), this "New Micro" approach focuses on information that is dispersed. Examples include the micro-level information used to construct aggregate variables like output, money demand, goods prices, and risk preferences, all of which are important for exchange rates in existing models. That private transaction flows might be correlated with micro-level information, thereby conveying this information to the market, is not considered under the traditional macro approach. As an empirical matter, countless bits of information need to be considered in determining exchange rates. Understanding the nature of this information and the role that transaction flows play in transmitting it to exchange rates is the essence of the New Micro approach.[2]

To understand why this shift in perspective toward dispersed information is truly a qualitative shift, consider the following idea, one of the deepest in financial economics: transaction quantities play two distinct roles – they clear markets and they convey information (Grossman 1976; Kyle 1985; Glosten and Milgrom, 1985). Only the first of these roles is reflected in the "more buyers than sellers" view of why flows affect prices. The second is the more insightful (and too often missed). It arises in contexts where information is dispersed because transaction flows affect people's expectations (about future fundamentals and prices). That this second role is empirically relevant to foreign exchange is clear from, for example: (1) findings that transactions have different effects on price, dollar for dollar, depending on the institution type behind them and (2) findings that transaction flows in one currency market have price effects in other currency markets, despite not occurring in those other markets. In addition, not only is transaction flow not "just demand" in the sense above that it plays more than the first role, it is also not just demand in that it reflects *transactions*. Shifts in demand that move prices need not involve any transactions whatsoever. For example, in traditional macro models of exchange rates, demand shifts move price before any transactions can occur at the old price, and at the new price people are indifferent again between buying and selling. In these models, then, there is no relation whatsoever between transaction flows and price movements.

This presentation summarizes key lessons from the New Micro approach with respect to explaining and forecasting exchange rates.[3] Section 1 provides an executive summary. Section 2 addresses three overarching topics that regularly receive attention in this literature. These include: (1) differences between various transaction flow measures, (2) how to think about causality between transaction flows and returns, and (3) strategies for pinning down underlying flow drivers. Section 3 addresses main findings.

2. A clearing-house for research by over 50 authors in this area is available at faculty.haas.berkeley.edu/lyons. That website also provides other resources in New Micro exchange rate economics (e.g., data, slides for a doctoral course in New Micro, etc.). New Micro is based largely on models and techniques developed in microstructure finance, but departs from that field in that it focuses on questions of a more macroeconomic nature.

3. Most of the work I have done in this area is joint work with Martin Evans (Georgetown University).

EXECUTIVE SUMMARY

- Interbank transaction flows have substantial power to explain concurrent 1-day returns.
  - R^2 statistics using direct interbank flows for a host of major currencies in 60-80 percent range (Evans and Lyons 2002c).
  - Though daily frequency, these results apply to long-horizon exchange rates because nominal rates at the daily frequency and lower are (very nearly) random walks, so each increment is an increment to the long-horizon level.
- End-user flows have substantial power to explain concurrent 1-month returns.
  - Long samples for this type of flow is what permits analysis at the monthly frequency.
  - Exploiting the different information in different end-user segments boosts power considerably.
- End-user flows have substantial power to forecast subsequent 1-month returns.
  - Exploiting different information in different end-user segments boosts this power as well.
- Aggregating end-user flows over time increases their power to explain concurrent returns.
  - This is consistent with forecasting power, i.e., correlations that are non-concurrent.
- Different flow segments affect the exchange rate at different horizons.
  - Longer horizon price moves appear more closely tied to non-financial corporate flows.
  - Abrupt short-horizon price moves are more closely tied to financial institution flows.
- Flows are strongly linked to traditional macro anchors of currency value.
  - Flows forecast future macro variables.
  - 30-40% of the variance of segment flows (monthly) is accounted for by macro news arrival.
- Flows' forecasting power is not simply flow forecasting future flow.
  - Two thirds of forecasting power relates to subsequent price movements that are unrelated to subsequent flow.

THREE OVERARCHING TOPICS

Distinguishing different flow measures

The first of the overarching topics is an important one: what distinguishes the different flow measures used in the literature? There are basically three types: (1) interbank flow data, (2) futures flow data, and (3) spot end-user data.[4] We begin by discussing interbank data since that is the data type used most frequently in existing academic work. An important strength

4. For this discussion I am treating the two types of interbank flow—direct and brokered transactions—as a single "interbank" category. (Direct interbank transactions are bilateral, whereas brokered interbank transactions are multilateral in the sense that they are executed within multilateral quoting systems like EBS). For explaining concurrent returns, direct flow data are more powerful. (Compare, for example, the direct-flow results of Evans and Lyons, 2002c, with the brokered-flow results of Danielsson, Payne, and Luo, 2002. This same conclusion that direct flow data convey more information is reached via joint analysis of both data types by Bjonnes and Rime, 2000.

of interbank data relative to futures or end-user data is that they often cover a substantial fraction of total trading in the market (e.g., data from the Reuters 2000-1 system or from EBS). Shortcomings relative to those other data types include the following. First, the interbank data samples are much shorter (one year or less), which makes it hard to address longer-range price impact and forecasting. Second, the interbank data are not available (at least not to date) on a disaggregated basis, so one cannot address how the underlying segment structure affects concurrent price impact and forecasting. (We know from end-user flows that the underlying segment structure is quite important in this respect, as outlined below). Third, and perhaps most substantively, the time-series properties of interbank data are quite different from end-user data, most notably for forecasting purposes: daily interbank flows show no persistence, meaning that today's net interbank flow is uncorrelated with tomorrow's net flow (whereas end-user flows are positively correlated over time).[5] Consistent with this, interbank flows are also uncorrelated with tomorrow's currency return. Thus, though the interbank flows exhibit quite high concurrent correlations with currency returns, they are unlikely to be useful for forecasting at daily frequencies and lower.

Futures flow data is the second main flow measure. Advantages of the futures data include: they are available publicly (with delays), they are available in long samples (> 10 years), and they are rather highly correlated with currency returns in concurrent data. Shortcomings of futures data include the following. First, futures data are not publicly available at the same level of disaggregation as end-user flows (and, again, the disaggregated data are more powerful). The segment breakdown in futures data is limited to two basic categories, "hedgers" and "speculators". Second, not only do historical futures data have lower leading correlations with currency returns than end-user flows, we can expect that their public availability will soon eliminate any forecasting power through competition effects, if this has not already occurred. Third, futures market participants are not representative of the full universe of currency market end-users. For example, futures transactions do not include the subset of currency demands requiring immediate delivery.

The third competing data type is that from end-users. Sources include commercial and custodial banks. These data sets have the advantages of covering long sample periods (> 10 years) and many currencies (permitting cross-sectional as well as time-series analysis). Moreover, these datasets can typically be disaggregated. Commercial bank datasets generally have strong coverage of all end-user segments, both financial and non-financial. Which types of end-user flows convey the most information is, of course, an empirical question, one that we address below.

5. Why these two data types differ so clearly in this respect is not well understood. Part of the answer surely lies in spot marketmakers taking positions that are typically intraday only, whereas the end-user positions are typically maintained much longer (and are adjusted more gradually).

Causality between flow and returns

Though causality between flow and returns remains an open issue, there is an emerging consensus that exchange rate movements are driven predominantly by flow, at least in the major markets. (Note that this is not the same as saying that causality goes only one way; more below). That consensus is built from two basic facts. First, flow effects are empirically present and substantial: nobody serious in this field is arguing that flow has no causal effect on price, with causality running wholly in reverse. Of course, this first fact is not, in itself, enough to establish that price movements are driven predominantly by flow. It must also be true that non-flow-driven price movements are small relative to flow-driven movements. Enter the second fact underlying the consensus: the alternative to flow-driven prices – namely direct price effects from public information arrival – empirically accounts for a very small proportion of total exchange rate variance (less than 5 percent). In the end, it is only in the following more narrow sense that causality remains an open issue: the current debate is about the degree to which causality runs in reverse – from returns to flows – not whether it runs wholly in reverse. This is an important question: if there is also substantial reverse causality, then the coefficients estimated using standard procedures (e.g., OLS) are biased measures of the size of flow effects on price.

From my own experience with the causality issue, it is best to first open conceptual room for understanding why flows *should* cause price changes, and then address why alternatives to this flow-to-price causality are often faulty. That is the approach taken here. The theory of trading is quite helpful for opening conceptual room for why flows cause price changes. Models of trading show that flow causes price adjustment because it conveys information.[6] What kind of information could currency flows possibly convey? There are two basic kinds: (1) risk-neutral fundamentals and (2) risk fundamentals.[7] The term "risk-neutral fundamentals" refers to macro variables that would drive exchange rates even in a risk-neutral world (i.e., a world where the risk of taking on currency positions has no effect on prices). These variables are money supplies, GDPs, interest rates, and inflation – the variables at the center of macro exchange rate models.

The category of "risk fundamentals" is more subtle. It includes variables that take effect as one moves from a risk-neutral world to a risk-averse one in which risk affects position-taking and prices. Examples of risk fundamentals include risk preferences, hedging demands, and

---

6. This information perspective makes it clear that flow is a *proximate* cause of price movements, not an underlying cause. That is, dispersed information first drives flow, then once flow is observed it drives price (i.e., the chain of causality is one-directional, running from information to flow to price).

7. These two categories correspond to what I call "payoff information" and "discount rate information" in my book, Lyons (2001a), and what researchers in asset pricing (e.g., Campbell and Shiller, 1988) call "cash flows" and "expected returns." For people specialized in exchange rates, these other terms – which come from finance – are less familiar. Also, the terms "risk-neutral fundamentals" and "risk fundamentals" emphasize the fundamental nature of both categories as drivers of prices in rational markets. This contrasts with traditional use of the term fundamental in academic research as referring only to the risk-neutral category. Defining fundamentals in this narrow way produces a tendency to define variables not included under this definition as "non fundamental," even if these variables are perfectly sensible drivers of asset prices.

aggregate liquidity demands, all of which vary over time in ways that are not publicly observable. Even when independent of traditional macro fundamentals, these variables drive trades that can and should affect prices: trades motivated by shifts in risk fundamentals have to be absorbed by willing counterparties, and if these counterparties are risk averse then price changes in their favor is what makes them willing. The resulting price changes persist and can even be permanent (see, e.g., the model in Evans and Lyons 2002a).[8] That flows actually do affect exchange rates via this risk-fundamentals channel is obvious to most practitioners.

In contrast, most people are skeptical about the empirical relevance of flow conveying information about risk-neutral (macro) fundamentals. This skepticism arises in large part because inside information about macro fundamentals in the hands of one or a few traders is rather unlikely, at least in major markets. However, this skeptical view is being revised in the face of recent research (reviewed later) that shifts attention from inside information in the hands of a few to small bits of information dispersed throughout the whole economy. Consider some specific examples, variables like the employment report, the trade balance, and the CPI index. Official announcements of these variables occur much later than the underlying economic activity that they measure. In the meantime, that same underlying activity is affecting currency transactions (e.g., an import generating an FX trade). Importantly, it is not necessary for end-users to view themselves as trading strategically on these bits of macro-relevant information. What is necessary is that their FX trades be correlated with these bits. At higher levels of aggregation, like that in end-user flows, the macro information conveyed by these trades becomes significant. In contrast, the traditional view on how exchange rates impound this dispersed macro information is that the bits are first aggregated by official institutions, then a summary is conveyed to the whole market simultaneously in the form of a macro announcement. But as noted, research shows that public macro announcements account for disappointingly little exchange rate variation in total (despite being important within brief time windows). The modern view – backed by recent results using end-user flows – is that much macro information is in fact aggregated by markets rather than official institutions, with flow playing a central role. Under this modern view, the impounding of macro information in price is a continuous process of dispersed information being "symmetrized" rather than a lumpy process of announcement arrivals.

We move now from opening conceptual room for flow-to-price causality to considering some alternatives to flow-to-price causality. The principal alternative is that price changes cause flows, so-called reverse causality. Surely some reverse causality is present. Nevertheless,

8. Note too that even considered in isolation, this "risk fundamentals" channel for price impact does not reduce to the "more buyers than sellers" view of the FX market (that latter view being that a given sized trade has the same price impact regardless of source because it simply reflects demand, not information). The reason is that individuals' trades are not unrelated over time, as is clear from end-user flows data. For example, if a hedge fund takes a position, but the market knows that the position is likely to be reversed soon, then the trade would have less price impact through this risk-fundamentals channel than the trade of a real-money investor that is unlikely to be reversed soon (and might be correlated with future trades in the same direction).

often the evidence offered to support reverse causality is misinterpreted. For example, some have argued that when price movements precede flows this implies reverse causality (typically called feedback trading). Not necessarily true. Here is a counter-example. Suppose flows are positively correlated over time (net purchases followed by net purchases, on average, and vice versa) and are not in any way caused by price. (We know that major-bank end-user flows are positively correlated over time; not being caused by price is simply an assumption to make the example clear.) Now remember that exchange rates are forward-looking asset prices, and so effects from anticipatable flows should be fully discounted immediately. So, a positive shock to dollar buying today will increase the value of the dollar today and will be followed by further positive flow into dollars (i.e., price is forecasting these flows). Though the price increase occurs before the subsequent dollar purchases, there is no positive feedback trading: by construction causality is running wholly from flow to price. Bottom line: because the exchange rate is a forward-looking asset price, one cannot conclude that prices are causing flow just because prices sometimes precede flow.

In addition, reverse causality is often asserted without supporting empirical evidence. For example, some people argue that trading strategies like "the trend is your friend" or "momentum" produce positive feedback that can explain the strong positive correlation found in aggregate data between flows and price. But it is not enough for *some* people to follow these strategies; it must be true that the market on average follows these strategies (on average because we are speaking of aggregate flows here). This is a higher hurdle. In fact, empirical work on this topic in foreign exchange finds the opposite result: the direction of feedback trading in aggregate interbank data is, if anything, negative (Evans and Lyons, 2002b). The empirical case that feedback from price to flow is causing more than a small share of the positive correlation is, at present, not compelling.

The last main alternative to causality running from flow to price is that there is no causality between flow and price and that a third factor is driving both. Consider the following seemingly plausible example. Suppose public macro news arrives that is good for the dollar. Such news would cause the dollar to appreciate and at the same time cause positive flow into the dollar, generating the positive correlation without any causal link. But this view is flawed: it is inconsistent with rational markets (or, in the jargon, inconsistent with "rational expectations"). Information that is truly public, in the traditional sense that everyone observes it simultaneously and interprets it the same way, should induce price adjustment that is instantaneous, i.e., before any transactions can occur at old prices. At the new price, which impounds all the new information, the incentive to buy is no greater than the incentive to sell, so on average positive transaction flow into dollars does not result from the good dollar news (i.e., transaction flow does not systematically correlate with the price adjustment). Now, if one allows instead for public news whose exchange rate implications are not agreed upon by all – a realistic departure from traditional specifications – then the causal link running from flow to price is restored: flow is how price-setting marketmakers learn about the resulting market interpretations. (More precisely, marketmakers learn from flow insofar as their own observation of the public news does not inform them fully about all participants' views). Flow data provide some nice lessons along these lines (reviewed in the subsection below on Flow Links to Macro Fundamentals).

What drives flows? Empirical strategies

Empirical work is concerned with both the causes and consequences of flows. The consequences are addressed later. At this stage we will outline strategies for analyzing flow's causes (thereby helping to frame some of the empirical results below). What specific information is driving flow? The emerging literature offers many approaches to this question (though thus far implementation is based primarily on interbank data; see Lyons 2001b). Those include focusing on: (1) Macro news: does the arrival of public information induce informative flow? (2) Other market states: does intervention induce private flow? Does volatility induce flow? (3) Source: does the type of agent behind the flow matter, and if so, what types of information might specific agent-types have? (4) Cross market signals: are $/¥ trades informative for the $/¥ rate? Does the pattern of cross market effects identify the nature of the underlying information (e.g., is it $ specific)? And (5) Expectation proxies: might flow provide a signal of heterogeneous time-varying expectations regarding future macro variables? And if so, can flow forecast those macro variables? All these strategies are being pursued by researchers in this field. As these results emerge we should have an increasingly clear picture of the underlying sources of exchange rate shocks.

Main findings on specific topics

This section is divided into six topic overviews. Those topics address the following: (A) flows' ability to explain concurrent returns, (B) flow links to macro fundamentals, (C) flows' ability to forecast returns, (D) the sustainability of flow-based trading advantages, (E) effects of segment flows at different time horizons, and (F) forecasting scenarios for the future of FX market structure.

Flows' ability to explain concurrent returns

In the late 1990s, the empirical literature using FX flows shifted from estimating structural models of individual dealer behavior (e.g., Lyons, 1995) toward estimating price determination at the market-wide level. The data used for this work on market price determination were interbank flows (e.g., Payne, 1999; Rime, 2000; Evans and Lyons, 2002a,c; data for replicating the 60-80 percent R^2 statistics from regressions of daily exchange rates on concurrent interbank flows are available from my website[9]). More recently, research in this area has turned to end-user flows and whether this type of flow has comparable explanatory power.

The remainder of this subsection focuses on this recent work on end-user flows. Specifically, to what extent do end-user flows account for exchange rate movements and how is this

9. http://faculty.haas.berkeley.edu/lyons/.

affected by extending the horizon? The issue of extending the horizon is important because it bears on critical questions like the persistence of price effects. For addressing longer-horizon effects, end-user flows provide a critical advantage over extant interbank data.

TABLE 1 presents regression results for the relation between realized excess currency returns in $/€ and concurrent end-user flow (from Evans and Lyons, 2003b; Citibank data from 1993-2000; see also the published tables in my book, *The Microstructure Approach to Exchange Rates*, MIT Press, 2001, pages 254-255). The regressions include flows from all 6 of the main end-user segments: "Corporations" refers to non-financial corporations, "Traders" refers to hedge funds and other leveraged traders, and "Investors" refers to unleveraged asset managers such as mutual funds and pension funds (these three are broken into US and Non-US parts). (TABLE 1 reports only those end-user segments that were significant at that 1 percent level for all three horizons.) Two points emerge. First, the price impact of trades (dollar for dollar) from different sources is quite different, implying different information content. Second, the explanatory power of flows for concurrent returns rises substantially with horizon: at the daily horizon the regression with all 6 flow segments produces an R^2 of only 8 percent, whereas the same regression at the 1-month horizon produces an R^2 of 30 percent.[10] Using end-user flow data from a custodial bank (State Street), Froot and Ramadorai (2002) also find higher correlations between end-user flows and returns as the horizon is extended to 1 month; their flow measure is institutional investors, which consolidates the "Traders" and "Investors" segments used here.

Table 1 - Do different end-user flows have different information potency, dollar for dollar?

	Corporations: Non-US	Investors: Non-US	R^2
1 day (k = 1)	–0.214**	0.353**	0.08
	(0.064)	(0.059)	
2 weeks (k = 10)	–0.390**	0.493**	0.25
	(0.092)	(0.107)	
1 month (k = 20)	–0.376**	0.583**	0.30
	(0.102)	(0.130)	

Notes: The table reports OLS estimates at the daily frequency of the β_j coefficients in the regression:

$$er_{t,k} = \beta + \sum_j \beta_j \Delta^k x_t^j + \varepsilon_t ,$$

where $er_{t,k}$ is the log excess return on the euro between t–k and t, defined as $p_t - p_{t-k} + i^*_{t-k,k} - i_{t-k,k}$, where p_t is the log $/€ rate at time t and $i^*_{t-k,k} - i_{t-k,k}$ is the k-period nominal interest rate differential at time t–k (the non-dollar inter est rate denoted with "*" is the euro after January 1, 1999, and the DM before January 1, 1999). The variable $\Delta^k x_t^j$ is the flow between t–k and t for end-user flow j. (Only those end-user segments that are significant at the 1 percent level for all three horizons are reported.) Standard errors in parentheses are asymptotic and corrected for heteroskedasticity. For the 2-week and 1-month results, standard errors are also corrected for the induced MA(k–1) process in ε_t from overlapping observations. "**" denotes significance at the 1% level.

10. Though results for the 1-day horizon may seem irrelevant at long horizons, recall that major floating rates at daily and lower frequencies are very nearly a random walk. Every change in the level of a random walk is a *permanent* change (i.e., is relevant for that process' level at the infinite horizon).

Among the many questions that the table might motivate, four are particularly interesting. First, why does explanatory power rise with horizon? This is indicative of non-concurrent correlation between flows and returns, so that when data are time-aggregated, more of this correlation shows up as concurrent. I document non-concurrent correlation in the forecasting analysis below. Second, why should trades' information content depend on the transacting party's location? The discussion above regarding risk-neutral fundamentals versus risk fundamentals provides some insight into this question. Think, for example, of the idea that flow helps aggregate dispersed information about aggregate real output or the aggregate risk preferences of financial institutions. We might expect that variations in these variables are in many instances country specific. Thus, the trades originating in a particular country/region are likely to be more highly correlated with changes in these variables from that same country/region, implying different information content than trades from other countries/regions. Third, why are some of these coefficients negative, which seems to have the weird implication that the purchase of euros is bad news for the euro?[11] This is less weird than it might seem. To understand why, consider the negative coefficient from a forecasting perspective. When the market at large learns from a given bank's interbank trades that that bank has received end-user orders, the market cannot separate underlying segment identities. In contrast, an econometrician using the underlying bank data can separate the identities. *Ex post*, then, the initial price response in the market to that bank's interbank trades is too large for some underlying segments and too small for others. When the initial response is too large, the relation of that segment's flows to subsequent returns will be negative, as the market corrects its unavoidable mistake. (Note that this explanation accounts for negative forecasting coefficients, not negative concurrent coefficients.) The fourth interesting question that arises is whether end-user flows convey information beyond that in interbank flows? The answer is yes: when both types of flow are included in a regression of daily excess returns, the null that the coefficients on end-user flows are zero is rejected at the 1 percent level (based on the 4-month sample of daily Reuters 2000-1 interbank flows used in Evans and Lyons, 2002a).

Flow links to macro fundamentals

As noted, evidence from end-user flows indicates that much macro information (i.e., risk-neutral fundamentals) is in fact aggregated by markets rather than official institutions, with flow playing a central role. What exactly is this evidence? There are three different types: (1) evidence that end-user flows forecast future macro variables, (2) evidence that the exchange rate, too, forecasts future macro variables, and (3) evidence that macro news arrivals account for a significant fraction of the monthly variance in segment flows. I address each of these in turn.

If indeed dispersed information about future fundamentals is being aggregated through flow, then flow should help forecast those fundamentals. To address this one can test whe-

11. It is not the case that the negative covariances are "conditional" in that they arise only because all 6 segment flows are included in the same regression: the same pattern of coefficient signs arises in univariate regressions.

ther flows Granger cause macro variables. Consider a set of variables considered fundamental across a wide range of macro modeling traditions: money growth, output growth, and inflation. Evans and Lyons (2003b) show that flows have more forecasting power than the exchange rate itself, a rather striking result. Flows Granger cause US money growth, US output growth and German inflation, whereas the exchange rate Granger causes only US output growth. (That only inflation is Granger caused for the German variables is perhaps less surprising, given the inflation focus of the Bundesbank over the bulk of the 1993-2000 sample.) There does indeed appear to be information in end-user flows useful for forecasting future fundamentals. This result highlights the qualitatively new issues that end-user data allow one to address. Note too that because the market at large does not have access to these disaggregated end-user flows, the finding that flows are an even more powerful predictor than the exchange rate is not a violation of any simple efficiency criterion.

We turn now to a second approach for examining whether flow conveys macro information. The basic idea of this approach is that if flow is conveying this information, then the information should be impounded in price. Accordingly, the exchange rate should forecast future macro variables. Effectively, this is what the Granger causality tests address, though they do so in a rather particular way. Our second approach does so in a completely different (complementary) way. The starting point is an influential paper by Mark (AER, 1995). Mark estimates regressions of the following form:

$$p_{t+k} - p_t = \alpha + \beta\,(f_t - p_t) + \varepsilon_{t+k}$$

where $p_{t+k} - p_t$ is the log change in the spot rate over k periods and f_t denotes fundamentals measured at time t (measured from money supply and GDP data – see paper for details). The idea is that exchange rates may be "catching up" to fundamentals over time, i.e., if price is below fundamental value today ($f_t > p_t$), then on average it should rise over time – a positive coefficient β. (This is essentially an error-correction model.) Mark shows that as the horizon is lengthened, the fit of the regression gets better and better. At horizons in the 3-5 year range, the model fits the data quite well, suggesting that ultimately fundamentals are getting into exchange rates, but it takes considerable time.

An alternative approach turns the basic logic of the Mark regression on its head. The underlying idea motivating the Mark regression is that exchange rates do a poor job of tracking fundamentals over shorter horizons. Under this view the FX market is not working particularly well in impounding macro information. The view that order flow is important in aggregating macro information is in some sense quite the opposite: order flow is conveying information to the market today that is relevant to the paths of macro variables in the future. Thus, under this view, order flow acts to telescope the market's forward-looking expectations and impound changes in those expectations into today's price. This is a view that markets are instead working quite well, and that the problem with the earlier analysis in the literature is that the variable f_t (money supply and real output, but only their current values) is a poor measure of the expected future fundamentals that the market is discounting in price.

This opposite view to that underlying the work of Mark can be captured by the following variation on his regression:

$$f_{t+k} - f_t = \alpha + \beta\,(p_t - f_t) + \varepsilon_{t+k}$$

The idea now is that if price today is forecasting strong future fundamentals ($p_t > f_t$), then those fundamental macro variables should tend to strengthen – a positive coefficient β. What one finds, in a nutshell, is that the supporting evidence from this version of the equation is as least as strong as that found by Mark. This suggests that exchange rates are indeed impounding forward-looking information as one would expect of an asset market.

I do not report those results here for three reasons. First, the lengthening of the horizon k together with overlapping observations makes the test statistics for evaluating the significance of β and the regression fit tenuous. Indeed, several papers were published in the years following the Mark (1995) paper that were dedicated to showing how tenuous inference really is in this particular setting. Second, a paper by Engel and West (2002) effectively addresses the same question. (They also found that exchange rates have some power to forecast macro variables, but still not so much as theory would predict.) Third, remember that this approach makes no use of end-user flows whatsoever. We embarked on this path because at the time we believed that flows could not possibly be forecasting macro aggregates if exchange rates were not forecasting macro aggregates. But I no longer believe this: there is no good reason why all the information in end-user flows is necessarily impounded in price, so that price makes end-user flows redundant as a forecasting variable. Indeed, the forecasting results below (forecasting returns using flows) imply that the current exchange rate does not already impound all the information in end-user flows. Given this, it is not a stretch to believe that current exchange rates do not already impound all the information in end-user flows that is specifically macro related.

The third type of evidence we produced regarding links between end-user flows and macro information is based more specifically on the arrival of macro news. This analysis borrows heavily from the methodology introduced in Evans and Lyons (2003a). The basic idea is the following. We can measure daily rates of macro news flow by measuring the macro news items that arrive each day over the Reuters Money Market Headline News screen. (The median number of arrivals per day over the end-user flows sample relating to US or German macroeconomics is 11, with considerable variation around this median.) We then ask the question: is the variance of segment flows higher on days with a lot of macro news flow? The answer to this question is very definitely yes, and it is yes for all 6 of the segments introduced in TABLE 1. We then model flow determination in a more structured way to determine what fraction of flow volatility is accounted for by macro news arrival (in the jargon, we do a variance decomposition). The answer is that in the $/€ market (the only one for which the analysis has been done thus far), macro news arrival accounts for 30-40 percent of each of the six segment flow variances (monthly frequency). Though not direct evidence, in itself, that end-user flow segments convey macro information, these results indicate that end-user flows are linked to the process of macro information flow.

Flows' ability to forecast returns

Forecasting asset prices is hard. Indeed, in efficient markets it should be hard: unless one has an informational advantage of some kind it shouldn't be possible in an efficient market (by which I mean a market where price reflects all public information). Foreign exchange is no exception. The informational advantage we are considering here is explicit: information on end-user currency flows that are not available to the market at large (and not available to any central bank either). *A priori*, efficient markets logic is not enough in this instance to presume that these data have no forecasting power. Put differently, the question is an empirical one, since theory provides no strong guidance. Let me add that I would not say the same thing about flow data from currency futures. Futures flow data are provided by the exchanges to the public. Because they are available to the market at large, in an efficient market they should not be valuable for forecasting. Accordingly, the argument that futures flow data have no forecasting power, and therefore the data being considered here shouldn't either, is not sound.

Let us shift from the explanatory regressions in TABLE 1 (concurrent flow and return data) to the corresponding predictive regression (see TABLE 2 notes, next page, for regression model). The forecasting results in Panel A of TABLE 2 are really quite striking: the coefficients on US corporate and investor flows are highly statistically significant and the R^2 statistic is 19 percent.[12] (Only coefficients with forecasting power at the 1 percent level are reported, the same threshold used in TABLE 1.) For comparison, the big forecasting puzzle in the empirical literature is the so-called forward bias puzzle, the well documented result that monthly forecasting using the beginning-of-period forward discount produces significant coefficients. But despite this being the best forecasting variable identified in academic research, in that case R^2 statistics are generally only 2-4 percent.

Going a step further, recall that an important objective is to clarify the nature of the underlying dispersed information driving flows. For example, it is possible that the forecasting power in end-user flow data arises because flow forecasts risk-neutral fundamentals, i.e., the variables that enter a monetary model with no risk premium (money supplies, outputs, interest rates, price levels). That flows do indeed have some forecasting power in this narrow respect is indicated by the Granger causality results. But it is also possible that flows are instead forecasting risk fundamentals, e.g., determinants of foreign exchange risk premia. The most obvious example is the possibility that flows forecast future flows, which will have concurrent price impact when realized due to simple market absorption effects. (If the flows forecasting future flows are not public information, and end-user flows are not, then there is no market efficiency argument that requires all the price impact to occur with the initial flow realiza-

12. From Evans and Lyons (2003b). Over horizons less than one month, end-user flows have less forecasting power (as measured by R^2). Note that the sample includes all data available to us, so we are unable to study the model's out-of-sample performance. Why the significance of segment flows differs across the concurrent and forecasting regressions is being addressed in subsequent work (e.g., in the concurrent regressions non-US corporations have a stronger link to returns than US corporations, whereas this ordering is reversed in the forecasting regressions).

tion.) This would produce a correlation between current flows and future returns, but it would have nothing to do with forecasting future macro fundamentals.

As a first cut on disentangling these two distinct views on why flows forecast returns (i.e., on the type of information flows convey), consider a decomposition of the excess return regressions. The basic idea is the following. We know that flows forecast future returns. We also know that flows are an important concurrent determinant of returns. If flow is forecasting future flow, then the part of the future returns that flow should forecast is the part that is determined by those future flows. Accordingly, one can regress returns on concurrent flows and use the fitted values from the regression to separate the return series into a flow explained part (the fitted values) and the flow-unrelated part (the regression residual). One can then determine whether flow tends to forecast the flow-explained part of the return or the flow-unrelated part. The former would suggest the "flow forecasting future flow" channel. The latter is consistent with the "flow forecasting narrow fundamentals" channel.

Table 2 - **Forecasting excess returns 1 month ahead with end-user flows**

Panel A: 1-Month forecasting regressions			
Dependent variables	Corporations: US	Investors: US	R^2
er_{t+k}	0.984**	–0.763**	0.19
	(0.259)	(0.242)	
Panel B: Forecast decomposition (same regressors)			
$\hat{E}[er_{t+k} \mid \Delta^k x^j_{t+k}]$	0.202	–0.239**	0.07
	(0.169)	(0.098)	
$\hat{\varepsilon}_{t+k}$	0.710**	–0.490**	0.13
	(0.243)	(0.185)	

Notes: The table reports OLS estimates of the β_j coefficients in regressions of the form

$$y_{t+k} = \beta + \sum_j \beta_j \Delta^k x^j_t + \varepsilon_{t+k},$$

where $\Delta^k x^j_t$ is the lagged flow between $t–k$ and t for end-user segment j. In the Panel A, the dependent variable er_{t+k} is the log excess return on the euro between t and $t+k$, (i.e., $er_{t+k,k}$). Estimates are from daily data with k = 20 (roughly one trading month). (Only those end-user segments that are significant at the 1 percent level are reported). Standard errors are in parentheses, corrected for heteroskedasticity and the MA(k–1) process in ε_t induced by overlapping observations. Panel B addresses which component of the subsequent realized return is being forecasted. We use the estimates from table 1 (with k = 20) to write $er_{t+k,k} = \hat{E}[er_{t+k,k} \mid \Delta^k x^j_{t+k}] + \hat{\varepsilon}_{t+k}$ where $\hat{E}[er_{t+k,k} \mid \Delta^k x^j_{t+k}]$ is the fitted value from the regression of excess returns on concurrent flows, and $\hat{\varepsilon}_{t+k}$ is the regression residual. Panel B reports OLS estimates of the β_j coefficients from the regression with $y_{t+k} = \hat{E}[er_{t+k,k} \mid \Delta^k x^j_{t+k}]$ (upper row) and $y_{t+k} = \hat{\varepsilon}_{t+k}$ (lower row), with standard errors corrected for heteroskedasticity and the induced MA(k–1) process. "**" denotes significance at the 1% level.

The punch-line of the Panel-B results is that flow has power to forecast both components. That said, roughly two-thirds of this forecasting power is unrelated to future flow effects on price (the power to forecast the future regression residual). This is consistent with the Granger causality results: flow appears to be forecasting price components that are impounded in price without any relation to flow, or at least without any relation to Citibank's end-user flow. Whether this last statement generalizes to more encompassing definitions of flow is a question that requires richer flow data sets than are currently available.

The sustainability of flow-based trading advantages

Whether the forecasting power of end-user flows is sustainable is an important question. Though a definitive answer is not possible, by taking a disciplined approach (i.e., using a conceptual framework), we can provide some insights. My conclusion is that forecasting power is more sustainable than most people believe.

There are two key elements for thinking about sustainability: segment structure and information structure. Consider segment structure first. By highlighting segment structure I mean to highlight the fact that the six segments used in the regressions above definitely do not share the same concurrent relation with price, dollar for dollar, nor the same forecast relation with price. This is good news for sustainability. If, in contrast, all flows had statistically indistinguishable links to price, then other players using even small subsets of end-user flows data would essentially have access to the same market signals. This would increase the competition in trading on these signals, and would rather rapidly render future prices independent of current flow information. This is not the case. Some end-user flow datasets do not include non-financial corporate flows, which are quite important (from TABLE 2). Similarly, the fact that US clients appear more valuable for forecasting (both US corporates and US investors) may be bad news for end-user datasets with low penetration among US clients (if this property of the data is robust).

The second key element for sustainability is what I call information structure. With this I am referring back to the organizing framework provided by the separation of risk-neutral fundamentals and risk fundamentals. Now, the correlation of flow signals received by different market participants is clearly important for sustainability. But rather than think about this correlation in a one-dimensional way, it is useful to consider the two dimensions provided by these two distinct types of information. The question then becomes the following: in each of these two separate categories, how correlated is a given set of flow signals with the flow signals of other participants? What this does is open conceptual room for understanding why a given institution's flow data might convey particular types of information well relative to data from other institutions, contributing to sustainability. Think, for example, of the category of risk fundamentals. Effectively, in that category what matters is the total net flow in the market and the effective time horizon over which the position changes underlying that flow are likely to be held. Suppose, in the extreme, that the flow received by each bank is an idiosyncratic realization of end-user demand shocks. In this setting, what matters for forecasting price is the sum of all these orders. And no matter how aggressively other banks capture and exploit their own flow data, the end-user flow received by a given bank cannot be inferred from the trades of others, and therefore cannot be impounded in price until the market gets some sense for them. (This is in fact the information structure in the model of Evans and Lyons 2002.) In this instance, a bank's information advantage would be sustainable (as would the separate information advantages of the other banks). My point is not that the "idiosyncratic flows" assumption is realistic. It is rather that a more disaggregated view of the different components of fundamentals can only enhance one's appreciation for the sustainability of individual banks' flow-based information advantages. The one-dimensional view of signal correlation is too simplistic.

Effects of segment flows at different time horizons

Whether different flow segments have price effects at different time horizons is a fascinating issue on which we have made some progress. Mintao Fan and I provide perspective on this issue in Fan and Lyons (2003). Here I summarize some of those findings.

We find that abrupt, short-horizon price moves (e.g., daily) are more closely tied to the trades of financial institutions. This is perhaps unsurprising, but there are some subtleties nonetheless. For example, we find that on that phenomenal day in the dollar-yen market in October 1998 (price fell from over 130 yen to the dollar to less than 120), it was the financial institutions (at least at Citibank) that were selling the dollar. The subtlety is that is was not the hedge funds who were selling dollars heavily. (The story in the popular press was that hedge funds were all trying to unwind their dollar-yen carry trades that day.) In fact, hedge funds bought dollars, net, that day. We find that the unleveraged ("real-money") investors were the ones selling the dollar. Like hedge funds, non-financial corporates were also net buyers of dollars that day, though their net buying was a good deal smaller than that of hedge funds. Turning to a sample that includes other daily rate moves that were among the largest, we find the same result that financial institutions were much more likely on those days to be trading in the direction of the rate than non-financial corporations. Beyond the October 1998 yen-dollar event, though, it is not systematically true that the real-money investors were more likely than the hedge funds to be the big net traders on those days.

For longer horizon price moves, in contrast, we find that non-financial corporate flows are more likely to be the driver. We offer two pieces of evidence. First, if one looks at the major downward trend in the dollar-euro rate from 1999 to 2002, one finds that only the non-financial corporates track the long trend in the exchange rate in terms of cumulative euro selling. This is not a statistical test, but it certainly does make it difficult to understand how the financial institution flows could account for the cumulative net change in the rate. Second, and more statistically based, note that the only coefficient in the forecasting results (TABLE 2) that is positive is for corporates (US corporates in particular). The financial segments in that table show coefficients that are, on balance, negative, meaning that if they are buying euros this month then the euro tends to depreciate next month. As noted above in the discussion of how to understand negative forecasting coefficients, it may simply be that the market initially over-reacted to these financial-institution flows, and that is why the subsequent effect is so strongly negative. In any event, in the process of time-aggregating the date underlying TABLE 2, we found that as one extends from 1 week, to 2 weeks, to 3, and then 4, the coefficients on corporates remain positive and get progressively larger, whereas the coefficients on investors get progressively more negative. This too is consistent with positive corporate-flow effects on price that take longer to impound in price, but are ultimately stronger than financial-institution flow effects (which appear to decay after about a year; see Froot and Ramadorai, 2002).

These lower frequency relationships are absolutely crucial to most practitioners so let me add a couple more comments. First, that corporate flows should be relatively important at longer horizons makes sense from a theoretical perspective. International macro models put great emphasis on the underlying real economy as a low-frequency driver of exchange rates. These models also highlight the longer-horizon importance of balance-of-payments sustainability, which is intimately linked to the corporate activity that produces the current account. Second, there remains much work to be done with respect to corporate flows, e.g., linking them to data sources for public flow such as those within official balance of payments statistics.

Forecasting scenarios for the future of FX market structure

To understand how banks are likely to fit into the future food chain of the market, it is important to have a sense for how the structure might evolve. In Lyons (Brookings, 2002), I address the market's future using two organizing ideas. The first is that the structure of currency markets is deeply dependent on the management of credit risk. This contrasts with the drivers of market structure focused on by academic work (those being management of market risk, attenuation of asymmetric information, and entry barriers). The second organizing idea is that price variation in spot currency markets is driven primarily by dispersed information. I argue that these two ideas are vital to understanding this market's future and propose three scenarios for that evolution. The scenario I consider most likely is one in which the current dealer structure is maintained through top dealing banks' cross-subsidizing their liquidity provision by using gains generated from their flow information.

Of the two organizing ideas, the first is the more provocative, so let me clarify why I believe that the structure of currency markets is indeed deeply dependent on credit risk. An important factor in considering whether current FX customers would want to participate in an open, electronic auction market (with all participants on equal footing) is whether they would be willing to assume the counterparty credit risk arising from trading bilaterally with other non-bank customers. Banks themselves are careful not to do this without tight management of the attendant exposures. Non-financial institutions, or even financial institutions like pension funds, hedge funds, and mutual funds, are in general quite averse to assuming such risk. True, these institutions are familiar with trading currency futures, for which clearinghouses have been designed to mitigate this risk. But a clearinghouse system of sufficient size to handle the worldwide spot market and at the same time offer flexible access is unlikely to be economically practicable given the different national legal jurisdictions and concomitant legal risk.

Commercial banks, on the other hand, are particularly well suited to play this role. They have developed significant relationships with institutional customers. Indeed, there is a long line of research in corporate finance based on banks' comparative advantage in monitoring the credit-worthiness and performance of their corporate customers. These relationships, together with banks' expertise in managing credit risk, gives them an advantage that is, in my judgment, the key driver of the current bank-intermediated dealership structure.[13] It contributes substantially to bank dealers' ability to compete with the new, auction-based alternatives recently offered to FX customers by new entrants.[14]

In the paper I offer three scenarios for the future of FX market institutions. These three scenarios are not only the most likely (in my judgment), when considered jointly they also highlight the competitive forces at work. The time frame I have in mind is 10 years.

Future 1: Emergence of a centralized auction for customers that will entail an unbundling of liquidity provision and credit-risk management.

Banks are providing (at least) two distinct services when fulfilling their dealership function: liquidity provision (i.e., low transaction costs) and counterparty credit-risk management.[15] But these services need not be bundled. As has been happening across the banking industry for decades, unbundling is a viable alternative. Under the Future-1 scenario, FX customers (and former dealers) would provide one another liquidity within a centralized auction structure. At a scale sufficient to benefit from network effects in liquidity, this structure should have a comparative advantage in providing liquidity. At the same time, all bilateral deals struck within this auction-market structure would then be routed through commercial banks that act as legal counterparties to both sides of the transaction (and are paid for the service). In principle, this arrangement concentrates provision of the separate services where comparative advantage is highest.

Future 2: Continuation of dealer market with banks cross-subsidizing their disadvantage in liquidity provision (disadvantage relative to the centralized auction with full network effects).

Top dealing banks have a vested interest in maintaining the current dealer structure. One way to do so is to keep customer transaction costs at such a level that auction-based alternatives are unable to establish a significant foothold, and are therefore unable to develop the network effects for lower-cost liquidity provision. Indeed, for the major currency markets, this second scenario is already emerging: bid-offer spreads are so thin in the major FX markets that it is difficult to make much money directly from intermediation.

---

13. An important source of this expertise in managing credit risk may be banks' management of their own credit risk portfolios. Some have suggested that this credit management advantage may also stem from government guarantees in the commercial banking sector.

14. Another factor that puts the new auction-based entrants at a disadvantage is the classic one: they lack sufficient scale to benefit from network effects from concentrating liquidity. My own view is that the credit-risk factor is the more primitive: if it were not present then the network externalities would be much more easily achieved.

15. Though less fundamental than these two basic services, there is a third service that is important to many commercial bank customers—"transaction management." This includes advising customers on how to time trades to reduce execution costs and how to reduce the risk of transaction cost overruns. Whether this service contributes substantially to the current dominance of the dealership structure remains to be seen. (After all, this service, too, can be unbundled and priced separately.)

But if banks are making little money, why are they so adamant about hanging on to the liquidity-provision part of the business? Why not just let Future 1 play itself out? Enter the indirect profitability of FX liquidity provision. Large FX banks now recognize an important source of indirect profits from intermediating flows, namely their ability to exploit information contained in those flows. Indeed, one of the factors that is depressing spreads so much for the largest, most coveted customers is the intense bidding for this flow information by dealing banks. The banks are, in a very real sense, investing in flow information by cross-subsidizing their comparative disadvantage in liquidity provision.[16]

Future 3: Emergence of centralized auction due to the opening of IDBs (e.g., EBS) to customers.

Futures 1 and 2 have already been set into motion. Future 1 does not appear to be winning, in part due to electronic initiatives by top banks and in part due to aggressive pricing of liquidity provision (as in Future 2). Nevertheless, if top dealing banks were to sense a Future-1 victory, in spite of their best efforts to encourage Future 2, I do not believe that Future 1 would be allowed to predominate. Almost surely, faced with such a threat the banks that own EBS would choose to open it to customers. For the following reasons, this "open EBS" would likely gather network effects very quickly. The customer relationships are there: the banks that own EBS are the same banks that have customer relationships via their dealing services. The technology is not a major hurdle: one could ape the auction technology that has been developed in the US bond market for secondary-market customer trading. From the EBS perspective, it would be essential to maintain the network liquidity effects in its favor – if the market were going in the direction of centralized customer trading, EBS could not afford to wait.

R. K. L.

References

Bjonnes, G., Rime, D., 2000. Dealer behavior and trading systems in foreign exchange markets, typescript, Bank of Norway (at www.sifr.org/dagfinn.html). *Journal of Financial Economics*, forthcoming.

Campbell, J.Y., Shiller, R., 1988. The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies*, January.

Cohen, B., Shin, H., 2002. Positive feedback trading under stress: Evidence from the US treasury securities market, typescript, London School of Economics, August.

16. It is tempting to argue that this information would still be available to banks under Future 1 because they would still be privy to the trade and its particulars (*via* their ultimate role as legal counterparty). I find this uncompelling for two reasons. First, banks would find it difficult under Future 1 to maintain their existing FX customer bases purely on the basis of pricing the credit-risk management fee (i.e., existing relationship capital would be destroyed). Second, though banks would surely have access under Future 1 to trade details like size and price, they would not have access to other relevant information that they rely on when using order flow information (e.g., whether the transaction was an outright transaction or part of a swap).

Danielsson, J., Payne, R., Luo, J., 2002. Exchange rate determination and inter-market order flow effects, typescript, London School of Economics.

Engel, C., West, K.D., 2002. Exchange Rates and Fundamentals, typescript, University of Wisconsin, Madison, August.

Evans, M.D.D., Lyons, R.K., 2002a. Order flow and exchange rate dynamics, *Journal of Political Economy* 110(1), February, 170-80.

Evans, M.D.D., Lyons, R.K., 2002b. Time-varying liquidity in foreign exchange, *Journal of Monetary Economics* 49(5), July, 1025-51.

Evans, M.D.D., Lyons, R.K., 2002c. Informational integration and FX trading, *Journal of International Money and Finance* 21 (6), November, 807-31.

Evans, M.D.D., Lyons, R.K., 2003a. How is macro news transmitted to exchange rates?, NBER Working Paper 9433, January.

Evans, M.D.D., Lyons, R.K., 2003b. New micro exchange rate economics, typescript, U.C. Berkeley, January.

Fan, M., Lyons, R.K., 2003. Customer trades and extreme events, in Mizen, P. (Ed), *Foreign Exchange, Monetary History, Exchange Rates and Financial Markets: Essays in Honor of Charles Goodhart*, Edward Elgar, 160-179.

Froot, K., Ramadorai, T., 2002. Currency returns, institutional investor flows, and exchange rate fundamentals, NBER Working Paper 9101, August.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *Journal of Financial Economics* 14: 71-100.

Grossman, S. 1976. On the efficiency of competitive stock markets where traders have diverse information, *Journal of Finance* 31, 573-585.

Kyle, A.S., 1985. Continuous auctions and insider trading, *Econometrica* 53 (6), novembre, 1315-35.

Lyons, R.K., 1995. Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321-351.

Lyons, R.K., 2001a. *The Microstructure Approach to Exchange Rates*, MIT Press: Cambridge, MA, 2001.

Lyons, R.K., 2001b. New perspective on FX markets: order-flow analysis, *International Finance* 4 (2), summer 2001, 303-20.

Lyons, R.K., 2002a. Foreign exchange: macro puzzles, micro tools, *Federal Reserve Bank of San Francisco Economic Review*, June 2002, 51-69.

Lyons, R.K., 2002b. The Future of the foreign exchange market, in Litan, R., Herring, R. (Eds.), *The Future of Securities Markets*, Brookings-Wharton Papers on Financial Services, Brookings Institution Press: Washington, DC.

Mark, N.C., 1995. Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review* 85 (1), March, 201-218.

Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies, *Journal of International Economics* 14 (1-2), February, 3-24.

Osler, C., 2002. Currency orders and exchange-rate dynamics: An explanation for the predictive success of technical analysis, *Journal of Finance*, forthcoming.

Payne, R., 1999. Informed trade in spot foreign exchange markets: An empirical investigation, typescript, London School of Economics, January. *Journal of International Economics*, forthcoming.

Rime, D., 2000. Private or public information in foreign exchange markets? An empirical analysis, typescript, Bank of Norway, March.

Tien, D., 2002. Hedging demand and foreign exchange risk premia, typescript, Haas School of Business, U.C. Berkeley, January.