Foreign Exchange: Macro Puzzles, Micro Tools*

Richard K. Lyons

*Professor of Economics and Finance, Haas School of Business, University of California, Berkeley,
National Bureau of Economic Research, and
Visiting Scholar, Federal Reserve Bank of San Francisco*

This paper reviews recent progress in applying information-theoretic tools to long-standing exchange rate puzzles. I begin by distinguishing the traditional public information approach (e.g., monetary models, including new open economy models) from the newer dispersed information approach. (The latter focuses on how information is aggregated in the trading process.) I then review empirical results from the dispersed information approach and relate them to two key puzzles, the determination puzzle and the excess volatility puzzle. The dispersed information approach has made progress on both.

To repeat a central fact of life, there is remarkably little evidence that macroeconomic variables have consistent strong effects on floating exchange rates, except during extraordinary circumstances such as hyperinflations. Such negative findings have led the profession to a certain degree of pessimism vis-à-vis exchange-rate research.

Frankel and Rose (1995, p. 1,709)

1. Introduction

Does the foreign exchange market aggregate information? Surely it does: so many of the variables that drive pricing are dispersed throughout the economy (e.g., individuals' risk preferences, firms' productivities, individuals' money demands, individuals' hedging demands, etc.). Indeed, aggregating dispersed information is one of asset markets' central functions.[1] Yet models of exchange rate determination abstract completely from information aggregation. These models (e.g., monetary models, portfolio balance models, new open economy macro models) posit an information environment in which all relevant information is publicly known. This approach is sensible if the abstraction misses little, i.e., if dispersed information is rapidly summarized in the public macro variables we rely on to estimate our models. Only recently has this common assumption received any attention.

My thesis is that abstracting from information aggregation when analyzing exchange rates misses quite a lot. The argument rests on two main points. First, empirically the public information approach fares poorly (see, e.g., Meese and Rogoff 1983, Frankel, et al. 1996, and the surveys by Frankel and Rose 1995 and Taylor 1995). Meese (1990) describes the explanatory power of these models (for monthly or quarterly exchange rates) as "essentially zero." More recent models within this approach also fare poorly (Bergin 2001). In sum, there is general agreement that the public information approach is deficient; the open question is why.

My second main point is more positive: recent empirical work on exchange rates using what I call the "dispersed information approach" has enjoyed some success. This work relies on micro models of how, specifically, asset markets accomplish information aggregation. When coupled with the poor performance of public information models, these positive results imply that the above assumption—that dispersed information is rapidly summarized in public information—is dubious.

The remainder of this article focuses on positive results from the dispersed information approach and relates them to fundamental exchange rate puzzles. Section 2 provides an overview of order flow as an information aggregator. Section 3 addresses the determination puzzle—why the explanatory power of concurrent macro variables is so low.

*I thank the following for helpful comments: Richard Portes, Michael Melvin, Michael Moore, Helene Rey, and Andrew Rose. I also thank the National Science Foundation for financial assistance and the Federal Reserve Bank of San Francisco for support as a visiting scholar.

1. Nobel laureate Friedrich Hayek (1945) provides an early and powerful articulation of this point: "the problem of rational economic order is determined precisely by the fact that knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus a problem of the utilization of knowledge not given to anyone in its totality" (p. 519).

Section 4 addresses the excess volatility puzzle—why floating rates are more volatile than measured fundamentals predict. Section 5 concludes by providing directions for further research.

2. Order Flow: An Information Aggregator

2.1. Introduction and Definition

When one moves from the public information approach to the dispersed information approach, a variable that plays no role in the former takes center stage: order flow. Order flow is a term from the field of microstructure finance.[2] Understanding it is essential for appreciating how the dispersed information approach departs from the public information approach. Order flow is transaction volume that is *signed* according to whether the transaction is initiated from the buy side (+) or the sell side (–). For example, if you decide to sell a dealer (marketmaker) 10 units (shares, euros, etc.), then transaction volume, which is unsigned, is 10, but order flow is –10.[3] Over time, order flow is measured as the sum of signed buyer-initiated and seller-initiated orders. A negative sum means net selling over the period.

Order flow is a variant of another important term, "excess demand." It is a variant rather than a synonym for two reasons, the first relating to the excess part and the second relating to the demand part. For the former, note that excess demand equals zero in equilibrium by definition—there are two sides to every transaction. This is not true of order flow: in markets organized like foreign exchange (FX), orders are initiated against a marketmaker, who (if properly compensated) stands ready to absorb imbalances between buyers and sellers. These "uninitiated" trades of the marketmaker drive a wedge between the two concepts, excess demand and order flow.[4] The second reason the concepts differ is that order flow is in fact distinct from demand itself. Order flow measures actual transactions, whereas demand shifts need not induce transactions. For example, the demand shifts that move prices in traditional exchange rate models (e.g., monetary models) are caused by the flow of public information, which moves rates without transactions ever needing to occur.

In dispersed information models, information processing has two stages, the second of which depends on order flow. The first stage is the analysis or observation of dispersed fundamentals by nondealer market participants (mutual funds, hedge funds, individuals with special information, etc.). The second stage is the dealer's—i.e., the price setter's—interpretation of the first-stage analysis, which comes from reading the order flow. Dealers set prices on the basis of this reading.

Order flow conveys information about dispersed fundamentals because it contains the trades of those who analyze/observe those fundamentals. It is a transmission mechanism. Naturally, though, these informative trades may be mixed with uninformative trades, making the task of "vote counting" rather complex. In some dispersed information models, the dealer learns nothing about fundamentals that she does not learn from order flow. As a practical matter, this is clearly too strong. The dealer's dependence on learning only from order flow arises in some models because all of the relevant information is dispersed. When information is publicly known, dealers do not need to learn from order flow. In practice, although some information relevant to FX is publicly known, some is not, so learning from order flow can be important. The empirical models I describe in Section 3 admit both possibilities.

Consider such a "hybrid" model from a graphical perspective. The top panel of Figure 1 illustrates the connection between fundamentals and prices under the public information approach. Under this approach, not only is information about fundamentals publicly known, but so, too, is the mapping from that information to the price. Consequently, price adjustment is direct and immediate. The middle panel shows the dispersed information approach. The focus in that case is on fundamental information that is not publicly known. In those models, information is first transformed into order flow. This order flow becomes a signal to the price setter that the price needs to be adjusted. The bottom panel presents the hybrid view. Here, the model accommodates both possibilities: information that affects prices directly and information that

2. Microstructure finance has two main strands: market design and information processing. The dispersed information approach to exchange rates borrows heavily from the second of these strands.

3. Measuring order flow is slightly different when trading takes place via a "limit order book" rather than through dealers. (An example of a limit order is "buy 10 units for me if the market reaches a price of 50.") Limit orders are collected in an electronic "book," and the most competitive of those orders define the best available bid and offer prices. When measuring order flow, limit orders are the passive side of any transaction, just as the quoting dealer is always on the passive side when trading involves dealers. When orders arrive that require immediate execution (e.g., an order to "sell 10 units now at the best available price"), these orders—called market orders—generate the signed order flow.

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

FIGURE 1
APPROACHES TO PRICE

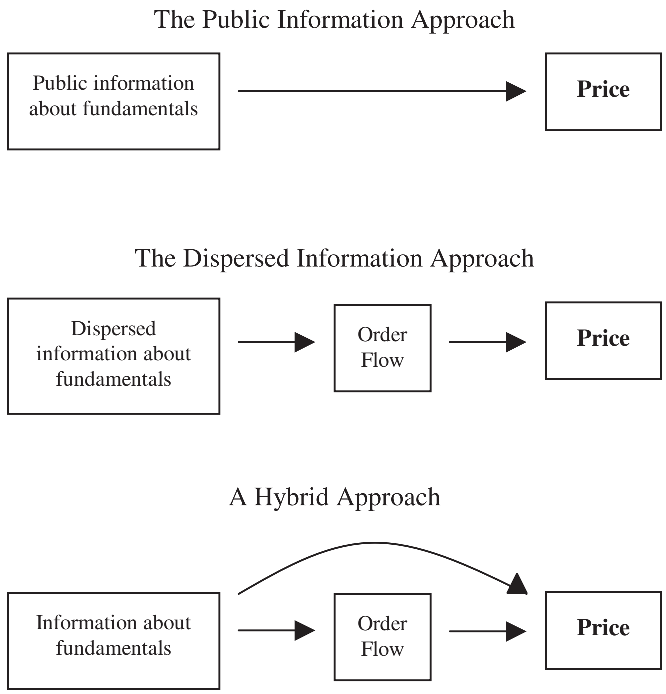

affects prices via order flow. With models that allow for both, the data can determine which possibility accounts for more exchange rate variation.

2.2. Order Flow and Exchange Rates over Long Horizons

Although empirical work in microstructure finance is generally applied to high frequency events, this does not imply that microstructure tools are irrelevant to lower-frequency, resource relevant phenomena. Indeed, there are ample tools within the micro approach for addressing lower-frequency phenomena. And new tools continue to emerge, thanks in part to a recognition within the broader microstructure literature that resource allocation warrants greater attention.

Regarding long-lived effects, the most important point to recognize is that when order flow conveys information, its effect on prices *should* be long-lived. Indeed, a common assumption in empirical work for distinguishing information from pricing errors is that information's effects on prices are permanent, whereas pricing errors are transitory (French and Roll 1986, Hasbrouck 1991). These long-lived effects are borne out in the data, in equity markets, bond markets, and FX markets. In FX, for example, Evans (1997, 2001), Evans and Lyons (2002), Payne (1999), and Rime (2000) show that order flow has significant effects on exchange rates that persist. Indeed, statistically these effects appear to be permanent. Among microstructure's long-lived implications, this "information" channel is definitely the most fundamental.

An analogy may be helpful. The dispersed information approach may speak to longer-horizon exchange rates in much the same way that microscopes speak to pathologies with macro impact. In medicine, microscopes provide resolution at the appropriate level—the level at which the phenomenon emerges. This is true irrespective of whether the phenomenon also has macro impact. Resolution at this level is the key to our understanding. Similarly, tools from the dispersed information approach provide resolution at the level where its "phenomenon" emerges—the level where prices are determined. What information do dealers have available to them, and what are the forces that influence their pricing decisions? (Whether we like it or not, it is a stubborn fact that in the major currency markets, *there is no exchange rate other than the prices these people set.*) Answering these questions does indeed help explain exchange rates over longer horizons, as the next section shows.

2.3. Applying Microstructure Tools to Exchange Rate Puzzles

What about the big puzzles in exchange rate economics? Two of the biggest puzzles are:[5]

(1) The determination puzzle: exchange rate movements are virtually unrelated to macroeconomic fundamentals (at least over periods of less than about two years); and

(2) The excess volatility puzzle: exchange rates are excessively volatile relative to our best measures of fundamentals.

---

5. Within international finance more broadly, there are four main puzzles, the two listed plus the "forward bias" and "home bias" puzzles. (Forward bias refers to conditional bias—potentially due to a risk premium—in forward exchange rates, whereas home bias refers to investors underinvesting internationally.) For applications of the dispersed information approach to these other puzzles, see Lyons (2001).

These four puzzles have analogues in other markets. For equities, papers that address the puzzles include Roll (1988) on determination, Shiller (1981) on excess volatility, Mehra and Prescott (1985) on equity risk premia, and Coval and Moskowitz (1999) on home bias. (The equity market version of the forward bias puzzle—the so-called equity premium puzzle—is a much looser analogue than the others: the large risk premium on equity is rather stable over time and remains positive, whereas the large risk premium in FX changes over time, including frequent changes in sign.) Microstructure tools are just beginning to be applied to those major equity puzzles (see, for example, Easley, et al. 1999).

Figure 2
Four Months of Exchange Rates and Order Flow, May 1 to August 31, 1996

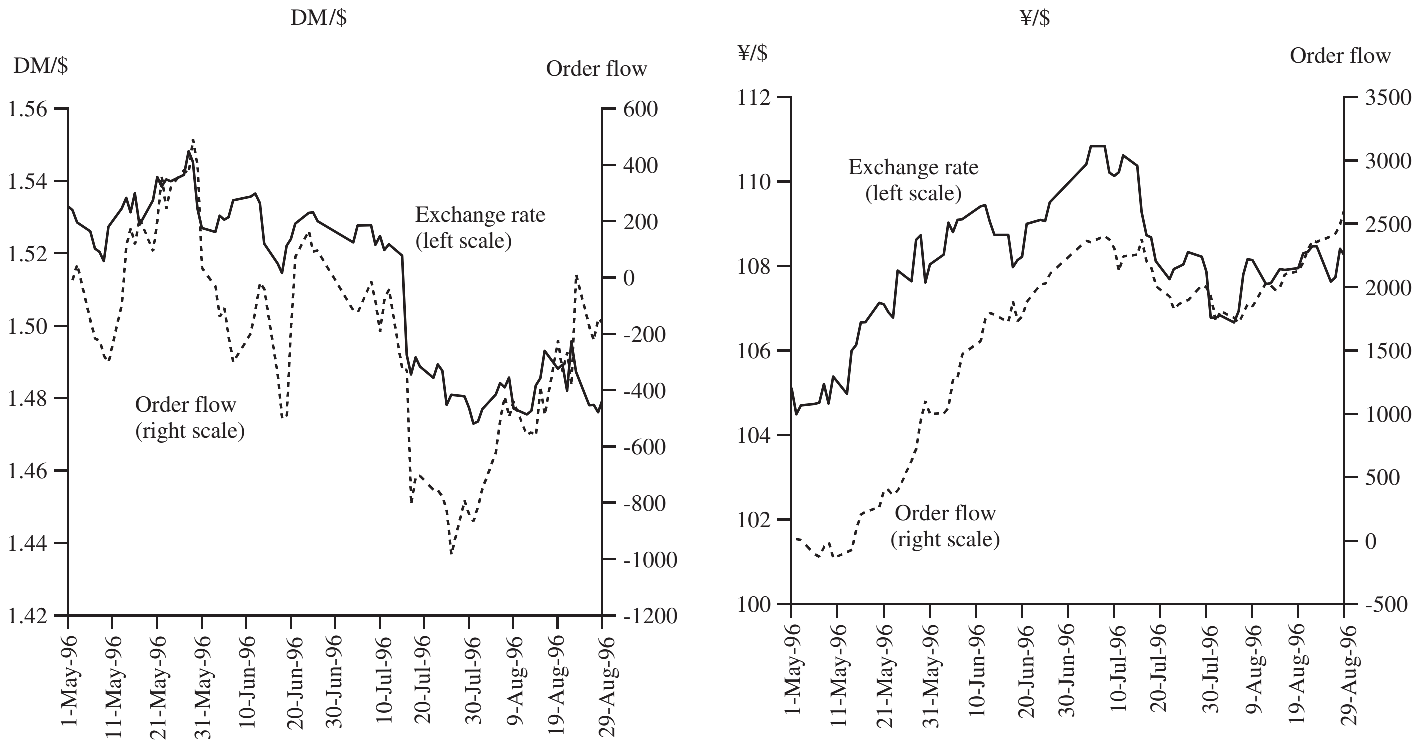

Source: Transactions data from Evans (1997).

The dispersed information approach links these puzzles to one another via expectations formation, i.e., how market participants form their expectations of future fundamentals. It makes this link without departing from rational expectations. Rather, the microstructure approach grounds expectations formation more directly in a richer, information-economic setting. The focus is on information *types* (such as public versus dispersed) and *how* information maps into expectations (e.g., whether the aggregation of order flow "votes" is efficient). The issues of information type and mapping to expectations are precisely where tools from microstructure finance provide resolving power.[6]

2.4. A First Look at the Data

Figure 2 provides a convenient summary of order flow's explanatory power. The solid lines represent the spot rates of the deutsche mark and yen against the dollar over the four-month sample of the Evans (1997) data set. The dashed lines represent cumulative order flow for the respective currencies over the same period. Order flow is the sum of signed trades (starting from the beginning of the sample) between foreign exchange dealers worldwide.[7] Cumulative order flow and nominal exchange rate levels are strongly positively correlated (prices increase with buying pressure). This result is intriguing. Order flow appears to matter for exchange rate determination, and the effect appears to persist (otherwise the exchange rate's level would reflect only concurrent or very recent order flow and not *cumulative* order flow). This persistence is an important property, one that I examine more closely below. For order flow to be helpful in resolving big exchange rate puzzles, its effects have to persist over horizons that match those puzzles (monthly, at a minimum).[8]

6. Of course, the dispersed information approach also has its drawbacks, an important one being the lack of publicly available order flow data over long periods.

7. Because the Evans (1997) data set does not include the size of every trade, this measure of order flow is in fact the number of buys minus sells. That is, if a dealer initiates a trade against another dealer's DM/$ quote, and that trade is a $ purchase (sale), then order flow is +1 (–1). These are cumulated across dealers over each 24-hour trading day (weekend trading—which is minimal—is included in Monday).

8. Readers familiar with the concept of cointegration will recognize that it offers a natural means of testing for a long-run relationship. In Section 4, I present evidence that cumulative order flow and the level of the exchange rate are indeed cointegrated, indicating that the relationship between order flow and prices is not limited to high frequencies. I also show in that section why a long-run relationship of this kind is what one should expect.

3. The Determination Puzzle

This section and the next examine traditional exchange rate puzzles, showing how tools from microstructure finance are used to address them. They are not intended to put these puzzles to rest: the puzzles wouldn't be traditional if they weren't stubborn. My intent is to provide a sense for how to address macro issues by looking under the "micro lamppost."

As noted, textbook models do a poor job of explaining monthly exchange rate changes. In their survey, Frankel and Rose (1995) summarize as follows:[9]

> The Meese and Rogoff analysis at short horizons has never been convincingly overturned or explained. It continues to exert a pessimistic effect on the field of empirical exchange rate modeling in particular and international finance in general…such results indicate that no model based on such standard fundamentals like money supplies, real income, interest rates, inflation rates, and current account balances will ever succeed in explaining or predicting a high percentage of the variation in the exchange rate, at least at short- or medium-term frequencies. (pp. 1,704, 1,708)

This is the determination puzzle. Immense effort has been expended to resolve it.[10]

If determinants are not macro fundamentals like interest rates, money supplies, and trade balances, then what are they? Two alternatives have attracted a lot of attention among macroeconomists. The first is that exchange rate determinants include extraneous variables. These extraneous variables are typically modeled as speculative bubbles. (A bubble is a component of an asset's price that is nonfundamental. A bubble can cause prices to rise so fast that investors are induced to buy, even though the bubble may burst at any time; see, e.g., Meese 1986 and Evans 1986.) On the whole, however, the empirical evidence on bubbles is not supportive; see the survey by Flood and Hodrick (1990). A second alternative to macro fundamentals is irrationality. For example, exchange rates may be determined, in part, from avoidable expectational errors (Dominguez 1986, Frankel and Froot 1987, and Hau 1998). On a priori grounds, many financial economists find this second alternative unappealing. Even if one is sympathetic, however, there is a wide gulf between the presence of irrationality and accounting for exchange rates empirically.[11]

This section addresses the determination puzzle using the dispersed information approach, drawing heavily from work presented in Evans and Lyons (2002). One advantage of this approach is that it directs attention to variables that have escaped the attention of macroeconomists. A telling quote along these lines appears in Meese (1990):

> Omitted variables is another possible explanation for the lack of explanatory power in asset market models. However, empirical researchers have shown considerable imagination in their specification searches, so it is not easy to think of variables that have escaped consideration in an exchange rate equation. (p. 130)

Among the variables escaping consideration, order flow may be the most important.

3.1. A Hybrid Model with Both Macro and Micro Determinants

To establish a link between the micro and macro approaches, Figure 1 introduced a "hybrid" model with components from both. The hybrid model in that figure could be written as follows:

$$\Delta P_t = f(i, m, z) + g(X, I, Z) + \varepsilon_t \,, \tag{1}$$

where the function $f(i,m,z)$ is the macro component of the model and $g(X,I,Z)$ is the microstructure component. The driving variables in the function $f(i,m,z)$ include current and past values of home and foreign nominal interest rates i, money supply m, and other macro determinants, denoted here by z. The driving variables in the function $g(X,I,Z)$ include order flow X (signed so as to indicate direction), a measure of dealer net positions (or inventory) I, and other micro determinants, denoted by Z. An important take-away from the relevant literatures is that $f(i,m,z)$ and $g(X,I,Z)$ depend on more than just current and past values of their determinants—they also depend, crucially, on expectations of the future values of their determinants. This stands to reason: rational markets are forward-looking, so these expectations are important for setting prices today.

9. At longer horizons, e.g., longer than two years, macro models begin to dominate the random walk (e.g., Chinn 1991 and Mark 1995). But exchange rate determination remains a puzzle at horizons shorter than two years (except in cases of hyperinflation, in which case the inflation differential asserts itself as a driving factor, in the spirit of purchasing power parity).

10. The determination puzzle exists in equity markets as well—see Roll (1988). Roll can account for only 20 percent of daily stock returns using traditional equity fundamentals, a result he describes as a "significant challenge to our science."

11. Another alternative to traditional macro modeling is the recent "new open economy macro" approach (e.g., Obstfeld and Rogoff 1995). I do not address this alternative here because, as yet, the approach has not produced empirical exchange rate equations that alter the Meese-Rogoff (1983) conclusions (see Bergin 2001).

Though I have split this stylized hybrid model into two parts, the two parts are not necessarily independent. This will depend on the main micro determinant—order flow *X*—and the type of information it conveys. In fact, order flow conveys two main information types: *payoff information* and *discount rate information.* In macro models, information about future payoffs translates to information about future *(i,m,z)*. One way order flow can convey information about future *(i,m,z)* is by aggregating the information in people's expectations of *(i,m,z)*. (Recall that as a measure of expectations, order flow reflects people's willingness to back their beliefs with money; and like actual expectations, this measure evolves rapidly, in contrast to measures derived from macro data.) To fix ideas, write the price of foreign exchange, P_t, in the standard way as a function of current and expected future macro fundamentals: $P_t = g(f_t, f^e_{t+1})$. In dispersed information models, price setters learn about changes in f^e_{t+1} by observing order flow. Thus, when order flow conveys payoff information, macro and micro determinants are interdependent: order flow acts as a proximate determinant of prices, but standard macro fundamentals are the underlying determinant.[12]

If order flow *X* conveys discount rate information only, then the two sets of determinants *(i,m,z)* and *(X,I,Z)* can indeed be independent. To understand why, suppose the discount rate information conveyed by order flow *X* is about portfolio balance effects (e.g., persistent changes in discount rates, due to changing risk preferences, changing hedging demands, or changing liquidity demands under imperfect substitutability).[13] Now, consider the two monetary macro models (flexible and sticky-price). Portfolio balance effects from order flow *X* are unrelated to these models' specifications of *f(i,m,z)*. This is because the monetary models assume that different-currency assets are perfect substitutes (i.e., they assume that uncovered interest parity holds: assets differing only in their currency denomination have the same expected return). Thus, effects from imperfect substitutability are necessarily independent of the *f(i,m,z)* of these monetary models. In the case of the macro portfolio balance model, in contrast, portfolio balance effects from order flow *X* are quite likely to be related to the determining variables *(i,m,z)*. Indeed, in that model, price effects from imperfect substitutability are the focus of *f(i,m,z)*.

Before describing the hybrid model estimated by Evans and Lyons (2002), let me address some front-end considerations in modeling strategy. First, the determination puzzle concerns exchange rate behavior over months and years, not minutes. Yet most empirical work in microstructure finance is estimated at the transaction frequency. The first order of business is to design a trading model that makes sense at lower frequencies. Several features of the Evans-Lyons model contribute to this (as will be noted specifically below, as the features are presented). Second, because in actual currency markets interdealer order flow is more transparent than customer-dealer order flow, it is more immediately relevant to FX price determination. The hybrid model should reflect this important institutional feature. Third, the model should provide a vehicle for understanding the behavior of interdealer order flow in Figure 2. That figure presents cumulative interdealer flow in the DM/\$ and ¥/\$ markets over the four-month Evans (1997) data set, the same data set used by Evans and Lyons (2002). A puzzling feature is the persistence: there is no obvious evidence of mean reversion in cumulative order flow. How can this be consistent with the fact that individual dealer inventories have a very short half-life (i.e., their positions revert to zero rapidly)? The Evans-Lyons model accounts for this seeming incongruity.

3.2. The Evans-Lyons Model

Consider an infinitely lived, pure exchange economy with two assets, one riskless and one with stochastic payoffs representing foreign exchange. The periodic payoff on foreign exchange, denoted R_t, is composed of a series of increments, so that

$$R_t = \sum_{\tau=1}^{t} \Delta R_\tau . \qquad (2)$$

The increments ΔR_t are i.i.d. Normal$(0,\sigma^2_R)$ and represent the flow of public macroeconomic information—the macro component of the model *f(i,m,z)*. For concreteness, one can think of this abstract payoff increment ΔR_t as changes in interest rates. Periodic payoffs are realized at the beginning of each day.

The foreign exchange market is organized as a decentralized dealership market with *N* dealers, indexed by *i*, and a continuum of nondealer customers (the public), indexed by $z \in [0,1]$. Within each period (day) there are three rounds of trading:

12. If order flow is an informative measure of macro expectations, then it should forecast surprises in important variables (like interest rates). New order flow data sets that cover up to six years of FX trading—such as the data set examined by Fan and Lyons (2001)—provide enough statistical power to test this. The Evans (1997) data set used by Evans and Lyons (2002) is only four months, so they are not able to push in this direction.

13. Lyons (2001) introduces two subcategories of discount rate information: information about inventory effects and information about portfolio balance effects. I do not consider information about inventory effects here because inventory effects are transitory and are therefore unlikely to be relevant for longer-horizon macro puzzles.

FIGURE 3
DAILY TIMING IN THE EVANS AND LYONS (2002) MODEL

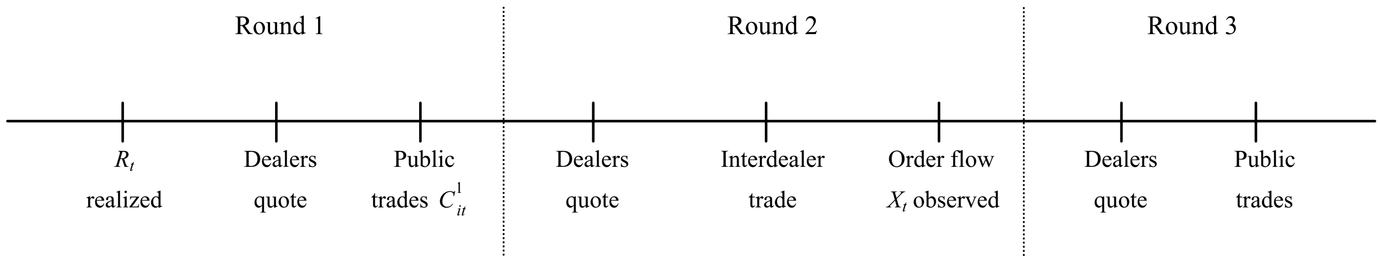

Round 1: dealers trade with the public.
Round 2: dealers trade among themselves to share risk.
Round 3: dealers trade again with the public to share risk more broadly.

The timing within each day is summarized in Figure 3. Dealers and customers all have identical negative exponential utility (constant absolute risk aversion).

Per Figure 3, after observing R_t each dealer sets a quote for his public customers. These quotes are scalar two-way prices, set simultaneously and independently.[14] Denote this dealer *i* quote in Round 1 of day *t* as P^1_{it}. Evans and Lyons show that, in equilibrium, all dealers choose to quote the same price, denoted P^1_t (implied by no arbitrage). Each dealer then receives a customer-order realization C^1_{it} that is executed at his quoted price P^1_t, where C^1_{it}<0 denotes a customer sale (dealer *i* purchase). Each of these *N* customer order realizations is distributed

(3) $$C^1_{it} \sim \text{Normal}(0, \sigma^2_C),$$

and they are uncorrelated across dealers. Importantly, the C^1_{it} realizations are not publicly observable. For later discussion of the model's intuition, it is useful to define the aggregate public demand in Round 1 as the sum of the customer orders received by the *N* dealers:

(4) $$C^1_t = \sum_{i=1}^{N} C^1_{it}.$$

One important choice in specifying the model is the correlation between customer orders C^1_{it} and the stream of payoff increments ΔR_t. This choice determines whether the macro and micro components of the model—the *f(i,m,z)* and *g(X,I,Z)*—are interdependent. If there is no correlation, it is not possible for order flow to convey payoff information. Because Evans and Lyons (2002) have only four months of order flow data, they are unable to determine empirically whether order flow conveys payoff information, discount rate information, or both. They choose to model the customer orders C^1_{it} as distributed independently of the payoff stream R_t—arguably, a less controversial choice. This means that, in their model, the only kind of information that order flow can convey is discount rate information. And because their model rules out inventory effects at the daily frequency (as we shall see below), the discount rate information in their model corresponds to what macroeconomists call portfolio balance effects.

Round 2 is the interdealer trading round. Each dealer simultaneously and independently quotes a scalar two-way price to other dealers P^2_{it}. These interdealer quotes are observable and available to all dealers in the market. Evans and Lyons show that, as in Round 1, all dealers choose to quote the same price, denoted P^2_t. Each dealer then simultaneously and independently trades on other dealers' quotes. (Orders at a given price are split evenly across any dealers quoting that price.) Let T_{it} denote the interdealer trade initiated by dealer *i* in Round 2 of day *t* (negative for dealer *i* sales).

Importantly, at the close of Round 2 all dealers observe the order flow from interdealer trading that day:

(5) $$X_t = \sum_{i=1}^{N} T_{it}.$$

This order flow information is important to the model because it conveys the size and sign of the public order flow in Round 1. To understand why, consider the interdealer trading rule derived by Evans and Lyons:

(6) $$T_{it} = \alpha C^1_{it},$$

where α is a constant coefficient. Each dealer's trade in Round 2 is proportional to the customer order he receives in Round 1. This implies that when dealers observe the interdealer order flow $X_t = \Sigma_i T_{it} = \alpha C^1_t$, they can infer the aggregate public order flow C^1_t in Round 1.

14. Introducing a bid-offer spread (or price schedule) in Round 1 to endogenize the number of dealers is a straightforward extension.

In Round 3, dealers share overnight risk with the nondealer public. This feature is important in distinguishing this model from models focused on intraday trading. Unlike Round 1, the public's trading in Round 3 is nonstochastic. To start the round, each dealer simultaneously and independently quotes a scalar two-way price P_t^3 (also common across dealers). These quotes are observable and available to the public.

A crucial assumption made by Evans and Lyons is that dealers set prices in Round 3 such that the public willingly absorbs all dealer inventory imbalances, so that each dealer ends the day with no net position.[15] As an empirical matter, it is common practice for FX dealers to end each day with no net position, and this squares with the empirical findings (Lyons 1995, Yao 1998). Note too that this assumption rules out inventory effects on prices at the daily frequency (because dealers do not hold overnight positions that require compensation). The Round 3 price that dealers actually quote to induce public absorption of these imbalances depends on the Round 2 interdealer order flow X_t: this interdealer order flow informs dealers of the size of the total position that the public needs to absorb (as noted, $X_t = \alpha C_t^1$).

More precisely, to determine the Round 3 price, dealers need to know two things: the total position that the public needs to absorb (which they learn from X_t), and the public's risk-bearing capacity. Regarding the latter, the public's capacity for bearing foreign exchange risk is assumed less than infinite; i.e., Evans and Lyons assume that foreign- and domestic-currency assets are not perfect substitutes. This is a key assumption: it makes room in the model for portfolio balance effects on prices. Consistent with negative exponential utility, the public's total demand for foreign exchange in Round 3, denoted C_t^3, is a linear function of its expected return conditional on public information:

$$(7) \qquad C_t^3 = \gamma E\left[\Delta P_{t+1}^3 + R_{t+1} \middle| \Omega_t^3\right].$$

The positive coefficient γ captures the aggregate risk-bearing capacity of the public: a larger γ means the public is willing to absorb a larger foreign exchange position for a given expected return. Ω_t^3 is the public information available at the time of trading in Round 3 (which includes all past R_t and X_t).

Evans and Lyons (2002) show that the price at the end of day t is:

$$(8) \qquad P_t = \beta_1 \sum_{\tau=1}^{t} \Delta R_\tau \quad + \quad \beta_2 \sum_{\tau=1}^{t} X_\tau \ .$$

The change in prices from the end of day t–1 to the end of day t can therefore be written as:

$$(9) \qquad \Delta P_t = \beta_1 \Delta R_t + \beta_2 X_t \,,$$

where β_2 is a positive constant (that depends on γ and α).[16] It is not surprising that this price change includes the payoff increment ΔR_t: upon realization, the increment ΔR_t becomes a known (i.e., risk-free) component of the continuing daily payoff R_t, and its discounted value is impounded in the price (β_1).

Let me provide some intuition for the portfolio balance effect—the $\beta_2 X_t$ term. This term is the required price adjustment that induces reabsorption of the random order flow C_t^1 that occurred at the beginning of the day. The value of the parameter β_2 ensures that at the Round 3 price

$$C_t^1 + C_t^3 = 0,$$

i.e., that the dealers have no net overnight position. To understand the link to order flow, recall that the Round 3 price depends on two things: the public's risk-bearing capacity (summarized by γ) and the total position that the public needs to absorb. As noted, dealers learn about the total position the public needs to absorb from order flow X_t. This produces the relation between the interdealer order flow and the subsequent price adjustment.

Let's walk through an example. Consider the price at the close of day t, as described by equation (8). The next day's increment to the daily payoff R, ΔR_{t+1}, is uncertain, but all previous realizations of the payoff increment ΔR are known and are impounded in the price. (Expectations of future realizations do not enter equation (8) due to the simple specification of ΔR_t and C_t^1 as independently distributed across time with mean zero.) To understand the portfolio balance term, $\beta_2 \Sigma_\tau X_\tau$, recall that:

$$X_t \equiv \sum_{i=1}^{N} T_{it} = \alpha C_t^1 \,.$$

Therefore, we can write:

$$\sum_{\tau=1}^{t} X_\tau \ \propto \ \sum_{\tau=1}^{t} C_\tau^1 \,.$$

The sum of the portfolio shifts C_t^1 represents changes in "effective" asset supply, in the sense that these stochastic shifts out of FX are an increase in the net supply that the remainder of the public must absorb. (I couch this in terms of

15. This is tantamount to assuming that—when it comes to bearing overnight risk—the dealers' capacity is small relative to the capacity of the whole public.

16. This model can also be used to generate multiple equilibria. Introducing multiple equilibria obscures the essential portfolio balance logic, however, so I do not pursue this direction here.

supply to connect with traditional portfolio balance intuition.) The total increase in net supply is the sum of past portfolio shifts out of FX,

$$\text{Increase in net supply} = -\sum_{\tau=1}^{t} C_\tau^1 .$$

As is standard in portfolio balance models, increases in supply lower prices, and decreases in supply raise prices. This is why a negative cumulative X_t in equation (8) lowers prices: if cumulative X_t is negative, this implies that cumulative C_t^1 is also negative, which is an increase in net supply, requiring a decrease in prices to clear the market. X_t is the variable that conveys this information about the decrease in net supply (C_t^1 is unobservable). P_t depends on the sum of the X_t because each additional decrease in supply C_t^1 requires a *persistent* incremental increase in prices.

Before moving on to the Evans-Lyons results, I want to address another of their model's important features. Recall that one of their modeling objectives is to clarify the behavior of order flow in Figure 2. Specifically, cumulative order flow is puzzlingly persistent: there is no obvious evidence of mean reversion in cumulative order flow, yet, empirically, individual dealer inventories have a short half-life. How can these two facts be consistent? The Evans-Lyons model provides an explanation. First, note that dealer inventories in the Evans-Lyons model are short-lived: no dealer carries an inventory longer than one day. At the same time, cumulative interdealer order flow in their model is persistent—in fact, it follows a random walk (i.e., there is no mean reversion whatsoever). Equations (5) and (6) hold the key to this random walk result. Interdealer order flow each day is proportional to the public order flow that occurs at the beginning of that day. Because this public order flow is i.i.d. across dealers and time, cumulative interdealer order flow follows a random walk. In the end, these seemingly incongruous facts are consistent because, ultimately, dealers can only decumulate inventory by trading with the public, so aggregate decumulation is not reflected in interdealer flow.[17]

3.3. Evans-Lyons Results

The equation Evans and Lyons actually estimate is the following:

$$\text{(10)} \qquad \Delta p_t = \beta_0 + \beta_1 \Delta(i_t - i_t^*) + \beta_2 X_t + \eta_t \ ,$$

where Δp_t is the change in the log spot rate (DM/\$ or ¥/\$) from the end of day t–1 to the end of day t, $\Delta\left(i_t – i_t^*\right)$ is the change in the overnight interest differential from day t–1 to day t (* denotes DM or ¥), and X_t is the interdealer order flow from the end of day t–1 to the end of day t (negative denotes net dollar sales).

There are two changes in this equation relative to equation (9). First, the public information payoff ΔR_t in equation (9) represents the macro component, or *f(i,m,z)*. To estimate the model, Evans and Lyons have to take a stand on what to include in the regression for ΔR_t. They choose to include changes in the nominal interest differential; i.e., they define $\Delta R_t = \Delta(i_t – i_t^*)$, where i_t is the nominal dollar interest rate and i_t^* is the nominal nondollar interest rate (DM or ¥). As a measure of variation in macro fundamentals, the interest differential is obviously incomplete. The reason Evans and Lyons do not specify a full-blown macro model is because other macro variables (e.g., money supply, output, etc.) are not available at the daily frequency. Accordingly, one should not view their model as fully accommodating both the macro and micro approaches. At the same time, if one were to choose a single macro determinant that needs to be included, interest rates would be it: innovations in interest differentials are the main engine of exchange rate variation in macro models (e.g., the sticky-price monetary model).[18] Moreover, using the change in the interest differential rather than the level is consistent with monetary macro models: in monetary models, shocks to prices are driven by unanticipated changes in the differential.[19]

The second difference in equation (10) relative to (9) is the replacement of the change in the price ΔP_t with the change in the log price Δp_t. This difference makes their estimation more directly comparable to previous macro specifications, since those specifications use the log change (which is approximately equal to a percentage change). As an empirical matter, using Δp_t is inconsequential: the two

17. Consider an example. Starting from $X_t = 0$, an initial customer sale to a dealer does not move X_t from zero because X_t measures interdealer order flow only. After the customer sale (say of one unit), when dealer i unloads the position by selling to another dealer, dealer j, X_t drops to –1. A subsequent sale by dealer j to another dealer, dealer k, reduces X_t further to –2. If a customer happens to buy dealer k's position from him, then the process comes to rest with X_t at –2. In this simple scenario, order flow measured only from trades between customers and dealers would have reverted to zero: the concluding customer trade offsets the initiating customer trade, putting a stop to the hot potato. The interdealer order flow, however, does not revert to zero.

18. Cheung and Chinn (2001) corroborate this empirically: their surveys of foreign exchange traders show that the importance of individual macroeconomic variables shifts over time, but "interest rates always appear to be important."

19. As a diagnostic, though, Evans and Lyons also estimate the model using the level of the differential, à la uncovered interest parity, and find similar results.

different measures for the change in prices produce nearly identical results.

Table 1 presents estimates of the Evans-Lyons model (equation (10)) using daily data for the DM/$ and ¥/$ exchange rates. The coefficient β_2 on order flow X_t is correctly signed and significant, with t-statistics above 5 in both equations. To see that the sign is correct, recall from the model that net purchases of dollars—a positive X_t—should lead to a higher DM price of dollars. The traditional macro fundamental—the interest differential—is correctly signed, but is only significant in the ¥ equation. (The sign should be positive because, in the sticky-price monetary model, for example, an increase in the dollar interest rate i_t induces an immediate dollar appreciation—increase in DM/$.)

The overall fit of the model is striking relative to traditional macro models, with R^2 statistics of 64 percent and 45 percent for the DM and ¥ equations, respectively. Moreover, the explanatory power of these regressions is almost wholly due to order flow X_t: regressing Δp_t on $\Delta(i_t – i_t^*)$ alone, plus a constant, produces an R^2 statistic of less than 1 percent in both equations and coefficients on $\Delta(i_t – i_t^*)$ that are insignificant at the 5 percent level.[20] That the interest differential regains significance once order flow is included, at least in the ¥ equation, is consistent with omitted variable bias in the interest-rates-only specification.

The size of the order flow coefficient is consistent with estimates based on single-dealer data. The coefficient of 2.1 in the DM equation of Table 1 implies that a day with 1,000 more dollar purchases than sales induces an increase in the DM price by 2.1 percent. Given an average trade size in the sample of $3.9 million, this implies that:[21]

- $1 billion of net dollar purchases increases the DM price of a dollar by 0.54 percent.

Table 1
OLS Estimates of the Evans-Lyons Model
$\Delta p_t = \beta_0 + \beta_1 \Delta(i_t - i_t^*) + \beta_2 X_t + \eta_t$

	β_1	β_2	R^2
DM	0.52	2.10	0.64
	(1.5)	(10.5)	
¥	2.48	2.90	0.45
	(2.7)	(6.3)	

Note: *t*-statistics are shown in parentheses (constant not reported). In the case of the DM equation, the *t*-statistics are corrected for heteroskedasticity; there is no evidence of heteroskedasticity in the ¥ equation, and no evidence of serial correlation in either equation. The dependent variable Δp_t is the change in the log spot exchange rate from 4 p.m. GMT on day *t*–1 to 4 p.m. GMT on day *t* (DM/$ or ¥/$). The regressor $\Delta(i_t - i_t^*)$ is the change in the one-day interest differential from day *t*–1 to day *t* (* denotes DM or ¥, annual basis). The regressor X_t is interdealer order flow between 4 p.m. GMT on day *t*–1 and 4 p.m. GMT on day *t* (negative for net dollar sales, in thousands of transactions). The sample spans four months (May 1 to August 31, 1996), which is 89 trading days. (Saturday and Sunday order flow—of which there is little—is included in Monday.)

Equivalently, at a spot rate of 1.5 DM/$, $1 billion of net dollar purchases increases the DM price of a dollar by 0.8 pfennig.

Turning now to estimates at the single-dealer level, these show that information asymmetry induces the dealer to increase the price by 1/100th of a pfennig (0.0001 DM) for every incoming buy order of $10 million (Lyons 1995). That translates to 1 pfennig per $1 billion, versus the 0.8 pfennig per $1 billion found by Evans and Lyons. Though linearly extrapolating the single-dealer estimate (based on individual order sizes around $10 million) to $1 billion of order flow is certainly not an accurate description of single-dealer price elasticity, with multiple dealers it may be a good description of price elasticity marketwide.

3.4. Robustness Checks

To check robustness, Evans and Lyons examine several obvious variations on the model. For example, they include a constant in the regression, even though the model does not call for one; the constant is insignificant for both currencies and has no substantive effect on the other coefficients. Second, in the spirit of uncovered interest parity, they include the level of the interest differential in lieu of its change; the level of the differential is insignificant in both cases. Third, they test for simple forms of nonlinearity, such as adding a squared order flow term, or testing for piece-wise linearity. Though the squared order flow term is insignificant in both equations, and though they find no evidence of piece-wise linearity in the DM equation, they do find some evidence of piece-wise linearity in the ¥ equation (there is a greater sensitivity of the ¥/$ price to order

20. There is a vast empirical literature that attempts to increase the explanatory power of interest rates in exchange rate equations (by introducing individual interest rates as separate regressors, by introducing nonlinearities, etc.). Because these efforts have not been successful, it is very unlikely that variations on the interest rate specification could alter the relative importance of order flow.

21. One of the shortcomings of the Evans (1997) data set is that it does not include the size of each trade, so that order flow is measured as the number of buys minus the number of sells. (The data set does include the total volume over the sample, however, so that an average trade size can be calculated.) This shortcoming must be kept in perspective, however: if the Evans-Lyons results were negative, then data concerns would be serious indeed—the negative results could easily be due to noisy data. But their results are quite positive, which noise alone could not produce. Indeed, that there is noise in the data only underscores the apparent strength of the order flow/price relation.

flow in the downward direction, though estimates for both directions are positive and significant). Fourth, they test whether the order flow/price relation depends on the gross level of activity. They find that it does: in the DM equation, the order flow coefficient is lowest on days when the number of transactions is at a middling level (i.e., the pattern is U-shaped); in the ¥ equation, they find that the order flow coefficient is lowest on days when the number of transactions are at a low level (i.e., the coefficient increases with activity level). Their model is not rich enough to account for these coefficient variations. Fifth, Evans and Lyons decompose contemporaneous order flow into expected and unexpected components (by projecting order flow on past flow). In their model, all order flow X_t is unexpected, but this need not be the case in the data. They find, as one would expect, that order flow's explanatory power comes from its unexpected component.

3.5. Isn't This Just Demand Increases Driving Price Increases?

At first blush, it might appear that the Evans-Lyons results are right out of Economics 101: of course when demand goes up, prices go up. But this misses the most important lesson. A (correct) premise of textbook exchange rate models is that order flow is not necessary to push prices around. Rather, when public information arrives, rational markets adjust prices instantaneously (i.e., excess demand from new information causes prices to adjust without trading—order flow—needing to take place). That order flow explains such a large percentage of price moves underscores the inadequacy of this public information framework. The information the FX market is aggregating is much subtler than textbook models assume. This we learn from our order flow regressions. To summarize, yes, it is demand, but it is demand of a nature very different from the demand in textbook models.

3.6. But What Drives Order Flow?

An important challenge for the microstructure approach is determining what drives order flow, i.e., the first link in the fundamentals/order flow/price chain (Figure 1). Here are three promising strategies for shedding light on this question. Strategy one is to disaggregate order flow. For example, interdealer order flow can be split into large banks versus small banks or investment banks versus commercial banks. Data sets on customer order flow can be split into nonfinancial corporations, leveraged financial institutions (e.g., hedge funds), and unleveraged financial institutions (e.g., mutual and pension funds). Do all these trade types have the same price impact? Someone believing that order flow is just undifferentiated demand would predict that they do. In fact, they do not: certain types of orders (e.g., those from financial institutions) convey more information and therefore have more price impact. People who view order flow as undifferentiated demand overlook this level of analysis, i.e., they overlook the fact that order flow is a vehicle for conveying information. Understanding the information intensity of different trade types brings us closer to this market's underlying information structure.

Strategy two for determining what drives order flow focuses on public information intensity. Consider, for example, periods encompassing scheduled macro announcements. Does order flow account for a smaller share of the price variation within these periods? Or is order flow an important driver of prices even at these times, perhaps helping to reconcile differences in people's mapping from public information to price? Work along these lines, too, will shed light on the forces driving order flow.[22]

Strategy three for determining what drives order flow focuses on discriminating payoff information from discount rate information. If order flow conveys payoff information, then it should forecast surprises in important macro variables like interest rates, money supplies, and trade balances. New order flow data sets that cover many years of FX trading—such as those used by Fan and Lyons (2001)—provide enough statistical power to test this. At a broad level, separating these two types of nonpublic information has implications for how we define the concept of "fundamentals." Order flow that reflects information about payoffs—like expectations of future interest rates—is in keeping with traditional definitions of exchange rate fundamentals. But order flow that reflects changing discount rates may encompass nontraditional exchange rate determinants (e.g., changing risk preferences, or changing hedging demands), calling perhaps for a broader definition.

3.7. Comments on Causality

Under the Evans-Lyons model's null hypothesis, causality runs strictly from order flow to prices. Accordingly, under the null, their estimation is not subject to simultaneity bias. (Unlike the classic supply-demand identification problem,

22. A direct role for macro announcements in determining order flow warrants exploring as well (see Evans and Lyons 2001). Another possible use of macro announcements is to introduce them directly into an Evans-Lyons-type model. This tack is not likely to be fruitful, however: there is a long literature showing that macro announcements are unable to account for exchange rate first moments (though they do help to account for second moments—see Andersen and Bollerslev 1998 and Andersen, et al. 2001).

Evans and Lyons are not simply regressing prices on quantity; quantity—i.e., volume—and order flow are fundamentally different concepts.) Within microstructure theory more broadly, this direction of causality holds in all the canonical models (i.e., the Kyle (1985) auction model and the Glosten and Milgrom (1985) sequential trade model), despite the fact that prices and order flow are determined simultaneously. In these models, price innovations are a function of order flow innovations, not the other way around. Put differently, order flow is indeed a cause of price changes, but only a proximate cause; the underlying cause is nonpublic information (about payoffs or about discount rates).

Although there is no simultaneity bias under the null hypothesis, alternative hypotheses do exist under which causality is reversed. Let me offer a few thoughts. First, in the FX market there is no momentum in daily returns, so it is difficult to rationalize momentum (i.e., feedback) trading strategies in this context. Second, work by Killeen, et al. (2001) shows that daily FX order flow Granger causes returns, but not vice versa. Finally, even in episodes where one would expect feedback trading, e.g., when institutions are distressed, the evidence is not there. The best example of this is the recent experience of October 1998 in the ¥/$ market (after the collapse of the hedge fund Long Term Capital Management). The dollar fell from about 130 ¥/$ to about 118 ¥/$ *in a single day.* The popular press view of this episode was that hedge funds attempting to stop their already substantial losses felt they had to sell into a falling market, thereby making it fall further. Using data on order flows over that period, I examine this special episode as a case study in Lyons (2001). In the end, the data do not support the popular story that distressed hedge funds were rushing from the dollar en masse that day. On the contrary, they provided liquidity (i.e., they bought the dollar on that day in aggregate). The selling came from other sources.

4. The Excess Volatility Puzzle

This section addresses the second puzzle: the excess volatility puzzle. By excess I mean that exchange rates are much more volatile than our best measures of fundamentals. Though other asset markets share this property (e.g., stock markets; see Shiller 1981), the puzzle in FX markets is in many ways distinctive.[23] Consider, for example, the fact that most exchange rates are not allowed to float freely but are managed through intervention by central banks. This fact allows one to address the volatility puzzle in ways not possible in other markets. To understand why, note first that exchange rates are generally less volatile when managed. Given this, one can compare regimes with different management intensities to identify why volatility differs, thereby shedding light on volatility's causes. This approach is common in the literature (e.g., Flood and Rose 1995 and Killeen, Lyons, and Moore 2000). The analysis I present here draws primarily on the empirical findings of Killeen, Lyons, and Moore (KLM).

Before reviewing the KLM findings, let me provide more perspective on the "cross-regime" approach to exchange rate volatility.[24] Why is it that similar macro environments produce more volatility when exchange rates float freely? There are two main approaches to this question, one theoretical and one empirical. The theoretical approach was pioneered by Dornbusch (1976) in his sticky-price monetary model. Dornbusch shows that when goods prices are sticky, but the exchange rate is free to jump, then economic shocks have a disproportionately large effect on the exchange rate—so-called overshooting. From the perspective of excess volatility, the sticky-price monetary model generates the kind of "amplification" that might explain why floating rates are more volatile than fundamentals. This theoretical explanation is not borne out empirically, however: the sticky-price model does not fit the data.

The second main approach to why floating rates are more volatile is empirical. A good example is Flood and Rose (1995), who put the cross-regime logic as follows:

> Intuitively, if exchange rate stability arises across regimes without corresponding variation in macroeconomic volatility, then macroeconomic variables will be unable to explain much exchange rate volatility. Thus existing models, such as monetary models, do not pass our test; indeed, this is also true of any potential model that depends on standard macroeconomic variables. We are driven to the conclusion that the most critical determinants of exchange rate volatility are not macroeconomic. (p. 5)

The central idea here starts with the Flood-Rose finding that managing rates does not change the volatility of funda-

23. Contrary to popular belief, in an absolute sense exchange rates are less volatile than stock prices: the annual standard deviation of exchange rate returns is in the 10–12 percent range for major currencies against the dollar, whereas the annual standard deviation of equity market returns is in the 15–20 percent range (and for individual stocks it is still higher).

24. Exchange rate regimes are not limited to floating and fixed. They fall along a spectrum. Ordered in terms of increasing commitment to the exchange rate target, these regimes include: (1) free float, (2) dirty float, (3) target zone, (4) peg—fixed or crawling, (5) currency board, and (6) monetary union. A dirty float involves some limited intervention. A currency board is an institutional commitment to dedicate monetary policy to the exchange rate target.

FIGURE 4
FF/DM EXCHANGE RATE LEVEL AND CUMULATIVE ORDER FLOW (BUYS MINUS SELLS)

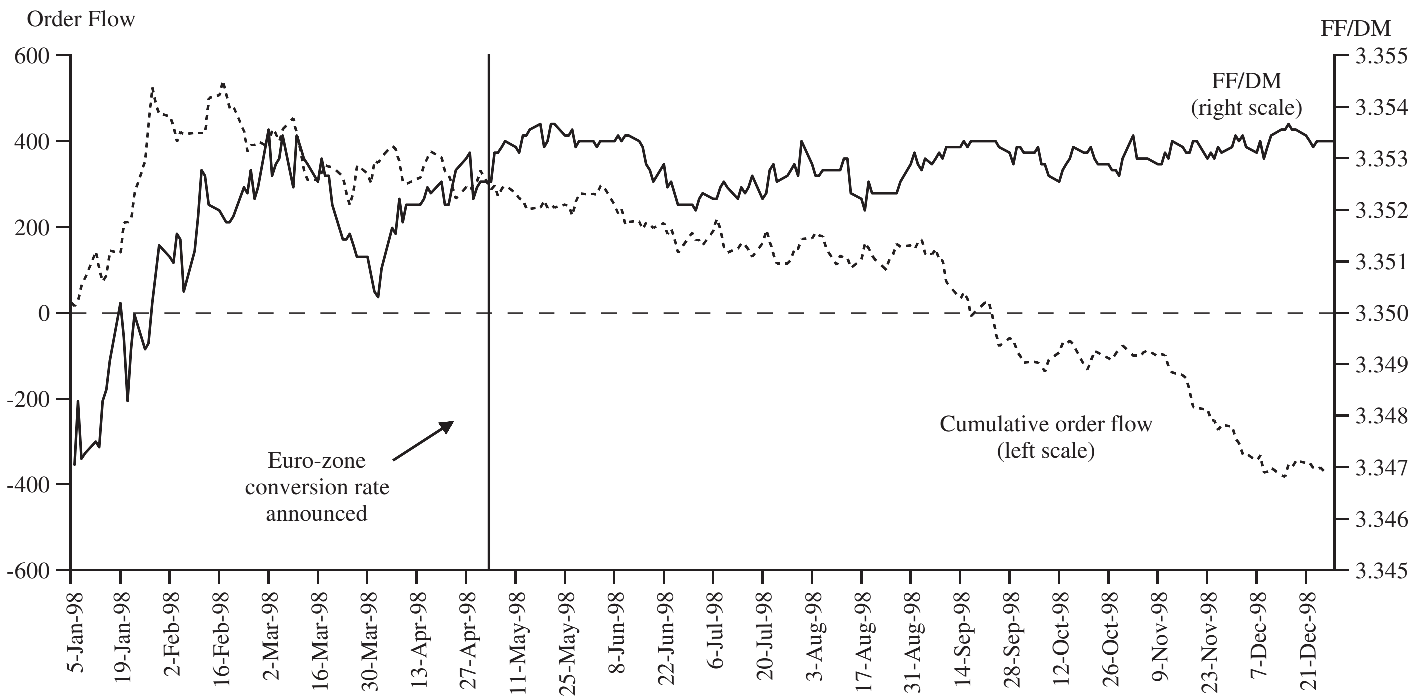

Note: A buy indicates a purchase of DM.
Source: EBS, Datastream.

mentals (fundamentals as described by the public information approach). So, if the volatility reduction from management is not coming from changed behavior of these fundamentals, then it is unlikely these are critical fundamentals. In a sense, then, the Flood-Rose conclusion deepens the puzzle.

KLM take a different tack—they exploit a natural experiment. The experiment is the switch from the European Monetary System (EMS) to European Monetary Union (EMU), which in terms of regimes is a switch from a target zone to a rigidly fixed rate.[25] Starting in January 1999, the euro-country currencies have been rigidly fixed to one another. Before January 1999, however—particularly before May 1998—there was still uncertainty about which countries would participate in the EMU. There was also uncertainty about the timing of interest rate harmonization (which had to occur among the countries adopting the euro).

25. The transition from EMS to EMU was indisputably a transition toward exchange rate fixity. KLM assume that EMU was perfectly credible after the weekend of May 2–3, 1998—the date the 11 "in" countries were selected and the date the internal conversion rates for the eurozone were determined. Extending their model to environments of imperfectly credible fixed rates is a natural direction for further research.

KLM's analysis of this experiment leads them to the following punch line: exchange rates are more volatile under flexible rates because of order flow. Order flow conveys more information under flexible rates, which increases volatility. Under fixed exchange rates, order flow is prevented from conveying information—as a driver of returns, it is "turned off." The intuition for why this happens is tied to demand elasticity. Under floating, the elasticity of public demand is (endogenously) low, due to higher volatility and aversion to the risk this higher volatility entails. This makes room for the types of portfolio balance effects that arise in the Evans-Lyons model and allows order flow to convey information about those effects. Under (perfectly credible) fixed rates, the elasticity of public demand is infinite: return volatility shrinks to zero, making the holding of foreign exchange effectively riskless. This eliminates portfolio balance effects and precludes order flow from conveying this type of information. Consequently, order flow as a return driver is shut down.

Figure 4 provides an initial, suggestive illustration of the KLM results. It shows the relationship between the FF/DM exchange rate and cumulative order flow (interdealer order flow from EBS—see Section 4.3.). The vertical line is May 4, 1998, the first trading day after the announcement of the conversion rates of the euro-participating currencies. The relationship between the two series before May 4 is clearly positive: the correlation is 0.69. After May 4, however, there is a sharp unwinding of long DM positions with no corresponding movement in the exchange rate. In fact, during the second period there is a *negative* correlation of –0.35. Though total variation in the exchange rate is small, the effect of order flow on the exchange rate appears to have changed from one of clear impact to one of no impact. The model KLM develop provides a more formal framework for addressing this issue (to which I now turn).

4.1. Model Sketch

The specification of trading within each day is identical to that of Evans and Lyons (2002). The key difference here is the presence of two trading regimes: a flexible-rate regime followed by a fixed-rate regime. The shift from flexible to fixed rates is a random event that arrives with constant probability p at the end of each trading day (after all trading).[26] Once the regime has shifted to fixed rates the central bank commits to setting $\Delta R_t = 0$ indefinitely.

KLM show that the resulting price level at the end of day t can be written as:

$$(11)\quad P_t = \begin{cases} \lambda_1 \sum_{\tau=1}^{t} \Delta R_\tau + \lambda_2 \sum_{\tau=1}^{t} X_\tau & \text{under flexible rates } (t \le T) \\ \lambda_1 \sum_{\tau=1}^{T} \Delta R_\tau + \lambda_2 \sum_{\tau=1}^{T} X_\tau + \lambda_3 \sum_{\tau=T+1}^{t} X_\tau & \text{under fixed rates } (t > T), \end{cases}$$

where T denotes the day on which the regime shifts from flexible to fixed rates. The message of this equation is important: it describes a cointegrating relationship between the level of the exchange rate, cumulative macro fundamentals, and cumulative order flow. (This long-run relationship between cumulative order flow and the level of the exchange rate is not predicted by any traditional exchange rate model.) The cointegrating vector is regime dependent, however.

Under flexible rates, the change in the exchange rate from the end of day t–1 to the end of day t can be written as:

$$(12)\quad \Delta P_t = \lambda_1 \Delta R_t + \lambda_2 X_t$$

where λ_1 and λ_2 are positive constants. The portfolio balance effects from order flow enter through λ_2, which depends inversely on γ—the elasticity of public demand with respect to expected return—and also on the variances σ_R^2 and σ_C^2.[27]

4.2. Differences across Trading Regimes

Understanding the effects of the different trading regimes—and the changing role of order flow—comes from the effect of the exchange rate regime on equations (11) and (12). Specifically, the parameter γ, which represents the elasticity of public demand, is regime-dependent. This comes from the regime-dependence of the return variance $\text{Var}\left[\Delta P_{t+1} + R_{t+1} \mid \Omega_t^3\right]$ (γ being proportional to the inverse of this variance). The elimination of portfolio balance effects under fixed rates reduces this variance, implying that:

$$(13)\quad \gamma_{\text{flexible}} < \gamma_{\text{fixed}}\ .$$

Public demand is therefore more elastic in the (credible) fixed-rate regime than in the flexible-rate regime. The implication for the price impact parameters λ_2 and λ_3 in equation (11)—henceforth $\lambda_{\text{flexible}}$ and λ_{fixed}, respectively—is the following:

$$(14)\quad \lambda_{\text{flexible}} > \lambda_{\text{fixed}}\ .$$

Thus, the exchange rate reacts more to order flow under flexible rates than under fixed rates. For perfectly credible fixed rates (i.e., for which $\text{Var}\left[\Delta P_{t+1} + R_{t+1} \mid \Omega_t^3\right] = 0$), we have:

$$(15)\quad \lambda_{\text{fixed}} = 0.$$

The exchange rate does not respond to order flow in this case. The intuition is clear: under perfect credibility, the variance of exchange rate returns goes to zero because public demand is perfectly elastic, and vice versa.

For intuition, consider P_{T+1}, the price at the close of the first day of the fixed-rate regime. Foreign exchange is a riskless asset at this point, with return variance equal to zero. A return variance of zero implies that the elasticity of

26. This formulation has two important advantages. First, the effective horizon over which foreign exchange is priced in the flexible-rate regime remains constant. Second, the parameter p provides a compact means of describing regime shifts as far or near. As an empirical matter, particularly in the context of the EMS-EMU transition, this specification serves as a convenient abstraction from reality.

27. The probability p of the regime shift adds a parameter to the Evans-Lyons solution that has no qualitative impact on the coefficients of interest here, namely λ_2 and λ_3.

the public's speculative demand is infinite, and the price impact parameter λ_3 in equation (11) equals zero. This yields a price at the close of trading (Round 3) on day T+1 of:

$$P_{T+1} = \lambda_1 \sum_{t=1}^{T} \Delta R_t + \lambda_2 \sum_{t=1}^{T} X_t \,.$$

The summation over the payoff increment ΔR_t does not include an increment for day T+1 because the central bank maintains ΔR_t at zero in the fixed regime. Though X_{T+1} is not equal to zero, this has no effect on prices because λ_3=0, as noted. This logic holds throughout the fixed-rate regime. Under flexible rates, the economics behind the price impact of order flow is the same as that under the Evans-Lyons model, adjusted only by the change in parameter values due to the possibility of regime switch.

4.3. *The KLM Data Set*

The KLM data set includes daily order flow in the FF/DM market for one year, 1998. The data are from EBS, the electronic interdealer broking system. (At that time, EBS accounted for nearly half of interdealer trading in the largest currencies, which translates into about a third of total trading in major currencies; the Evans-Lyons data reflect the other half of interdealer trading—the direct portion.) By KLM's estimate, their sample accounts for about 18 percent of trading in the DM/FF market in 1998. Daily order flow includes all orders passing through the system over 24 hours starting at midnight GMT (weekdays only).

The data set is rich enough to allow measurement of order flow X_t two ways: number of buys minus number of sells (à la Evans and Lyons 2002) and amount bought minus amount sold (in DM). KLM find that the two measures behave quite similarly: the correlation between the two X_t measures in the flexible-rate portion of the sample (the first four months) is 0.98. They also find that substituting one measure for the other in their analysis has no substantive effect on their findings.

Let me provide a bit more detail on EBS. As noted, EBS is an electronic broking system for trading spot foreign exchange among dealers. It is limit order driven, screen-based, and ex ante anonymous (ex post, counterparties settle directly with one another). The EBS screen displays the best bid and ask prices together with information on the cash amounts available for trading at these prices. Amounts available at prices other than the best bid and offer are not displayed. Activity fields on this screen track a dealer's own recent trades, including prices and amount, as well as the recent trades executed on EBS systemwide.

There are two ways that dealers can trade currency on EBS. Dealers can either post prices (i.e., submit "limit orders"), which does not ensure execution, or dealers can "hit" prices (i.e., submit "market orders"), which does ensure execution. To construct a measure of order flow, trades are signed according to the direction of the latter—the initiator of the transaction.

When a dealer submits a limit order, she is displaying to other dealers an intention to buy or sell a given cash amount at a specified price.[28] Bid prices (limit order buys) and offer prices (limit order sells) are submitted with the hope of being executed against the market order of another dealer—the "initiator" of the trade. To be a bit more precise, not all initiating orders arrive in the form of market orders. Sometimes, a dealer will submit a limit order buy that is equal to or higher than the current best offer (or will submit a limit order sell that is equal to or lower than the current best bid). When this happens, the incoming limit order is treated as if it were a market order and is executed against the best opposing limit order immediately. In these cases, the incoming limit order is the initiating side of the trade.

4.4. *Results*

The relationship between cumulative order flow and the exchange rate is illustrated in Figure 4. We saw that the effect of order flow on the exchange rate appears to have changed from one of clear impact to one of no impact. The results that follow address this more formally, based on the KLM model's testable implications.

The analysis proceeds in two stages. First, KLM address whether there is evidence of a cointegrating relationship between order flow and prices, as the model predicts. This first stage also examines the related issues of stationarity and long-run coefficient sizes. The second stage addresses the degree to which order flow is exogenous (as assumed in their model). This stage includes a test for reverse Granger causality, i.e., statistical causality running from the exchange rate to order flow.

4.4.1. *Stage 1: Cointegration and Related Issues*

Let us begin by repeating equation (11) from the model, which establishes the relationship between the level of the exchange rate P_t, a variable summarizing public information ($\Sigma \Delta R_t$), and accumulated order flow (ΣX_t).

28. EBS has a prescreened credit facility whereby dealers can tell which prices correspond to trades that would not violate their counterparty credit limits, thereby eliminating the potential for failed deals because of credit issues.

$$(11)\quad P_t = \begin{cases} \lambda_1 \sum_{\tau=1}^{t} \Delta R_\tau + \lambda_2 \sum_{\tau=1}^{t} X_\tau \text{ under flexible rates } (t \leq T) \\ \lambda_1 \sum_{\tau=1}^{T} \Delta R_\tau + \lambda_2 \sum_{\tau=1}^{T} X_\tau + \lambda_3 \sum_{\tau=T+1}^{t} X_\tau \\ \qquad \text{under fixed rates } (t > T). \end{cases}$$

Like Evans and Lyons (2002), KLM use the interest differential as the public information variable (the Paris interbank offer rate minus the Frankfurt interbank offer rate).

The KLM model predicts that before May 4, 1998, all these variables are nonstationary and are cointegrated. After May 4, the model predicts that the exchange rate converges to its conversion rate and should be stationary. During this latter period (May to December), therefore, equation (11) makes sense only if the price impact coefficient, λ, goes to zero (as the model predicts), or if accumulated order flow becomes stationary. Otherwise, the regression is unbalanced, with some stationary variables and some nonstationary variables.

The first step is to test whether the relevant variables are nonstationary. KLM find that in the first four months of 1998, all variables are indeed nonstationary (inference based on Dickey-Fuller tests). In the remaining eight months, the exchange rate is stationary, as expected, but both cumulative order flow and the interest differential remain nonstationary. These results are consistent with a price impact parameter λ_3 in the latter period of zero. It is important to determine, however, whether equation (11) actually holds for the January to April period, i.e., whether the variables are cointegrated, as the model predicts.

KLM use the Johansen procedure to test for cointegration (Johansen 1992). The unrestricted vector autoregression (VAR) is assumed to consist of the three variables—the exchange rate, cumulative order flow, and the interest differential—as well as a constant and a trend. After testing various possible lag lengths, KLM find evidence that a lag length of 1 is appropriate.

The cointegration tests show that there is indeed one cointegrating vector. (The null of no cointegrating vectors is rejected in favor of the alternative of at least one cointegrating vector. But the null of one cointegrating vector cannot be rejected in favor of the alternative of at least two.) This implies that a linear combination of the three variables is stationary, as the KLM model predicts.

KLM go one step further and implement the test for cointegration without the interest differential. They find evidence of one cointegrating vector in that case, too, now between the exchange rate and cumulative order flow. The finding of one cointegrating vector in both the bivariate and trivariate systems suggests that the interest differential enters the trivariate cointegrating vector with a coefficient of zero. When KLM estimate the parameters of the cointegrating vector directly, this is exactly what they find: they cannot reject that the interest differential has a coefficient of zero. By contrast, the coefficient on cumulative order flow is highly significant and correctly signed. (The size of the coefficient implies that a 1 percent increase in cumulative order flow moves the spot rate by about five basis points.)[29] These findings of cointegration and an order flow coefficient that is correctly signed are supportive of their model's emphasis on order flow, even in the long run. At the same time, the lack of explanatory power in the interest differential suggests that this specialization of the payoff increment ΔR_t is deficient (in keeping with the negative results of the macro literature more generally).

4.4.2. Exogeneity of Order Flow

An important question facing the dispersed information approach is the degree to which causality can be viewed as running strictly from order flow to the exchange rate, rather than running in both directions. The KLM framework provides a convenient way to address this question. In particular, if a system of variables is cointegrated, then it has an error-correction representation (see Engle and Granger 1987). These error-correction representations provide clues about the direction of causality. Specifically, the error-correction representation allows one to determine whether the burden of adjustment to long-run equilibrium falls on the exchange rate, on cumulative order flow, or both. If adjustment falls at least in part on order flow, then order flow is responding to the rest of the system (i.e., it is not exogenous in the way specified by the Evans-Lyons and KLM models).

The KLM findings suggest that causality is indeed running strictly from order flow to prices and not the other way around. KLM test this by estimating the error-correction term in both the exchange rate and order flow equations. They find that the error-correction term is highly significant in the exchange rate equation, whereas the error-correction term in the order flow equation is insignificant. This implies that adjustment to long-run equilibrium is occurring via the exchange rate. More intuitively, *when a gap opens in the long-run relationship between cumulative order flow and the exchange rate, it is the exchange rate that adjusts to reduce the gap, not*

29. In their sample, the mean value of cumulative order flow is DM1.38 billion.

cumulative order flow. In the parlance of the literature, the insignificance of the error-correction term in the order flow equation means that order flow is weakly exogenous. Further, KLM show that there is no evidence of Granger causality running from the exchange rate to order flow (i.e., feedback trading is not taking place). This combination of weak exogeneity and the absence of Granger causality implies that cumulative order flow is strongly exogenous. Finally, the KLM error-correction estimates suggest that about one-third of departures from long-run equilibrium is dissipated each day.

To summarize, the KLM analysis addresses the excess volatility puzzle on two fronts, one theoretical and one empirical. On the theoretical front, they provide a new approach—based on order flow—for why volatility is high when exchange rates float freely. The punch line of their approach is that an important source of volatility is order flow or, more precisely, the information order flow conveys. Under floating, the elasticity of public demand is (endogenously) low, due to higher volatility and aversion to the risk that higher volatility entails. This makes room for the portfolio balance effects that arise in the Evans-Lyons model and allows order flow to convey information about those effects. Under (perfectly credible) fixed rates, the elasticity of public demand is infinite: return volatility shrinks to zero, making the holding of foreign exchange effectively riskless. This eliminates portfolio balance effects and precludes order flow from conveying this type of information. Thus, under fixed rates, order flow as a return driver is shut down.

A nice feature of the KLM approach to excess volatility, relative to other approaches, is that its implications can be brought to the data. There are many fine theoretical papers on excess exchange rate volatility (see, e.g., Hau 1998 and Jeanne and Rose 1999, and references to earlier work contained therein). But, in general, little of the existing theoretical work is easily implemented empirically. The order flow focus of the KLM approach makes it readily implementable. That said, the specific results that KLM offer are only suggestively supportive of their particular story. Much more empirical analysis along these lines remains to be done.

Two of the KLM empirical findings are especially relevant to interpreting work on order flow more generally. First, they find that Granger causality runs from order flow to the exchange rate, but not vice versa. True, Granger causality is not the same as economic causality. Nevertheless, the result does help assuage concern. Second, they find that gaps in the relationship between cumulative order flow and the level of the exchange rate engender an exchange rate response but not an order flow response. This result, too, helps assuage concern about the direction of causality between these two variables.

One might be tempted to conclude that data for only four months are not enough to produce reliable analysis of cointegration. An important aspect of the KLM results should assuage this concern, however. Recall that KLM find rapid adjustment back to the cointegrating relationship (their error-correction estimates suggest that about one-third of departures from long-run equilibrium is dissipated each day). The half-life of these departures is therefore only about two days. Data for four months are enough to cover about 45 of these half-lives, quite a lot in the context of estimating cointegrating relationships. For comparison, estimates of adjustment back to the cointegrating relationship of purchasing power parity generate half-lives of around 5 years. One would need over 200 years of data to estimate PPP error correction with as many half-lives in the sample.

Note, too, that the KLM model provides a different perspective on exchange rate credibility. In their model, a credible fixed rate is one in which the private sector, not the central bank, willingly absorbs innovations in order flow.[30] The textbook treatment of fixed-rate regimes, in contrast, is centered on the willingness of the central bank to buy and sell domestic currency at a predetermined price; i.e., it is the central bank that absorbs the order flow. If the central bank needs to intervene, the fixed exchange rate regime is already in difficulty because the private sector's demand for order flow is no longer perfectly elastic. It may be useful to revisit analysis of currency crises with this possibility in mind.

Finally, to recap, the KLM model provides a new explanation for the excess volatility puzzle. Shocks to order flow induce volatility under flexible rates because they have portfolio balance effects on price, whereas under fixed rates the same shocks do not have portfolio balance effects. These effects arise in one regime and not the other because the elasticity of speculative demand for foreign exchange is (endogenously) regime-dependent: under flexible rates, low elasticity magnifies portfolio balance effects; under credibly fixed rates, the elasticity of speculative demand is infinite, eliminating portfolio balance effects.

30. This is a theoretical point. Empirically, it appears that there was little intervention by the national central banks or the European Central Bank in the period from May to December, 1998 (these banks are not terribly forthcoming with intervention data over this period).

5. Conclusion

I have argued that abstracting from information aggregation when analyzing exchange rates misses quite a lot. The argument commonly offered in support of this abstraction—that dispersed information is rapidly summarized in public macro variables—is untenable. The abstraction would be easier to defend if either (1) both the public and dispersed information approaches performed well empirically or (2) both approaches performed poorly. In reality, the dispersed information approach performs rather well (e.g., Payne 1999 and Evans and Lyons 2002) while the public information approach does not.

How, specifically, can one identify the information that determines order flow? The notion of order flow as an intermediate link between information and prices suggests several strategies for answering this question, all of which are part of ongoing research. Three in particular are outlined here.

One strategy for linking order flow to underlying determinants starts by decomposing order flow. (That it can be decomposed is one of its nice properties.) Fan and Lyons (2001) test whether all parts of the aggregate order flow have the same price impact. They do not: the price impact of FX orders from financial institutions (e.g., mutual funds and hedge funds) is significantly higher than the price impact of orders from nonfinancial corporations. This suggests that order flow is not just undifferentiated demand. Rather, the orders of some participants are more informative than the orders of others. Analyzing order flow's parts gives us clues as to the underlying information structure.

A second strategy for linking order flow to underlying determinants is based on the view that order flow measures individuals' changing expectations. As a measure of expectations, it reflects a willingness to back one's beliefs with money—the backed-by-money expectational votes, if you will. Expectations measured from macro data, on the other hand, are slow-moving and imprecise. If order flow is serving as an expectations proxy, then it should forecast surprises in important macroeconomic variables (like interest rates). New order flow data sets that cover up to six years of FX trading provide enough statistical power to test this. Note too that this line of research offers a possible explanation of the Meese and Rogoff (1983) findings. To understand why, write the price of foreign exchange, P_t, in the standard way as a function of current and expected future macro fundamentals: $P_t = g(f_t, f^e_{t+1})$. If (big if) the macro variables that order flow is forecasting are largely beyond the one-year horizon, then the empirical link between exchange rates and *current* macro variables f_t will be loose. That macro empirical results are more positive at horizons beyond one year is consistent with this "anticipation" hypothesis.

A third strategy for determining what drives order flow focuses on public information intensity. Consider, for example, periods encompassing scheduled macro announcements. Does order flow account for a smaller share of price variation within these periods? Or is order flow an important driver of prices even at these times, perhaps helping to reconcile differences in people's mapping from public information to prices? Work along these lines, too, will shed light on the forces driving order flow (see, e.g., Evans and Lyons 2001).

References

Andersen, T., and T. Bollerslev. 1998. "Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies." *Journal of Finance* 53, pp. 219–266.

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega. 2001. "Micro Effects of Macro Announcements: Real-time Price Discovery in Foreign Exchange." Typescript, Northwestern University (September).

Bergin, P. 2001. "Putting the 'New Open Economy Macroeconomics' to a Test." Typescript, Department of Economics, University of California, Davis (September).

Cheung, Y., and M. Chinn. 2001. "Currency Traders and Exchange Rate Dynamics: A Survey of the U.S. Market." *Journal of International Money and Finance* 20, pp. 439–471.

Chinn, M. 1991. "Some Linear and Non-linear Thoughts on Exchange Rates." *Journal of International Money and Finance* 10, pp. 214–230.

Coval, J., and T. Moskowitz. 1999. "Home Bias at Home: Local Equity Preference in Domestic Portfolios." *Journal of Finance* 54, pp. 2,045–2,074.

Dominguez, K. 1986. "Are Foreign Exchange Forecasts Rational? New Evidence from Survey Data." *Economics Letters* 21, pp. 277–281.

Dornbusch, R. 1976. "Expectations and Exchange Rate Dynamics." *Journal of Political Economy* 84, pp. 1,161–1,176.

Easley, D., S. Hvidkjaer, and M. O'Hara. 1999. "Is Information Risk a Determinant of Asset Returns?" Typescript, Cornell University (November).

Engle, R., and C. Granger. 1987. "Cointegration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55, pp. 251–276.

Evans, G. 1986. "A Test for Speculative Bubbles in the Sterling-Dollar Exchange Rate." *American Economic Review* 76, pp. 621–636.

Evans, M. 1997. "The Microstructure of Foreign Exchange Dynamics." Typescript, Georgetown University (November).

Evans, M. 2001. "FX Trading and Exchange Rate Dynamics." NBER Working Paper 8116 (February). Forthcoming in *Journal of Finance.*

Evans, M., and R. Lyons. 2001. "Why Order Flow Explains Exchange Rates." Typescript, University of California, Berkeley (October).

Evans, M., and R. Lyons. 2002. "Order Flow and Exchange Rate Dynamics." *Journal of Political Economy* 110, pp. 170–180 (Long version available at www.haas.berkeley.edu/~lyons).

Fan, M., and R. Lyons. 2001. "Customer-Dealer Trading in the Foreign Exchange Market." Typescript, University of California, Berkeley (July).

Flood. R., and R. Hodrick. 1990. "On Testing for Speculative Bubbles." *Journal of Economic Perspectives* 4, pp. 85–101.

Flood, R., and A. Rose. 1995. "Fixing Exchange Rates: A Virtual Quest for Fundamentals." *Journal of Monetary Economics* 36, pp. 3–37.

Frankel, J., and K. Froot. 1987. "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations." *American Economic Review* 77, pp. 133–153.

Frankel, J., G. Galli, and A. Giovannini. 1996. "Introduction." In *The Microstructure of Foreign Exchange Markets.* Chicago: University of Chicago Press, pp. 1–15.

Frankel, J., and A. Rose. 1995. "Empirical Research on Nominal Exchange Rates." In *Handbook of International Economics,* eds. G. Grossman and K. Rogoff. Amsterdam: Elsevier Science, pp. 1,689–1,729.

French, K., and R. Roll. 1986. "Stock Return Variance: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics* 17, pp. 99–117.

Glosten, L., and P. Milgrom. 1985. "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents." *Journal of Financial Economics* 14, pp. 71–100.

Hasbrouck, J. 1991. "Measuring the Information Content of Stock Trades." *Journal of Finance* 46, pp. 179–207.

Hau, H. 1998. "Competitive Entry and Endogenous Risk in the Foreign Exchange Market." *Review of Financial Studies* 11, pp. 757–788.

Hayek, F. 1945. "The Use of Knowledge in Society." *American Economic Review* (September).

Jeanne, O., and A. Rose. 1999. "Noise Trading and Exchange Rate Regimes." NBER Working Paper #7104 (April). Forthcoming in *Quarterly Journal of Economics.*

Johansen, S. 1992. "Cointegration in Partial Systems and the Efficiency of Single Equation Analysis." *Journal of Econometrics* 52, pp. 389–402.

Killeen, W., R. Lyons, and M. Moore. 2001. "Fixed versus Flexible: Lessons from EMS Order Flow." NBER Working Paper #8491 (September).

Kyle, A. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53, pp. 1,315–1,335.

Lyons, R. 1995. "Tests of Microstructural Hypotheses in the Foreign Exchange Market." *Journal of Financial Economics* 39, pp. 321–351.

Lyons, R. 2001. *The Microstructure Approach to Exchange Rates.* Cambridge, MA: MIT Press (November) (chapters available at www.haas.berkeley.edu/~lyons/NewBook.html).

Mark, N. 1995. "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability." *American Economic Review* 85, pp. 201–218.

Meese, R. 1986. "Testing for Bubbles in Exchange Markets." *Journal of Political Economy* 94, pp. 345–373.

Meese, R. 1990. "Currency Fluctuations in the Post-Bretton Woods Era." *Journal of Economic Perspectives* 4, pp. 117–134.

Meese, R., and K. Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies." *Journal of International Economics* 14, pp. 3–24.

Mehra, R., and E. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15, pp. 145–161.

Obstfeld, M., and K. Rogoff. 1995. "Exchange Rate Dynamics Redux." *Journal of Political Economy* 103, pp. 624–660.

Payne, R. 1999. "Informed Trade in Spot Foreign Exchange Markets: An Empirical Investigation." Typescript, London School of Economics (January).

Peiers, B. 1997. "Informed Traders, Intervention, and Price Leadership: A Deeper View of the Microstructure of the Foreign Exchange Market." *Journal of Finance* 52, pp. 1,589–1,614.

Rime, D. 2000. "Private or Public Information in Foreign Exchange Markets? An Empirical Analysis." Typescript, Norwegian School of Management (April) (available at www.uio.no/~dagfinri).

Roll, R. 1988. "R^2." *Journal of Finance* 43, pp. 541–566.

Shiller, R. 1981. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71, pp. 421–436.

Taylor, M. 1995. "The Economics of Exchange Rates." *Journal of Economic Literature* 33, pp. 13–47.

Yao, J. 1998. "Market Making in the Interbank Foreign Exchange Market." Working Paper #S-98-3, New York University Salomon Center.