

# Order Flow and Exchange Rate Dynamics

Martin D. D. Evans\*

Richard K. Lyons

This draft: 29 January 2001

## Abstract

Macroeconomic models of nominal exchange rates perform poorly. The proportion of monthly exchange rate changes that these models can explain is essentially zero. This paper presents a model of a new kind. Instead of relying exclusively on macroeconomic determinants, the model includes a determinant from the field of microstructure—order flow. Order flow is the proximate determinant of price in all microstructure models. This is a radically different approach to exchange rate determination. It is also strikingly successful in accounting for realized rates. Our model of daily changes in log exchange rates produces  $R^2$  statistics above 60 percent. For the DM/\$ spot market, we find that \$1 billion of net dollar purchases increases the DM price of a dollar by about 0.5 percent.

### Correspondence

Richard K. Lyons  
Haas School of Business, UC Berkeley  
Berkeley, CA 94720-1900  
Tel: 510-642-1059, Fax: 510-643-1420  
lyons@haas.berkeley.edu  
www.haas.berkeley.edu/~lyons

---

\* Respective affiliations are Georgetown University and NBER, and UC Berkeley and NBER. We thank the following for valuable comments: two anonymous referees, Massimiliano Croce, Menzie Chinn, Peter DeMarzo, Frank Diebold, Petra Geraats, Eric Jondeau, Robert McCauley, Richard Meese, Michael Melvin, Peter Reiss, Andrew Rose, Mark Taranto, Ingrid Werner, Alwyn Young, and seminar participants at Chicago, Wharton, Columbia, MIT, Iowa, Houston, Stanford, UC Berkeley, the 1999 NBER Summer Institute (IFM), the December 1999 NBER program meeting in Microstructure, and the August 2000 BIS workshop on Market Liquidity. Lyons thanks the National Science Foundation for financial assistance.

# Order Flow and Exchange Rate Dynamics

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.

Richard Meese (1990)

## 1. Motivation: Microstructure Meets Exchange Rate Economics

Since the landmark papers of Meese and Rogoff (1983a, 1983b), exchange rate economics has been in crisis. It is in crisis in the sense that current macroeconomic approaches to exchange rates are empirical failures: the proportion of monthly exchange rate changes that current models can explain is essentially zero. In their survey, Frankel and Rose (1995) write “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.”<sup>1</sup>

Which direction to turn is not obvious. Flood and Rose (1995), for example, are “driven to the conclusion that the most critical determinants of exchange rate volatility are not macroeconomic.” If determinants are not macro fundamentals like interest rates, money supplies, and trade balances, then what are they? Two alternatives have attracted attention. The first is that exchange rate determinants include extraneous variables. These extraneous variables are typically modeled as rational speculative bubbles (Blanchard 1979, Dornbusch 1982, Meese 1986, and Evans 1986, among others). Though the jury is still out, Flood and Hodrick (1990) conclude that the bubble alternative remains unconvincing. 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, among others). On *a priori* grounds,

---

<sup>1</sup> The relevant literature is vast. Recent surveys include Frankel and Rose (1995), Isard (1995), and Taylor (1995).

many economists find this second alternative unappealing. Even if one is sympathetic to the presence of irrationality, there is a wide gulf between its presence and accounting for exchange rates empirically. Until it can produce an empirical account, this too will remain an unconvincing alternative.

Our paper moves in a new direction: the microeconomics of asset pricing. This direction makes available a rich set of models from the field of microstructure finance. These models are largely new to exchange rate economics, and in this sense they provide a fresh approach. For example, microstructure models direct attention to new variables, variables that have “escaped the consideration” of macroeconomists (borrowing from the opening quote). The most important of these variables is order flow.<sup>2</sup> Order flow is the proximate determinant of price in all microstructure models. (That order flow determines price is therefore robust to differences in market structure, which makes this property more general than it might seem.) Our analysis draws heavily on this causal link from order flow to price. One level deeper, microstructure models also provide discipline for thinking about how order flow itself is determined. Information is key here—in particular, information that currency markets need to aggregate. This can include traditional macro fundamentals, but is not limited to them. In sum, our microeconomic approach provides a new type of alternative to the traditional macro approach, one that does not rely on extraneous information or irrationality.<sup>3</sup>

Turning to the data, we find that order flow does indeed matter for exchange-rate determination. By “matter” we mean that order flow explains most of the variation in nominal exchange rates over periods as long as four months. The graphs below provide a convenient summary of this explanatory power. The solid lines are the spot rates of the DM and Yen against the Dollar over our four-month sample (May 1 to August 31, 1996). The dashed lines are marketwide order-flow

---

<sup>2</sup> Order flow is a measure of buying/selling pressure. It is the net of buyer-initiated orders and seller-initiated orders. In a dealer market such as spot foreign exchange, it is the dealers who absorb this order flow, and they are compensated for doing so. (In an auction market, limit orders absorb the flow of market orders.)

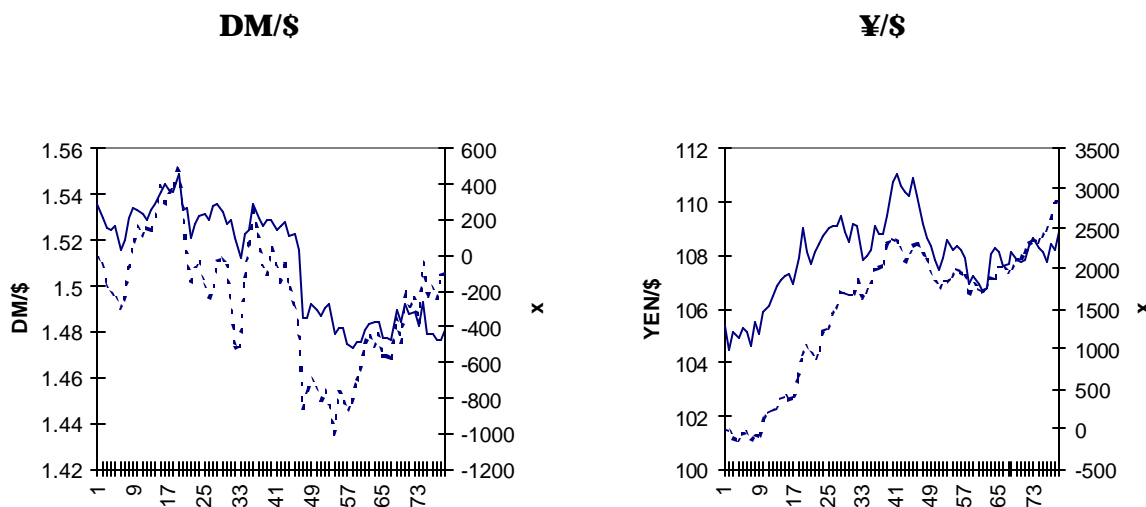
<sup>3</sup> Another alternative to traditional macro modeling is the recent “new open-economy macro” approach (e.g., Obstfeld and Rogoff 1995). We do not address this approach in this paragraph because, as yet, it has not produced an empirical literature.

for the respective currencies. Order flow, denoted by  $x$ , is the sum over time of signed trades between foreign exchange dealers worldwide.<sup>4</sup>

### Figure 1

Four Months of Exchange Rates (solid) and Order Flow (dashed)

May 1-August 31, 1996



Order flow and nominal exchange rates are strongly positively correlated (price increases with buying pressure). Macroeconomic exchange rate models, in contrast, produce virtually no correlation over periods as short as four months.

To address this more formally, we develop and estimate a model that includes both macroeconomic determinants (e.g., interest rates) and a microstructure determinant (order flow). Our estimates verify the significance of the above correlation. The model accounts for about 60 percent of daily changes in the DM/\$ exchange rate. For comparison, macro models rarely account for even 10 percent of monthly changes. Our daily frequency is noteworthy: though our model draws

<sup>4</sup> For example, 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). In spot foreign exchange, roughly 75% of total volume is between dealers (25% is between dealers and non-dealer customers).

from microstructure, it is not estimated at the transaction frequency. Daily analysis is in the missing middle between past microstructure work (tick-by-tick data) and past macro work (monthly data). Bridging the two helps clarify how lower-frequency exchange rates emerge from the market's operation in real time.

To complement these in-sample results, we also examine the model's out-of-sample forecasting ability. Work by Meese and Rogoff (1983a) examines short-horizon forecasts (1 to 12 months). They find that a random walk model out-forecasts the leading macro models, even when macro-model "forecasts" are based on realized future fundamentals. Subsequent work lengthens the horizon beyond 12 months and finds that macro models begin to dominate the random walk (Meese and Rogoff 1983b, Chinn 1991, Chinn and Meese 1994, and Mark 1995). But results at shorter horizons remain a puzzle. Here we examine horizons of less than one month. (Transaction data sets that are currently available are too short to generate statistical power at monthly horizons.) We find that at horizons from one-day to two-weeks, our model produces better forecasts than the random-walk model (over 30 percent lower root mean squared error).

The relation we find between exchange rates and order flow is not inconsistent with the macro approach, but it does raise several concerns. Under the macro approach, order flow should not matter for exchange rate determination: macroeconomic information is publicly available—it is impounded in exchange rates without the need for order flow. More precisely, the macro approach typically assumes that: (1) all information relevant for exchange rate determination is common knowledge; and (2) the mapping from that information to equilibrium prices is also common knowledge. If either of these two assumptions is relaxed, however, then order flow will convey information about market-clearing prices. Relaxing the second assumption should not be controversial, given the failure of current exchange-rate models. Direct evidence, too, corroborates that order flow conveys relevant information (Lyons 1995, Yao 1997, Covrig and Melvin 1998, Ito, Lyons and Melvin 1998, Cheung and Wong 1998, Bjonnes and Rime 1998, Evans 1999, Naranjo and Nimalendran 1999, and Payne 1999.)<sup>5</sup>

---

<sup>5</sup> The standard example of order flow that conveys non-public information is orders from central bank intervention. (Within our four-month sample, however, the Fed never intervened.) Probably more

Note that order flow being a proximate determinant of exchange rates does not preclude macro fundamentals from being the underlying determinant. Macro fundamentals in exchange rate equations may be so imprecisely measured that order-flow provides a better “proxy” of their variation. This interpretation of order flow as a proxy for macro fundamentals is particularly plausible with respect to expectations: standard empirical measures of expected future fundamentals are obviously imprecise.<sup>6</sup> Orders, on the other hand, reflect a willingness to back one's beliefs with real money (unlike survey-based measures of expectations). Measuring order flow under this interpretation is akin to counting the backed-by-money expectational votes.

This paper has six remaining sections. Section 2 contrasts the micro and macro approaches to exchange rates. Section 3 develops a model that includes both micro and macro determinants. Section 4 describes our data. Section 5 presents our results. Section 6 provides perspective on our results. Section 7 concludes.

## **2. Models: Spanning the Micro-Macro Divide**

A core distinction between a microstructure approach to exchange rates and the traditional macro approach is the role of trades in price determination. In macro models, trades have no distinct role in determining price. In microstructure models, trades have a leading role—they are the proximate cause of price adjustment. It is instructive to frame this distinction by contrasting the structural models that emerge from these two approaches.

---

important on an ongoing basis is order flow that conveys information about “portfolio shifts” that are not common knowledge. A recent event provides a sharp example. Major banks attribute the yen/dollar rate's drop from 145 to 115 in Fall 1998 to “the unwinding of positions by hedge funds that had borrowed in cheap yen to finance purchases of higher-yielding dollar assets” (*The Economist*, 10/10/98). This unwinding—and the selling of dollars that came with it—was forced by the scaling back of speculative leverage in the months following the Long Term Capital Management crisis. These trades were not common knowledge as they were occurring. (See also section 6 below, and Cai et al. 1999.)

<sup>6</sup> One might argue that expectations measurement cannot be driving the negative results of Meese and Rogoff because they use the driving variables' *realized* values. However, if the underlying macro model is incomplete, then realized values still produce an incorrect expectations measure.

### Structural Models: Macro Approach

Exchange-rate models within the macro approach are typically estimated at the monthly frequency. When estimated in changes they take the form:

$$(1) \quad \Delta p_t = f(\Delta i, \Delta m, \dots) + \varepsilon_t.$$

where  $\Delta p_t$  is the change in the log nominal exchange rate over the month (DM/\$). The driving variables in the function  $f(\Delta i, \Delta m, \dots)$  include changes in home and foreign nominal interest rates  $i$ , money supply  $m$ , and other macro determinants, denoted here by the ellipsis.<sup>7</sup> Changes in these public-information variables drive price—there is no role for order flow. Any incidental price effects from order flow that might arise are subsumed in the residual  $\varepsilon_t$ . These models are logically coherent and intuitively appealing. Unfortunately, they account for almost none of the monthly variation in floating exchange rates.

### Structural Models: Microstructure Approach

Equations of exchange-rate determination within the microstructure approach are derived from the optimization problem faced by price setters in the market—the dealers.<sup>8</sup> These models are all variations on the following specification:

$$(2) \quad \Delta p_t = g(\Delta x, \Delta I, \dots) + v_t$$

Now  $\Delta p_t$  is the DM/\$ rate change over two transactions, rather than over a month as in the macro models. The driving variables in the function  $g(\Delta x, \Delta I, \dots)$  include

---

<sup>7</sup> The precise list of determinants depends on the model. Meese and Rogoff (1983a) focus on three models in particular: the flexible-price monetary model, the sticky-price monetary model, and the sticky-price asset model. Here our interest is simply a broad-brush contrast between the macro and microstructure approaches. For specific models see Frenkel (1976), Dornbusch (1976), and Mussa (1976), among many others.

<sup>8</sup> Empirical work using structural micro models includes Glosten and Harris (1988), Madhavan and Smidt (1991), and Foster and Viswanathan (1993), all of which address the NYSE. Structural models in a multiple-dealer setting include Snell and Tonks (1995) for stocks, Lyons (1995) for currencies, and Vitale (1998) for bonds.

order flow  $\Delta x$ , the change in net dealer positions (or inventory)  $\Delta I$ , and other micro determinants, denoted by the ellipsis. Order flow can take both positive and negative values because the counterparty either purchases (+) at the dealer's offer or sells at the dealer's bid (-). Here we use the convention that a positive  $\Delta x$  is net dollar purchases, making the theoretical relation positive: net dollar purchases drive up the DM price of dollars. It is interesting to note that the residual in this case is the mirror image of the residual in equation 1: it subsumes any price changes due to determinants in the macro model  $f(\Delta i, \Delta m, \dots)$ , whereas the residual in equation 1 subsumes price changes due to determinants in the micro model  $g(\Delta x, \Delta I, \dots)$ .

Microstructure models predict a positive relation between  $\Delta p$  and  $\Delta x$  because order flow communicates non-public information, and once communicated, it is reflected in price. For example, if there is an agent who has superior information about the value of an asset, and that information advantage induces the agent to trade, then a dealer can learn from those trades (purchases indicate good news about the asset's value, and vice versa). Empirically, estimates of a relation between  $\Delta p$  and  $\Delta x$  at the transaction frequency are uniformly positive and significant. This is true for many different markets, including stocks, bonds, and foreign exchange.

The relation in microstructure models between  $\Delta p$  and  $\Delta I$  is not our focus in this paper, but let us clarify nonetheless. This relation is referred to as the inventory-control effect on price. The inventory-control effect arises when a dealer adjusts his price to control fluctuation in his inventory. For example, if a dealer has a larger long position than is desired, he may shade his bid and offer downward to induce a customer purchase, thereby reducing his position. This affects realized transaction prices, which accounts for the relation. (These idiosyncratic inventory effects on individual dealer prices do not arise in the model developed in the next section.)

### Spanning the Micro-Macro Divide

---

To span the divide between the micro and macro approaches, we develop a model with components from both:<sup>9</sup>

$$(3) \quad \Delta p_t = f(\Delta i, \dots) + g(\Delta x, \dots) + \eta_t.$$

The challenge is the frequency mismatch: transaction frequency for the micro models versus monthly frequency for the macro models. In the next section we develop a model in the spirit of equation (3). We estimate the model at the daily frequency by using micro determinants that are time-aggregated. We focus in particular on order flow  $\Delta x$ . Our time-aggregated measure spans a much longer period than is addressed elsewhere within empirical microstructure.

---

<sup>9</sup> Goldberg and Tenorio (1997) develop a model for the Russian ruble market that includes both macro and microstructure components. Osler's (1998) trading model includes macroeconomic "current account traders" who affect the exchange rate in flow equilibrium.

### 3. Portfolio Shifts Model

#### Overview

One source of exchange rate variation in the model is portfolio shifts on the part of the public. These portfolio shifts have two important features. First, they are not common knowledge as they occur. Second, they are large enough that clearing the market requires adjustment of the spot exchange rate.

The first feature—that portfolio shifts are not common knowledge—provides a role for order flow. At the beginning of each day, public portfolio shifts are manifested in orders in the foreign exchange market. These orders are not publicly observable. Dealers take the other side of these orders, and then trade among themselves during the day to share the resulting inventory risk. The market learns about the initial portfolio shifts by observing this interdealer trading activity. By the end of the day, the dealers' inventory risk is shared with the public.

The second important feature is that the initial portfolio shifts, once absorbed by the public at the end of the day, are large enough to move price. This requires that the public's demand for foreign-currency assets is less than perfectly elastic. If the public's demand is less than perfectly elastic, different-currency assets are imperfect substitutes, and price adjustment is required to clear the market.<sup>10</sup> In this sense, our model is in the spirit of the portfolio balance approach to exchange rates. In another sense, however, our model is very different from that earlier approach. Portfolio balance models are driven by changes in asset supply. Asset supply is constant in our model. Rather, our model identifies two distinct components on the demand side. The first is driven by innovations in public information (standard macro fundamentals). The second is driven by non-public information. This non-public information takes the form of

---

<sup>10</sup> For evidence of imperfect substitutability across stocks, see Scholes (1972), Shleifer (1986), Bagwell (1992), and Kaul et al. (2000), among others. Substitutability across currencies is likely to be lower than across same-currency stocks. Though direct evidence of this lower substitutability is lacking, some point to home bias in international portfolios as indirect evidence. Note, too, that the size of the order flows the DM/\$ spot market needs to absorb are on average more than 10,000 times those absorbed in a representative U.S. stock (e.g., the average daily volume on individual NYSE stocks in 1998 was about \$9 million, whereas the average daily volume in DM/\$ spot was about \$300 billion).

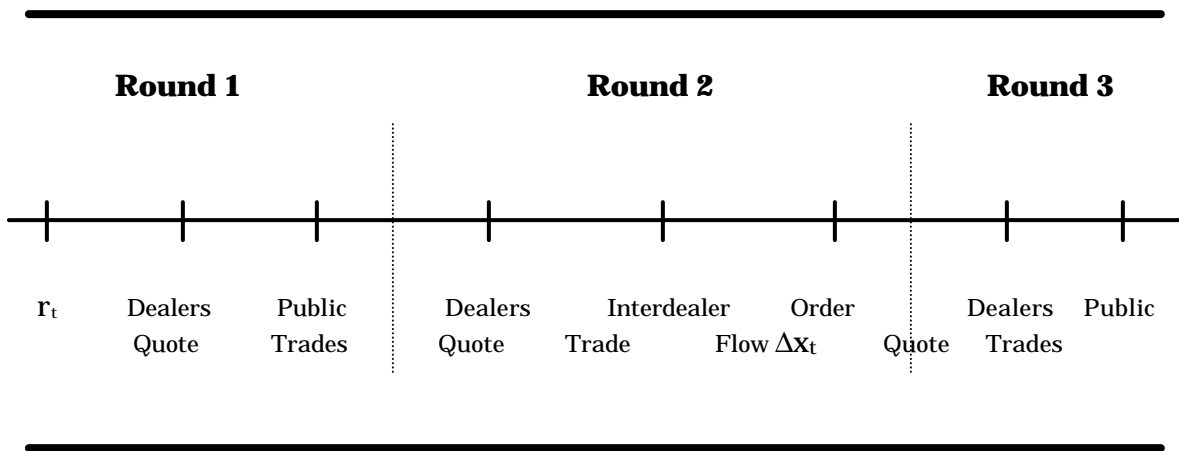
portfolio shifts. The model does not take a stand on the underlying determinants of these portfolio shifts (though we do address this issue in section 6).

Specifics

Consider an infinitely-lived, pure-exchange economy with two assets, one riskless, and one with stochastic payoffs representing foreign exchange.<sup>11</sup> The periodic payoff on foreign exchange, denoted  $R_t$ , is composed of a series of increments, so that  $R_t = \sum_{i=1}^t r_i$ . The increments  $r$  are i.i.d.  $\text{Normal}(0, \Sigma_r)$  and represent the flow of publicly available macroeconomic information over time (e.g., interest rates). Periodic payoffs are realized at the beginning of each period.

The foreign exchange market is organized as a decentralized dealership market with  $N$  dealers, indexed by  $i$ , and a continuum of non-dealer customers (the public), indexed by  $z \in [0, 1]$ . Within each period (day) there are three rounds of trading. In the first round dealers trade with the public. In the second round dealers trade among themselves to share the resulting inventory risk. In the third round dealers trade again with the public to share inventory risk more broadly. The timing within each day is:

**Daily Timing**



<sup>11</sup> An earlier version of this paper developed the portfolio shifts model in a finite-period setting. A subsequent paper by Rime (2000) extends the model to the infinite-period setting. The model presented below is a variation on the Rime extension.

Dealers and customers all have identical negative exponential utility defined over periodic wealth.

### Trading Round 1

At the beginning of each day  $t$ , all market participants observe the payoff increment  $r_t$  when the payoff  $R_t$  is realized. On the basis of this increment and other available information, each dealer simultaneously and independently quotes a scalar price to his customers at which he agrees to buy and sell any amount.<sup>12</sup> We denote this round-one price of dealer  $i$  as  $P_{i1}$ . (To ease the notational burden, we suppress the period subscript  $t$  when clarity permits.) Each dealer then receives a net customer-order realization  $c_{i1}$  that is executed at his quoted price  $P_{i1}$ , where  $c_{i1} < 0$  denotes a net customer sale (dealer  $i$  purchase). Each of these  $N$  customer-order realizations is distributed  $\text{Normal}(0, \Sigma_{c1})$ , and they are independent across dealers. (Think of these initial customer trades as assigned—or preferred—to a single dealer, resulting from bilateral customer relationships for example.) Customer orders are also distributed independently of the public-information increment  $r_t$ .<sup>13</sup> These orders represent portfolio shifts on the part of the non-dealer public. Their realizations are not publicly observable.

### Trading Round 2

Round 2 is the interdealer trading round. Each dealer simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell any amount. These interdealer quotes are observable and available to all dealers in the market. Each dealer then simultaneously and independently trades on other dealers' quotes. Orders at a given price are split evenly across any

---

<sup>12</sup> The sizes tradable at quoted prices in major FX markets are very large relative to other markets. At the time of our sample, the standard quote in DM/\$ was good for up to \$10 million, with a tiny bid-offer spread, typically less than four basis points. Introducing a bid-offer spread (or price schedule) in round one to endogenize the number of dealers is a straightforward—but distracting—extension of our model.

dealers quoting that price. Let  $T_{i2}$  denote the (net) interdealer trade initiated by dealer  $i$  in round two. At the close of round 2, all dealers observe the net interdealer order flow from that period:

$$(4) \quad \Delta x = \sum_{i=1}^N T_{i2}$$

Note that interdealer order flow is observed without noise, which maximizes the difference in transparency across trade types: customer-dealer trades are not publicly observed but interdealer trades are observed. In reality, FX trades between customers and dealers are not publicly observed. Though signals of interdealer order flow are publicly observed, it is not the case that these trades are observed without noise. Adding noise to Eq. (4), however, has no qualitative impact on our estimating equation, so we stick to this simpler specification.

### Trading Round 3

In round 3, dealers share overnight risk with the non-dealer public. Unlike round 1, the public's motive for trading in round 3 is non-stochastic and purely speculative. Initially, each dealer simultaneously and independently quotes a scalar price  $P_3$  at which he agrees to buy and sell any amount. These quotes are observable and available to the public at large.

The mass of customers on the interval  $[0,1]$  is large (in a convergence sense) relative to the  $N$  dealers. This implies that the dealers' capacity for bearing overnight risk is small relative to the public's capacity. Dealers therefore set prices so that the public willingly absorbs dealer inventory imbalances, and each dealer ends the day with no net position. These round-3 prices are conditioned on the round-2 interdealer order flow. The interdealer order flow informs dealers of the size of the total inventory that the public needs to absorb to achieve stock equilibrium.

---

<sup>13</sup> A natural extension of this specification is that customer orders reflect changing expectations of future  $r_t$ .

Knowing the size of the total inventory the public needs to absorb is not sufficient for determining round-3 prices. Dealers also need to know the risk-bearing capacity of the public. We assume it is less than infinite. Specifically, given negative exponential utility, the public's total demand for the risky asset in round-3, denoted  $c_3$ , is a linear function of the its expected return conditional on public information:

$$c_3 = \gamma(E[P_{3,t+1} + R_{t+1} | \Omega_3] - P_{3,t})$$

where the positive coefficient  $\gamma$  captures the aggregate risk-bearing capacity of the public, and  $\Omega_3$  is the public information available at the time of trading in round 3.

### Equilibrium

The dealer's problem is defined over four choice variables, the three scalar quotes  $P_{1t}$ ,  $P_{2t}$ , and  $P_{3t}$ , and the dealer's interdealer trade  $T_{12}$  (the latter being a component of the interdealer order flow  $\Delta x$ ). The appendix provides details of the model's solution. Here we provide some intuition. Consider the three scalar quotes. No arbitrage ensures that, within a given round, all dealers quote a common price. Given that all dealers quote a common price, this price is necessarily conditioned on common information only. Though the payoff increment  $r_t$  is common information at the beginning of round 1, order flow  $\Delta x_t$  is not observed until the end of round 2. The price for round-3 trading,  $P_3$ , therefore reflects the information in both  $r_t$  and  $\Delta x_t$ .

Whether  $\Delta x$  does influence price depends on whether it communicates any price-relevant information. The answer is yes. Understanding why requires a few steps. First, the appendix shows that it is optimal for each dealer to trade in round 2 according to the trading rule:

$$T_{12} = \alpha c_{1t}$$

with a constant coefficient  $\alpha$ . Thus, 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  $\Delta x = \sum_i T_{i2}$  at the end of round 2, they can infer the aggregate portfolio shift on the part of the public in round 1 (the sum of the  $N$  realizations of  $c_{i1}$ ). Dealers also know that the public needs to be induced to re-absorb this portfolio shift in round 3. This inducement requires a price adjustment. Hence the relation between the interdealer order flow and the subsequent price adjustment.

### The Pricing Relation

The appendix establishes that the change in price from the end of period  $t-1$  to the end of period  $t$  is:

$$(5) \quad \Delta P_t = \beta_1 r_t + \beta_2 \Delta x_t$$

where  $\beta_1$  and  $\beta_2$  are positive constants. The price change is proportional to the innovation in payoffs  $r_t$ : this is the perpetuity value of a permanent change in the periodic payoff. The  $\beta_2 \Delta x_t$  term is the portfolio shift term. This term reflects the price adjustment required to induce re-absorption of the public's portfolio shift from round 1. For intuition, note that  $\beta_2 \Delta x = \beta_2 \sum_i T_{i2} = \beta_2 \alpha \sum_i c_{i1}$ . The sum  $\sum_i c_{i1}$  is this total portfolio shift from round 1 (henceforth denoted  $c_1$ ). The public's total demand in round 3,  $c_3$ , is not perfectly elastic, and  $\beta_2$  insures that at the round-3 price  $c_3 + c_1 = 0$ .

### Empirical Implementation

Getting from equation (5) to an estimable model requires that we specialize the macro component of the model—the public-information increment  $r_t$ . We choose to specialize this component to capture changes in the nominal interest differential. That is, we define  $r_t \equiv \Delta(i_t - i_t^*)$ , where  $i_t$  is the nominal dollar interest rate and  $i_t^*$  is the nominal non-dollar interest rate (DM or Yen). This yields the following regression model:

$$(6) \quad \Delta P_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 \Delta x_t + \eta_t$$

Our choice of specialization has some advantages. First, this specification is consistent with monetary macro models in the sense that these models call for estimating innovations in  $\Delta P$  using the interest differential's change, not its level. (As a diagnostic, though, we also estimate the model using the level of the differential, a la Uncovered Interest Parity; see footnote 17.) Second, in asset-approach macro models like the Dornbusch (1976) overshooting model, innovations in the interest differential are the main engine of exchange rate variation.<sup>14</sup> Third, from a purely practical perspective, data on the interest differential are readily available at the daily frequency, which is certainly not the case for the other standard macro fundamentals (e.g., real output, nominal money supplies, etc.).

Naturally, this specification of our macro component of the model has some drawbacks. It is certainly true that, as a measure of variation in macro fundamentals, the interest differential is obviously incomplete. One can view it as an attempt to control for this key macro determinant in order to examine the importance of micro determinants. One should not view it as establishing a fair horse race between the micro and macro approaches.

#### **4. Data**

Our data set includes all the DM/\$ and ¥/\$ transactions that occurred from May 1 to August 31, 1996, on an interdealer trading system called Reuters Dealing 2000-1.<sup>15</sup> All trades on this system take the form of electronic bilateral

---

<sup>14</sup> Cheung and Chinn (1998) 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."

<sup>15</sup> In 1996, interdealer transactions account for about 75 percent of total volume in major spot markets (see BIS 1996). This 75 percent from interdealer trading breaks into two transaction types, direct and brokered. Direct trading accounts for about 60 percent of interdealer volume and brokered trading accounts for about 40 percent. Reuters Dealing 2000-1 is used for the direct portion. According to

conversations. A conversation is initiated when a dealer calls another dealer using the system to request a quote. If the calling dealer buys (sells) dollars then order flow from that trade is +1 (-1).<sup>16</sup> A dealer who has been called is expected to provide a fast two-way quote with a tight spread. Quotes are take-it-or-leave-it, and if not dealt or declined quickly (i.e., within seconds), the quoting dealer retracts the quote, ending the conversation. To settle disputes, Reuters keeps a record of all bilateral conversations. This record is the source of our data. (For more detail on the Dealing 2000-1 system, see Lyons 1995 and Evans 1997.)

Two features of the data are especially noteworthy. First, they provide transaction information for the whole interbank market over the full 24-hour trading day. This contrasts with earlier transaction data sets covering single dealers over some fraction of the trading day (Lyons 1995, Yao 1998, and Bjonnes and Rime 1998). Our comprehensive data set makes it possible, for the first time, to analyze order flow's role in price determination at the level of "the market." Though other data sets exist that cover multiple dealers, they include only brokered interdealer transactions (see Goodhart, Ito and Payne 1996, and Payne 1999). More important, these other data sets come from a particular brokered-trading system, one that accounts for a much smaller fraction of daily trading volume than the D2000-1 system covered by our data set. (There is also evidence that dealers attach more informational importance to direct interdealer order flow than to brokered interdealer order flow. See Bjonnes and Rime 1998.)

Second, our data cover a relatively long time span (four months) in comparison with other micro data sets covering the FX market. This is important because the longer time span allows us to address exchange-rate determination from more of an asset-pricing perspective than was possible with previous micro data spanning only days or weeks.

The three variables in our Portfolio Shifts model are measured as follows. The change in the spot rate (DM/\$ or ¥/\$),  $\Delta p_t$ , is the log change in the purchase

---

Reuters, over 90 percent of the world's direct interdealer transactions take place on this system (only dealers have access).

<sup>16</sup> Our data set does not identify the size of individual transactions. (See Jones, Kaul, and Lipson 1994 for evidence that trade size contains no information beyond that in the number of transactions.) Our

transaction price between 4 pm (GMT) on day  $t$  and 4 pm on day  $t-1$ . When a purchase transaction does not occur precisely at 4 pm, we use the subsequent purchase transaction (with roughly 1 million trades per day, the subsequent transaction is generally within a few seconds of 4 pm). When day  $t$  is a Monday, the day  $t-1$  price is the previous Friday's price. (Our dependent variable therefore spans the full four months of our sample, with no overnight or weekend breaks.) The daily order flow,  $\Delta x_t$ , is the difference between the number of buyer-initiated trades and the number of seller-initiated trades (in thousands), also measured from 4 pm (GMT) on day  $t-1$  to 4 pm on day  $t$  (negative sign denotes net dollar sales). The change in interest differential,  $\Delta(i_t - i_t^*)$ , is calculated from the daily overnight interest rates for the dollar, the deutschemark, and the yen (annual basis); the source is Datastream (typically measured at approximately 4 pm GMT).

---

data set does include total dollar volume over our sample, however, which allows us to calculate an average trade size; we use this below to interpret the estimated coefficients.

## 5. Empirical Results

Our empirical results are grouped in four sets. The first set addresses the in-sample fit of the portfolio shifts model. The second set addresses robustness issues. The third set addresses the direction of causality. The fourth set of results addresses the model's out-of-sample forecasting ability (in the spirit of Meese and Rogoff 1983a).

### 5.1 In-Sample Fit

Table 1 presents our estimates of the portfolio shifts model (equation 6) using daily data for the DM/\$ and ¥/\$ exchange rates. Specifically, we estimate the following regression:

$$(7) \quad \Delta p_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 \Delta x_t + \eta_t$$

where  $\Delta p_t$  is the change in the log spot rate (DM/\$ or ¥/\$) from the end of day  $t-1$  to the end of day  $t$ ,  $\Delta(i_t - i_t^*)$  is the change in the overnight interest differential from day  $t-1$  to day  $t$  (\* denotes DM or ¥), and  $\Delta x_t$  is the order flow from the end of day  $t-1$  to the end of day  $t$  (negative denotes net dollar sales).<sup>17</sup>

---

<sup>17</sup> Though the dependent variable in standard macro models is the change in the log spot rate, the dependent variable in the Portfolio Shifts model in equation (6) is the change in the spot rate without taking logs. These two measures for the dependent variable produce nearly identical results in all our tables ( $R^2$ s, coefficient significance, lack of autocorrelation, etc.). Here we present the log change results—equation 7—to make them directly comparable to previous macro specifications.

**Table 1: Model Estimates**

					Diagnostics		
		$\Delta(i_t - i_t^*)$	$\Delta x_t$	$i_{t-1} - i_{t-1}^*$	$R^2$	Serial	Hetero
DM	I	0.51 (0.26)	2.14 (0.29)		0.64	0.77 0.40	0.07 0.02
	II		2.15 (0.29)		0.63	0.73 0.45	0.05 0.03
	III	0.62 (0.77)			0.01	0.78 0.77	0.92 0.99
	IV		2.15 (0.29)	0.02 (0.01)	0.64	0.49 0.43	0.17 0.01
	V			0.02 (0.02)	0.00	0.04 0.24	0.83 0.98
Yen	I	2.47 (0.92)	2.86 (0.36)		0.46	0.06 0.44	0.92 0.74
	II		2.61 (0.36)		0.40	0.19 0.33	0.60 0.83
	III	0.57 (1.20)			0.00	0.85 0.81	0.13 0.67
	IV		2.78 (0.38)	0.02 (0.01)	0.42	0.00 0.03	0.66 0.72
	V			-0.01 (0.01)	0.00	0.12 0.46	0.18 0.79

The table reports OLS estimates of the Portfolio Shifts Model (specification I):

$$\Delta p_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 \Delta x_t + \eta_t$$

and four alternatives (specifications II – V). The dependent variable  $\Delta p_t$  is the change in the log spot exchange rate from 4 pm GMT on day t-1 to 4 pm 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  $\Delta x_t$  is interdealer order flow between 4 pm GMT on day t-1 and 4 pm GMT on day t (negative for net dollar sales, in thousands). Specifications IV and V replace  $\Delta(i_t - i_t^*)$  with the interest differential at 4 pm on day t-1. Standard errors are shown in parentheses (corrected for heteroskedasticity in the case of the DM). The sample spans four months (May 1 to August 31, 1996), which is 89 trading days. The Serial column presents the p-value of Breusch-Godfrey LM tests for residual serial correlation, first-order in the top row and fifth-order (one week) in

the bottom row. The Hetero column presents the p-values of Engle (1982) LM tests for ARCH in the residuals, first-order in the top row and fifth-order in the bottom row.

The coefficient on order flow  $\Delta x_t$  is correctly signed and significant, with t-statistics above 5 for both currency pairs. The positive sign indicates that net dollar purchases – a positive  $\Delta x_t$  – lead to a higher DM price of dollars. The traditional macro-fundamental – the interest differential – is correctly signed and significant. (The positive sign arises in the sticky-price monetary model, for example, because an increase in the dollar interest rate  $i$  requires immediate dollar appreciation – increase in DM/\$ – to make room for UIP-induced dollar depreciation.) The DM equation shows some evidence of heteroskedasticity, so we correct the standard errors in that case using a heteroskedasticity-consistent covariance matrix (White 1980).

The explanatory power of these regressions is due to order flow  $\Delta x_t$ : regressing  $\Delta p_t$  only on  $\Delta(i_t - i_t^*)$  produces an  $R^2$  statistic at 1 percent or lower for both currency pairs, and coefficients on  $\Delta(i_t - i_t^*)$  that are insignificant at the 5 percent level. (The same story emerges when the *level* of the interest differential is included rather than the change – specifications IV and V.)<sup>18</sup> Adding order flow is what increases the fit so dramatically:  $R^2$  statistics of 64 percent and 46 percent for the DM and Yen equations (specification I), respectively.

The size of the order flow coefficient is consistent with past estimates from single-dealer data. The coefficient of 2.1 in the DM equation implies that a day with 1000 more dollar purchases than sales increases the DM price of a dollar by 2.1 percent. Given an average trade size in our sample of \$3.9 million, \$1 billion of net dollar purchases increases the DM price by 0.54% (= 2.1/3.9), or, at a spot rate of 1.5 DM/\$, 0.8 pfennig. For comparison, at the single-dealer level Lyons (1995) finds that information asymmetry induces the dealer he tracks to increase price by 0.01 of a pfennig (0.0001 DM) for every incoming buy order of \$10 million. That is 1 pfennig per \$1 billion, quite close to the 0.8 pfennig found here. (Though

---

<sup>18</sup> There is a vast empirical literature that attempts to increase the explanatory power of interest rates in exchange rate equations by introducing interest rates as separate regressors, introducing non-linear specifications, etc. This literature has not been successful, so we do not pursue this line here. Note that the lack of explanatory power from traditional fundamentals is not unique to exchange rate economics: Roll (1988) produces  $R^2$ s of only 20% using traditional equity fundamentals to account for daily stock returns, a result he describes as a “significant challenge to our science.”

linearly extrapolating the single-dealer estimate is not an accurate description of individual dealers, it may be a good description of the market's aggregate elasticity.)<sup>19</sup>

Order flow's ability to account for the full four months of exchange rate variation is surprising, not only from the perspective of macro exchange rate economics, but also from the perspective of microstructure finance. Recall from section 2 that structural models within microstructure finance are typically estimated at the transaction frequency—they make no attempt to account for prices over the full 24-hour day. Our regression is at the daily frequency. One might have conjectured that the net impact of order flow over the day would be zero (each day accounts for about one million transactions). This conjecture would be consistent with a belief that cumulative order flow mean-reverts rapidly (e.g., within a day). But rapid mean reversion is clearly not the behavior displayed by cumulative order flow in Figure 1. This lack of mean reversion provides some room for the lower frequency relation we find here.

The lack of strong mean reversion in our measured order flow deserves further attention, particularly considering that half-lives of individual dealer positions can be as short as 10 minutes (Lyons 1998). The key lies in recognizing that our measure of order flow reflects interdealer trading, not customer-dealer trading. Consider a scenario that illustrates why our measure in Figure 1 can be so persistent. (Recall that Figure 1 displays cumulative order flow, defined as the sum of interdealer order flow,  $\Delta x_t$ , from 0 to  $t$ .) Starting the scenario from  $x_t=0$ , an initial customer sale does not move  $x_t$  from zero because  $x_t$  measures interdealer order flow only. After the customer sale, then when dealer  $i$  unloads the position by selling to another dealer  $j$ ,  $x_t$  drops to  $-1$ . A subsequent sale by dealer  $j$  to another dealer, dealer  $k$ , reduces  $x_t$  further to  $-2$ .<sup>20</sup> If a customer happens to buy dealer  $k$ 's position from him, then  $x_t$  remains at  $-2$ . In this simple scenario, order

---

<sup>19</sup> To check robustness, we examine several obvious variations on the model. For example, we also include a constant in the regression, even though the model does not call for one. The constant is insignificant for both currencies. Also, recall that our price series is measured from purchase transactions only. Results using 4 pm sale prices are identical. We address additional robustness issues in the next subsection.

<sup>20</sup> This repeated passing of dealer positions in the foreign exchange market is referred to as the "hot potato" phenomenon. See Burnham (1991) and Lyons (1997).

flow measured only from trades between customers and dealers would have reverted to zero—the concluding customer trade offsets the initiating customer trade. The interdealer order flow, however, does not revert to zero. Note, too, that this difference in the persistence of the two order-flow measures—customer-dealer versus interdealer—is also a property of the Portfolio Shifts model (indeed, clarifying this point is one of the model’s contributions). In the Portfolio Shifts model, customer order flow in round 3 always offsets the customer order flow in round 1. But the interdealer order flow, which only arises in round 2, does not net to zero. This non-zero  $\Delta x_t$  serves as a carrier of value in our estimating equation.

## 5.2 Robustness

In this section we address three robustness issues beyond those examined in the previous section. They correspond to the following three questions: (1) Might the order-flow/price relation be non-linear? (2) Does the relation depend on the gross level of activity? and (3) Does the relation depend on day of the week?

### Might the order-flow/price relation be non-linear?

The linearity of our Portfolio Shifts specification depends crucially on several simplifying assumptions, some of which are rather strong on empirical grounds. It is therefore natural to investigate whether non-linearities or asymmetries might be present.<sup>21</sup> A simple first test is to add a squared order-flow term to the baseline specification. The squared order-flow term is insignificant in both equations. We also test whether the coefficient on order flow is piece-wise linear, with a kink at  $\Delta x_t=0$ . If true, this means that buying pressure and selling pressure are not symmetric. A Wald test that the two slope coefficients are equal cannot be rejected for the DM equation. There is some evidence of different slopes in the Yen equation however: the test is rejected at the 4 percent marginal significance level. In that case, the point estimates show a greater sensitivity of price to order

---

<sup>21</sup> We pursue these (simple) non-linear specifications with the comfort that outliers are not driving our results—a fact that is manifest from Figure 1.

flow in the downward direction, though both estimates remain positive and significant.

#### Does the order-flow/price relation depend on the gross level of activity?

Another natural concern is whether the order-flow/price relation in Table 1 is state contingent in some way, perhaps depending on the market's overall activity level. Our data set provides a convenient measure of overall activity, namely the total number of transactions. As a simple test, we partition our sample of trading days into quartiles, from days with the fewest transactions to days with the most transactions. We then estimate separate order-flow coefficients for each of these four sample partitions. In both the DM and Yen equations, all four of the order-flow coefficients are positive. In the DM equation, the coefficients are slightly U-shaped (from fewest transactions to most, the point estimates for  $\beta_2$  are 2.7, 2.0, 1.9, and 3.3). In the Yen equation, the coefficients are monotonically increasing (from fewest transactions to most, the point estimates for  $\beta_2$  are 1.0, 1.1, 3.5, and 4.1).

In terms of theory, this result for the Yen is consistent with the “event-uncertainty” model of Easley and O’Hara (1992), but the DM result is not. The event-uncertainty model predicts that trades are more informative when trading intensity is higher. Key to understanding their result is that in their model, new information may not exist. If there is trading at time  $t$ , then a rational dealer raises her conditional probability that an information event has occurred, and lowers the probability of the “no-information” event. The upshot is that trades occurring when trading intensity is high induce a larger update in beliefs, and therefore a larger adjustment in price.

#### Does the order-flow/price relation depend on day of the week?

Another state-contingency that warrants attention is day-of-the-week effects.<sup>22</sup> To test whether day-of-the-week matters, we partition our sample into five sub-samples, one for each weekday (recall that weekends are subsumed in our Friday-to-Monday observations). In both the DM and Yen equations, all five of the resulting order-flow coefficients are positive. In the DM equation, the Tuesday coefficient is the largest, and the Wednesday coefficient is the smallest. The Yen equation also shows that Tuesday's coefficient is the largest, but in this case the Monday coefficient is smallest. More important, a Wald test that the coefficients are equal across the five days cannot be rejected at the 5 percent level in either equation (though in the case of the Yen, it can be rejected at the 10 percent level).

### 5.3 Causality

Under our model's null hypothesis, causality runs strictly from order flow to price. Accordingly, under the null, our estimation is not subject to simultaneity bias. (We are not simply "regressing price on quantity," as in the classic supply-demand identification problem. Quantity—i.e., volume—and order flow are fundamentally different concepts.) Within microstructure theory more broadly, this direction of causality is the norm: it holds in all the canonical models (Glosten and Milgrom 1985, Kyle 1985, Stoll 1978, Amihud and Mendelson 1980), despite the fact that price and order flow are determined simultaneously. The important point in these models is that price innovations are a function of order flow innovations, not the other way around.<sup>23</sup> That said, alternative hypotheses do exist under which causality is reversed. The following taxonomy frames the causality issue, and identifies specific alternatives under which causality is

---

<sup>22</sup> In terms of theory, the model of Foster and Viswanathan (1990) is a workhorse for specifying day-of-the-week effects. In their model, there is periodic variation in the information advantage of the informed trader. This advantage is assumed to grow over periods of market closure, in particular, over weekends, making order flow on Monday particularly potent.

<sup>23</sup> Put differently, order flow in these models is a proximate cause. The underlying driver of order flow is non-public information (information about uncertain demands, information about payoffs, etc.). Order flow is the channel through which this type of information is impounded in price.

reversed, so that the merits of these alternatives can be judged in a disciplined way.

### Theoretical Overview

The timing of the order-flow/price relation admits three possibilities, depending on whether order flow precedes, is concurrent with, or lags price adjustment. We shall refer to these three timing hypotheses as the *Anticipation* hypothesis, the *Pressure* hypothesis, and the *Feedback* hypothesis, respectively. As we shall see, reverse causality becomes a clear problem only under the Feedback hypothesis, and in that case only under certain conditions.

Within each of the three hypotheses—Anticipation, Pressure, and Feedback—there are also variations. Under the Anticipation hypothesis, for example, order flow can precede price adjustment because prices adjust fully only after order flow is commonly observed—in low-transparency markets like foreign exchange, order flow is not commonly observed when it occurs (Lyons 1996). Order flow might also precede price because price adjusts only after some piece of news anticipated by order flow is commonly observed (e.g., the short-lived private information in Foster and Viswanathan 1990). Under the Pressure hypothesis the two main variations correspond to microstructure theory’s canonical model types—information models and inventory models. In information models, observing order flow provides information about payoffs (Glosten and Milgrom 1985, Kyle 1985). In inventory models, order flow alters equilibrium risk premia (Stoll 1978, Ho and Stoll 1981).<sup>24</sup> Under the Feedback hypothesis, order flow lags price because of feedback trading. Negative-feedback trading is systematic selling in response to price increases, and buying in response to price decreases (e.g., Friedman’s celebrated “stabilizing speculators”). Positive-feedback trading is the reverse. Variations on the Feedback hypothesis are distinguished by whether this feedback trading is rational (an optimal response

---

<sup>24</sup> Within this inventory-model category, there is an additional distinction between price effects that arise at the marketmaker level (canonical inventory models) and price effects that arise at the marketwide level, due to imperfect substitutability (e.g., our Portfolio Shifts model). In the case of price effects at the marketmaker level, these effects are often modeled as changing risk premia. But

to return autocorrelation) or behavioral, meaning that it arises from systematic decision bias (DeLong et al. 1990, Jegadeesh and Titman 1993, Grinblatt et al. 1995).

Under the Pressure hypothesis, causality runs from order flow to price, despite their concurrent realization.<sup>25</sup> Causality also runs from order flow to price under most variations of the Anticipation hypothesis. This is true, for example, under the variation noted above where order flow is not commonly observed, and therefore affects price with a delay. One way to examine this variation of the Anticipation hypothesis is by introducing lagged order flow to our Portfolio Shifts model. Rows 1 and 3 of Table 2 present the results of this regression: lagged order flow is insignificant. At the daily frequency, lagged order flow is already embedded in price.<sup>26</sup>

Under the Feedback hypothesis, causality can go in reverse, that is, from price to order flow.<sup>27</sup> With negative-feedback trading, a la Friedman's stabilizing speculators for example, the direction of causality is reversed. However, in this case one would expect an order-flow/price relation that is *negative*; we find a positive relation, so this type of feedback trading cannot account for our results. If instead positive-feedback trading were present and significant, then one would expect order flow in period  $t$  to be positively related to the price change in period  $t-1$ . In daily data, this corresponds to  $\Delta x_t$  being explained, at least in part, by  $\Delta p_{t-1}$ . If our order-flow coefficient in Table 1 is picking up this daily-frequency positive feedback, then including lagged price change  $\Delta p_{t-1}$  in the Portfolio-Shifts regression should weaken, if not eliminate, the significance of order flow.

---

sometimes, largely for technical convenience, models are specified with risk-neutral marketmakers who face some generic "inventory holding cost."

<sup>25</sup> This does not imply that price cannot influence order flow. Price does influence order flow in micro-structure models (both for the usual downward sloping demand reason, and because agents learn from price). It is still the case that—in equilibrium—price innovations are functions of order flow innovations, not vice versa. Our Portfolio Shifts model is a case in point.

<sup>26</sup> As another check along these lines, we also decompose contemporaneous order flow into expected and unexpected components (by projecting it on past order flow). In our model, all order flow  $\Delta x$  is unexpected, but this need not be the case in the data. We find, as the model predicts, that order flow's explanatory power comes from its unexpected component.

<sup>27</sup> Note that the Feedback hypothesis does not imply that causality runs wholly in reverse. For example, the Feedback hypothesis does not rule out that feedback trading can affect prices.

Rows 2 and 4 of Table 2 present the results of this regression. Past price change does not reduce the significance of order flow, and is itself insignificant. These results run counter to the positive-feedback hypothesis at the daily frequency.

**Table 2**

Portfolio shifts model: Alternative specifications

$$\Delta p_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 \Delta x_t + \beta_3 \Delta x_{t-1} + \eta_t$$

$$\Delta p_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 \Delta x_t + \beta_3 \Delta p_{t-1} + \eta_t$$

	D( $i_t - i_t^*$ )	D $x_t$	D $x_{t-1}$	D $p_{t-1}$	Diagnostics		
					R <sup>2</sup>	Serial	Hetero
<b>DM</b>	0.40 (0.36)	2.16 (0.18)	0.29 (0.19)		0.65	0.76 0.48	0.39 0.03
	0.42 (0.35)	2.17 (0.18)		0.11 0.07	0.66	0.60 0.44	0.38 0.01
<b>Yen</b>	2.48 (0.91)	2.90 (0.36)	-0.20 (0.35)		0.47	0.07 0.41	0.55 0.84
	2.64 (0.91)	2.98 (0.36)		-0.13 (0.09)	0.48	0.21 0.63	0.52 0.81

The dependent variable  $\Delta p_t$  is the change in the log spot exchange rate from 4 pm GMT on day t-1 to 4 pm 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  $\Delta x_t$  is interdealer order flow between 4 pm GMT on day t-1 and 4 pm GMT on day t (negative for net dollar sales, in thousands). Estimated using OLS. Standard errors are shown in parentheses (corrected for heteroskedasticity in the case of the DM). The sample spans four months (May 1 to August 31, 1996), which is 89 trading days. The Serial column presents the p-value of a chi-squared test for residual serial correlation, first-order in the top row and fifth-order (one week) in the bottom row. The Hetero column presents the p-value of a chi-squared test for ARCH in the residuals, first-order in the top row and fifth-order in the bottom row.

## Empirical Reality

The theoretical overview above cannot resolve the fact that, in daily data, all three hypotheses—Anticipation, Pressure, and Feedback—may produce a relationship that appears contemporaneous. A concern therefore remains that the positive coefficient on order flow in Table 1 might be the result of positive-feedback trading that occurs intraday. We offer two additional types of evidence against this alternative interpretation of our results. The first is a set of three arguments why intraday positive feedback is an unappealing hypothesis in this context. The second is an explicit analysis of bias, designed to calibrate how extreme the positive feedback would have to be to account for the key moments of our data. (These moments include, but are not limited to, the moments that produce our order-flow coefficient in Table 1.)

There are three reasons, a priori, why the hypothesis of intraday positive-feedback trading is unappealing. First, direct empirical evidence does not support it: there is no evidence in the current literature of positive-feedback trading in the foreign exchange market. Second, if systematic positive-feedback trading were present, it would be irrational: intraday studies using transactions data find no evidence of the positive autocorrelation in price that would make positive-feedback an optimal response (Goodhart, Ito, and Payne 1996). Third, the fallback possibility of irrational positive-feedback trading is difficult to defend. Recall that the order flow we measure is interdealer order flow. Though systematic feedback trading of a behavioral nature (i.e., not fully rational) might be a good description of some market participants, dealers are among the most sophisticated participants in this market.

## Bias Analysis

To close this section on causality, let us consider what it would take for positive-feedback trading to account for our results. Specifically, suppose intraday positive-feedback trading is present—Under what conditions could it account for the key moments of our data? These moments include, but are not limited to, the moments that produce our positive order-flow coefficient in Table 1. We show below that these conditions are rather extreme. In fact, through a broad range of

underlying parameter values, feedback trading would have to be *negative* to account for the key moments of our data.

We start by decomposing measured order flow  $\Delta x_t$  into two components:

$$(8) \quad \Delta x_t = \Delta x_{t1} + \Delta x_{t2}$$

where  $\Delta x_{t1}$  denotes exogenous order flow from portfolio shifts (a la our model), with variance equal to  $\Sigma_{x1}$ , and  $\Delta x_{t2}$  denotes order flow due to feedback trading, where

$$(9) \quad \Delta x_{t2} = \gamma \Delta p_t$$

Suppose the true structural model can be written as:

$$(10) \quad \Delta p_t = \alpha \Delta x_{t1} + \varepsilon_t$$

where  $\varepsilon_t$  represents common-knowledge (CK) news, and  $\varepsilon_t$  is iid with variance  $\Sigma_\varepsilon$ . By CK news we mean that both the information *and* its implication for equilibrium price is common knowledge. If both conditions are not met, then order flow will convey information about market-clearing prices (recall the discussion in the introduction). If feedback trading is present ( $\gamma \neq 0$ ), then  $\alpha$  will be a reduced form coefficient that depends on  $\gamma$ . Note that under these circumstances, equation (10) is a valid reduced-form equation that could be estimated by OLS if one had data on  $\Delta x_{t1}$ .

With data on  $\Delta x_t$  and  $\Delta p_t$  only, suppose we estimate

$$(11) \quad \Delta p_t = \beta \Delta x_t + \varepsilon_t$$

If  $\gamma \neq 0$ , our estimates of  $\beta$  will suffer from simultaneity bias. To evaluate the size of this bias, consider the implications of equations (8) through (10) for the moments:

$$\beta = \text{Cov}(\Delta p_t, \Delta x_t) / \text{Var}(\Delta x_t)$$

$$\delta = \text{Var}(\Delta p_t) / \text{Var}(\Delta x_t)$$

From equations (8) through (10) we know that:

$$\Delta x_t = (1+\gamma\alpha)(\Delta x_{t1}) + \gamma\varepsilon_t$$

Solving for expressions for  $\text{Cov}(\Delta p_t, \Delta x_t)$ ,  $\text{Var}(\Delta p_t)$ , and  $\text{Var}(\Delta x_t)$ , we can write:

$$\beta = \text{Cov}(\Delta p_t, \Delta x_t) / \text{Var}(\Delta x_t) = (\alpha(1+\gamma\alpha)\Sigma_{x1} + \gamma\Sigma_\varepsilon) / ((1+\gamma\alpha)^2\Sigma_{x1} + \gamma^2\Sigma_\varepsilon)$$

$$\delta = \text{Var}(\Delta p_t) / \text{Var}(\Delta x_t) = (\alpha^2\Sigma_{x1} + \Sigma_\varepsilon) / ((1+\gamma\alpha)^2\Sigma_{x1} + \gamma^2\Sigma_\varepsilon)$$

Now, define an additional parameter:

$$\phi = \Sigma_\varepsilon / \Sigma_{x1}$$

This parameter represents the ratio of CK news to order-flow news. With this parameter  $\phi$  we can rewrite the key coefficients as:

$$\beta = (\alpha(1+\gamma\alpha) + \gamma\phi) / ((1+\gamma\alpha)^2 + \gamma^2\phi)$$

$$\delta = (\alpha^2 + \phi) / ((1+\gamma\alpha)^2 + \gamma^2\phi)$$

Using the sample moments for  $\text{Cov}(\Delta p_t, \Delta x_t)$ ,  $\text{Var}(\Delta p_t)$ , and  $\text{Var}(\Delta x_t)$ , we can solve for the implied values of the  $\alpha$  and  $\gamma$  for given values of  $\phi$ . The following table presents these implied values of  $\alpha$  and  $\gamma$ .

Note that even for values of  $\phi$  above 2, the feedback trading needed to generate our results is actually *negative*. Note too that the parameter  $\alpha$ —the order-flow-causes-price parameter—is not driven to zero until  $\phi$  reaches values well above 10. To invalidate our causality interpretation, then, CK news would have to be one to two orders of magnitude more important than order-flow news. In our judgment this is too extreme to be compelling.

To close this section on causality, it is not enough for the skeptical reader to assert simply that order flow and price are both “endogenous,” or that we are merely observing a “simultaneous relationship.” These points are true. But they are also true within the body of microstructure theory reviewed above. And within that body of theory, price innovations are still driven by order flow innovations. This section is our effort to bring some disciplined thinking to an otherwise superficial debate.

**Table 3**

Bias Analysis

$$\Delta x_{t2} = \gamma \Delta p_t$$

$$\Delta p_t = \alpha \Delta x_{t1} + \varepsilon_t$$

	$f = S_c / S_{x1}$	a	g
<b>DM</b>	0	2.4	-0.05
	0.1	1.2	-0.51
	1	1.9	-0.12
	2	2.1	-0.03
	10	2.0	0.16
	100	0.0	0.36
<b>Yen</b>	0	2.4	-0.18
	0.1	1.3	-0.58
	1	2.2	-0.23
	2	2.4	-0.15
	10	2.8	-0.02
	100	0.0	0.21

The table shows the values for the parameters  $\alpha$  (order-flow-causes-price) and  $\gamma$  (price-causes-order-flow) implied by the sample moments and given values for the parameter  $\phi$ . The parameter  $\phi$  is the ratio of common-knowledge news to order-flow news.

## 5.4 Out-of-Sample Forecasts

To control for the myriad specification searches conducted by empiricists, a tradition within exchange rate economics has been to augment in-sample model estimates with estimates of models' out-of-sample forecasting ability. Accordingly, we present results along these lines as well. The original work by Meese and Rogoff (1983a) examines forecasts from 1 to 12 months. Our four-month sample does not provide sufficient power to forecast at these horizons. Our horizons range instead from one day to two weeks. The Meese-Rogoff puzzle is why short-horizon forecasts do so poorly, and our focus is definitely on the short end (though not so short as to render the horizon irrelevant from a macro perspective).

Table 4 shows that the portfolio shifts model produces better forecasts than the random-walk (RW) model. The forecasts from our model are derived from recursive estimates that begin with the first 39 days of the sample. Like the Meese-Rogoff forecasts, our forecasts are based on realized values of the future forcing variables—in our case, realized values of order flow and changes in the interest differential. (Thus, they are not truly “out-of-sample forecasts.” We chose to stick with the Meese-Rogoff terminology.) The resulting root mean squared error (RMSE) is 30 to 40 percent lower than that for the random walk.

Note that our 89-day sample has very low power at the one- and two-week horizons. Even though our model's RMSE estimates are roughly 35 percent lower at these horizons, their out-performance is not statistically significant. With a sample this size, the one-week forecast would need to cut the RW model's RMSE by about 50 percent to reach the 5 percent significance level. (To see this, note that for the DM a two-standard-error difference at the one-week horizon is about 0.49, which is roughly half of the RW model's RMSE of 0.98). The two-week forecast would have to cut the RW model's forecast error by some 54 percent. More powerful tests at these longer horizons will have to wait for longer spans of transaction data.

**Table 4**

Out-of-sample forecasts errors  
 Root mean squared errors ( $\times 100$ )

		<u>RW</u>	<u>Portfolio Shifts</u>	<u>Difference</u>
<b>DM</b>	<u>Horizon</u>			
	1 day	0.44	0.29	0.15 (0.033)
	1 week	0.98	0.63	0.35 (0.245)
	2 weeks	1.56	0.96	0.60 (0.419)
<b>Yen</b>	1 day	0.40	0.32	0.08 (0.040)
	1 week	0.98	0.64	0.33 (0.239)
	2 weeks	1.34	0.90	0.45 (0.389)

The RW column reports the RMSE for the random walk model (approximately in percentage terms). The Portfolio Shifts column reports the RMSE for the model in equation (6). The Portfolio Shifts forecasts are based on realized values of the forcing variables. The forecasts are derived from recursive model estimates starting with the first 39 days of the sample. The Difference column reports the difference in the two RMSE estimates, and, in parentheses, the standard errors for the difference, calculated as in Meese and Rogoff (1988).

## 6. Discussion

The relation in our model between exchange rates and order flow is not easy to reconcile with the traditional macro approach. Under the traditional approach, information is common knowledge and is therefore impounded in exchange rates without the need for order flow. This apparent contradiction can

be resolved if either: (1) some information relevant for exchange rate determination is not common knowledge; or (2) some aspect of the mapping from information to equilibrium prices is not common knowledge. If either is relaxed then order flow conveys information about market-clearing prices.

Our portfolio shifts model resolves the contradiction by introducing information that is not common knowledge—information about shifts in public demand for foreign-currency assets. At a microeconomic level, dealers learn about these shifts in real time by observing order flow. As the dealers learn, they quote prices that reflect this information. At a macroeconomic level, these shifts are difficult to observe empirically. Indeed, the concept of order flow is not recognized within the international macro literature. (Transactions, if they occur at all, are strictly symmetric, and therefore cannot be signed to reflect net buying/selling pressure.)

If order flow drives exchange rates, then what drives order flow? From a valuation perspective, there are two distinct views. The first view is that order flow reflects new information about *valuation numerators* (i.e., future dividends in a dividend-discount model, which in foreign exchange take the form of future interest differentials). The second view is that order flow reflects new information about *valuation denominators* (i.e., anything that affects discount rates). Our portfolio shifts model is an example of the latter: order flow is unrelated to valuation numerators—the future  $r_t$ . This type of order flow can be rationalized with, for example, time-varying risk tolerance, time-varying hedging demands, or time-vary transactions demands. (In presenting the model, we did not take a stand on a specific rationalization.) An example consistent with the valuation-numerators view is the “proxy-for-expectations” idea introduced in the introduction. That is, an important source of innovations in exchange rates is innovations in expected future fundamentals, and in real time these may be well proxied by order flow.

Note that separating valuation numerators from valuation denominators has implications for the concept of “fundamentals.” Order flow that reflects information about valuation numerators—like expectations of future interest rates—is in keeping with traditional definitions of exchange-rate fundamentals.

But order flow that reflects valuation denominators encompasses nontraditional exchange-rate determinants, calling, perhaps for a broader definition. In any event, exploring these links to deeper determinants is a natural topic for future research. This will surely require a retreat back into intraday data.<sup>28</sup>

### The Practitioner View versus the Academic View

Another perspective on order flow emerges from the difference between academic and practitioner views on price determination. Practitioners often explain price increases with the familiar reasoning that “there were more buyers than sellers.” To most economists, this reasoning is tantamount to “price had to rise to balance demand and supply.” But these phrases may not be equivalent. For economists, the phrase “price had to rise to balance demand and supply” calls to mind the Walrasian auctioneer. The Walrasian auctioneer collects “preliminary” orders and uses them to find the market-clearing price. Importantly, the auctioneer’s price adjustment is immediate—no trading occurs in the transition. (In a rational-expectations model of trading, for example, this is manifested in all orders being conditioned on the market-clearing price.)

Many practitioners have a different model in mind. In the practitioner model there is a dealer instead of an abstract auctioneer. The dealer acts as a buffer between buyers and sellers. The orders the dealer collects are actual orders, rather than preliminary orders, so trading does occur in the transition to the new price. The dealer determines new prices from the new information about demand and supply that becomes available.

Can the practitioner model be rationalized? At first blush, it appears that trades are taking place out of equilibrium, implying irrational behavior. But this misses an important piece of the puzzle. Whether these trades are out-of-equilibrium depends on the information available to the dealer. If the dealer knows at the outset that there are more buyers than sellers (eventually pushing

---

<sup>28</sup> The role of macro announcements in determining order flow warrants exploring. This, too, requires the use of intraday data. A second possible use of macro announcements is to introduce them directly into our Portfolio Shifts specification, even at the daily frequency. This tack is not likely to be fruitful: there is a long literature showing that macro announcements are unable to account for exchange rate first moments (as opposed to second moments; see Andersen and Bollerslev 1998).

price up), then it may not be optimal to sell at a low interim price. If the buyer/seller imbalance is not known, however, then rational trades can occur through the transition. In this case, the dealer cannot set price conditional on all the information available to the Walrasian auctioneer. This is precisely the story developed in canonical microstructure models (Glosten and Milgrom 1985). Trading that would be irrational if the dealer could condition on the auctioneer's information can be rationalized in models with more limited (and realistic) conditioning information.

### Relation Between Our Model and the Flow Approach to Exchange Rates

Consider the relation between our model, with its emphasis on order flow, and the traditional “flow approach” to exchange rates. Is our approach just a return to the earlier flow approach? Despite their apparent similarity, the two approaches are distinct and, in fact, fundamentally different.

A key feature of our model is that order flow plays two roles. First, holding beliefs constant, order flow affects price through the traditional process of market clearing. Second, order flow also alters beliefs because it conveys information that is not yet common knowledge. That is:

$$\text{Price} = P(\Delta x, B(\Delta x, \dots), \dots)$$

Price  $P$  thus depends both directly and indirectly on order flow,  $\Delta x$ , where the indirect effect is via beliefs  $B$ . Early attempts to analyze equilibrium with differentially informed individuals ignored the information role—the effect of order flow on beliefs. Since the advent of rational expectations, models that ignore this information effect from order flow are viewed as less compelling.

This is the essential difference between the flow approach to exchange rates and the microstructure approach. Under the flow approach, order flow communicates no information back to individuals regarding others' views/information. All information is common knowledge, so there is no information that needs aggregating. Under the microstructure approach, order flow does

communicate information that is not common knowledge. This information needs to be aggregated by the market, and microstructure theory describes how that aggregation is achieved, depending on the underlying information type.

## 7. Conclusion

This paper presents a model of exchange rate determination of a new kind. Instead of relying exclusively on macroeconomic determinants, we draw on determinants from the field of microstructure. In particular, we focus on order flow, the variable within microstructure that is—both theoretically and empirically—the driver of price.<sup>29</sup> This is a radical departure from traditional approaches to exchange rate determination. Traditional approaches, with their common-knowledge environments, admit no role for information aggregation. Our findings suggest instead that the problem this market solves is indeed one of information aggregation.

Our Portfolio Shifts model provides an explicit characterization of this information aggregation problem. The model is also strikingly successful in accounting for realized rates. It accounts for more than 60 percent of daily changes in the DM/\$ rate, and more than 40 percent of daily changes in the Yen/\$ rate. Out of sample, our model produces better short-horizon forecasts than a random walk. Our estimates of the sensitivity of the spot rate to order flow are sensible as well, and square with past estimates at the individual-dealer level. We find that for the DM/\$ market as a whole, \$1 billion of net dollar purchases increases the DM price of a dollar by about 0.5 percent. This relation should be of particular interest to people working on central bank intervention (though care should be exercised in mapping central bank orders to subsequent interdealer trades).

---

<sup>29</sup> This primacy of order flow within microstructure should mitigate standard concerns about data snooping. (In the words of Richard Meese 1990, "At this point exchange rate modelers can be justly accused of in-sample data mining.") The variable we introduce—order flow—is *the* obvious a priori driving variable from microstructure theory.

Two issues raised by our measure of order flow deserve some remarks. First, though our measure captures a substantial share of total trading, it remains incomplete. As data sets covering customer-dealer trading and brokered inter-dealer trading become available, the order-flow picture can be completed (see, e.g., Payne 1999). A second interesting issue raised by our order-flow measure is whether its relation to price would change if order flow were observable to dealers in real time (i.e., if the market were more transparent). From a policy perspective, the effects of increasing order-flow transparency may be important: unlike most other financial markets, the FX market is unregulated in this respect. The welfare consequences are not yet well understood.

So where do the results of this paper lead us? In our judgment they point toward a research agenda that borrows from both the macro and microstructure approaches. It is not necessary to decouple exchange rates from macroeconomic fundamentals, as is common within microstructure finance. In this way, the approach is more firmly anchored in the broader context of asset pricing. (Though we freely admit that longer time series will be necessary to implement the macro dimension fully.) Nor is it necessary to treat exchange rates as driven wholly by public information, as is common within the macro approach. The information aggregation that arises when one reduces reliance on public information is well suited to microstructure: there are ample tools within the microstructure approach for addressing this aggregation. In the end, this two-pronged approach may help locate the missing middle in exchange rate economics—that disturbing space between our successful modeling of very short and very long horizons.

We close by addressing the obvious challenge for this agenda: What drives order flow? Here are two strategies, among many, for shedding light on this question. The first strategy involves disaggregating 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 non-financial 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? If not, whose trades are most informative? This will clarify the underlying sources of non-CK informa-

tion, which brings us closer to a specification of this market's information structure. The second strategy involves focusing on periods where we expect CK information to be most important, 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 channel for resolving uncertainty about the mapping from public information to price? Whatever the result, the important point is that the what-drives-order-flow question is not beyond our grasp.

## Appendix: Model Solution

Each dealer determines quotes and speculative demand by maximizing a negative exponential utility function defined over periodic wealth. Because returns are independent across periods, with an unchanging stochastic structure, the dealers' problem collapses to a series of independent trading problems, one for each period. Within a given period  $t$ , let  $W_{i\tau}$  denote the end-of-round  $\tau$  wealth of dealer  $i$ , where we use the convention that  $W_{i0}$  denotes wealth at the end of period  $t-1$ . (To ease the notational burden, we suppress the period subscript  $t$  when clarity permits.) With this notation, and normalizing the gross return on the riskless asset to one, we can write the dealers' problem as:

$$(A1) \quad \begin{aligned} & \text{Max} && E[-\exp(-qW_{i3} | \Omega_i)] \\ & && \{P_{i1}, P_{i2}, P_{i3}, T_{i2}\} \end{aligned}$$

s.t.

$$W_{i3} = W_{i0} + c_{i1}(P_{i1} - P'_{i2}) + (D_{i2} + E[T'_{i2} | \Omega_{i2}])(P_{i3} - P'_{i2}) - T'_{i2}(P_{i3} - P_{i2})$$

$P_{i\tau}$  is dealer  $i$ 's round- $\tau$  quote and a  $'$  denotes an interdealer quote or trade received by dealer  $i$ . The dealers' problem is defined over four choice variables: the three scalar quotes  $P_{i1}$ ,  $P_{i2}$ , and  $P_{i3}$ , and the dealer's outgoing interdealer trade in round 2,  $T_{i2}$ . This outgoing interdealer trade in round 2 has three components:

$$(A2) \quad T_{i2} = c_{i1} + D_{i2} + E[T'_{i2} | \Omega_{T_{i2}}]$$

where  $D_{i2}$  is dealer  $i$ 's speculative demand in round 2, and  $E[T'_{i2} | \Omega_{T_{i2}}]$  is the dealer's attempt to hedge against incoming orders from other dealers (this term is zero in equilibrium). The last three terms in  $W_{i3}$  capture capital gains/losses from round-1 customer orders  $c_{i1}$ , round-2 speculative demand  $D_{i2}$ , and the round-2 position disturbance from incoming interdealer orders  $T'_{i2}$ . The conditioning information  $\Omega_i$  at each decision node (3 quotes and 1 outgoing order) is summarized below.

$$\begin{aligned} \Omega_{P_{i1}} &\equiv \{ \{_{k=1}^t r_k \}, \{_{k=1}^{t-1} \Delta x_k \} \} \\ \Omega_{P_{i2}} &\equiv \{ \Omega_{P_{i1}}, c_{i1} \} \\ \Omega_{T_{i2}} &\equiv \{ \Omega_{P_{i2}} \} \\ \Omega_{P_{i3}} &\equiv \{ \Omega_{P_{i2}}, \Delta x_t \} \end{aligned}$$

### Conditional Variances

This appendix repeatedly uses several conditional return variances. These variances do not depend on conditioning variables' realizations (e.g., they do not depend on dealer  $i$ 's realization of  $c_{i1}$ ). These conditional variances are therefore common to all dealers and known in period one. (It is a convenient property of the normal distribution that realizations of conditioning variables affect the conditional mean but not the precision of the condition mean.) This predetermination of conditional variances is key to the derivation of optimal quoting and trading rules.

### Equilibrium

The equilibrium concept we use is Bayesian-Nash Equilibrium, or BNE. Under BNE, Bayes rule is used to update beliefs and strategies are sequentially rational given those beliefs.

Solving for the symmetric BNE, first we consider properties of optimal quoting strategies.

**PROPOSITION 1:** A quoting strategy is consistent with symmetric BNE only if the round-one and round-two quotes are common across dealers and equal to:

$$P_{1,t} = P_{2,t} = P_{3,t-1} + \beta_1 r_t$$

where  $P_{3,t-1}$  is the round-three quote from the previous period, and  $r_t$  is the payoff innovation at the beginning of period  $t$ .

**PROPOSITION 2:** A quoting strategy is consistent with symmetric BNE only if the common round-three quote is:

$$P_{3,t} = P_{2,t} + \beta_2 \Delta x_t$$

The constant  $\beta_2$  is strictly positive.

### Proof of Propositions 1 and 2

No arbitrage requires that all dealers post a common quote in all periods. (Recall from section 3 that all quotes are scalar prices at which the dealer agrees to buy/sell any amount, and trading with multiple partners is feasible.) Common prices require that quotes be conditioned on commonly observed information only. In rounds one and two, this includes the previous period's round-three price, plus the perpetuity value of the increment  $r$  to the periodic payoff  $R$ . (Dealer  $i$ 's round-two quote therefore cannot be conditioned on his realization of  $c_{i1}$ .)

The equations that pin down the levels of these three prices embed the dealer and customer trading rules. When conditioned on public information, these trading rules must be consistent with equilibrium price. This implies the following key relations:

$$(A3) \quad E[c_{i1} | \Omega_{P11}] + E[D_{i2}(P_{1,t}) | \Omega_{P11}] = 0$$

$$(A4) \quad E[c_{i1} | \Omega_{P_{i1}}] + E[D_{i2}(P_{2,t}) | \Omega_{P_{i1}}] = 0$$

$$(A5) \quad E[\Sigma_{iC_{i1}} | \Omega_{P_{i3}}] + E[\Delta c_3(P_{3,t}) | \Omega_{P_{i3}}] = 0$$

The first two equations simply state that, in expectation, dealers must be willing to absorb the demand from customers. The third equation states that, in expectation, the public must be willing at the round-3 price to absorb the day's aggregate portfolio shift (the change in  $c_3$  is the change in the public's total speculative demand). These equations pin down equilibrium price because any price except that which satisfies each would generate net excess demand in round-2 inter-dealer trading, which cannot be reconciled since dealers trade among themselves.

That  $P_{1,t}=P_{2,t}=P_{3,t-1}+\beta_1 r_t$  follows directly from the fact that expected value of  $c_{i1}$  conditional on public information  $\Omega_{P_{i1}}$  is zero, and expected speculative dealer demand  $D_{i2}$  is also zero at this public-information-unbiased price. To be more precise, this statement *postulates* that the dealer's demand  $D_{i2}$  has this property; we show below in the derivation of the optimal trading rule that this is the case.

That  $P_{3,t}=P_{2,t}+\beta_2 \Delta x_t$  follows from the fact that  $\Delta x_t$  is a sufficient statistic for the beginning-of-day aggregate portfolio shift  $\Sigma_{iC_{i1}}$  (which we denote  $c_1$ ). Given the aggregate portfolio shift must be absorbed by the public in round 3,  $P_{3,t}$  must adjust to induce the necessary public demand. Specifically, for market clearing the round-3 price must satisfy:

$$\Delta c_3(P_{3,t}) = -c_1$$

Given the optimal rule for determining  $T_{i2}$  (which we establish below), each day's  $c_1$  is proportional to that day's interdealer order flow  $\Delta x$ :

$$c_1 = (1/\alpha)\Delta x$$

and since the specification of  $c_3$  in the text is:

$$c_3 = g\left(E\left[P_{3,t+1} + R_{t+1} | \Omega_3\right] - P_{3,t}\right)$$

the change in the aggregate speculative holding,  $\Delta c_3$ , is  $c_3$  minus previous accumulated holdings, or:

$$\Delta c_3 = g\left(E\left[P_{3,t+1} + R_{t+1} | \Omega_3\right] - P_{3,t}\right) - \left(\sum_{t=1}^{t-1} -c_{1t}\right)$$

Because the last term in this expression is proportional to the sum of the previous day's  $\Delta x_t$ 's, this implies a market-clearing round-3 price of:

$$P_{3,t} = E\left[P_{3,t+1} + R_{t+1} \mid \Omega_3\right] + (\mathbf{ag})^{-1} \left( \sum_{t=1}^t \Delta x_t \right)$$

$$= \sum_{t=1}^t (\mathbf{b}_1 r_t + \mathbf{b}_2 \Delta x_t)$$

with  $\beta_2 = (\alpha\gamma)^{-1}$ , which is unambiguously positive. This sum is the perpetuity value of the expected periodic payoff  $R_t$  (i.e.,  $\beta_1$  times the sum of the  $r_t$  increments), adjusted for a risk premium determined by cumulative portfolio shifts (the  $\Delta x_t$  terms). This yields equation (5) in the text:

$$(5) \quad \Delta P_t = \beta_1 r_t + \beta_2 \Delta x_t$$

where  $\Delta P_t$  denotes the change in price from the end of round 3 in period  $t-1$  to the end of round 3 in period  $t$ . Note that conditional on information available at the end of period  $t$ , the expected change in price over the next period is zero; given the expected stream of payoffs  $R_t$ , the level of price is sufficient to maintain public absorption of all past portfolio shifts.

### Equilibrium Trading Strategies

An implication of common interdealer quotes  $P_{2,t}$  is that in round 2 each dealer receives a share  $1/(N-1)$  of every other dealer's interdealer trade. This order corresponds to the position disturbance  $T_{i2}$  in the dealer's problem in equation (A1).

Given the quoting strategy described in propositions 1 and 2, the following trading strategy is optimal and corresponds to symmetric linear equilibrium:

PROPOSITION 3: The trading strategy profile:

$$T_{i2} = \alpha c_{i1}$$

$\forall i \in \{1, \dots, N\}$ , with  $\alpha > 0$ , conforms to a Bayesian-Nash equilibrium.

### Proof of Proposition 3: Optimal Trading Strategies

As noted above, because returns are independent across periods, with an unchanging stochastic structure, the dealers' problem collapses to a series of independent trading problems, one for each period. Because there are only  $N$  dealers, however, each dealer acts strategically in the sense that his speculative demand depends on the impact his trade will have on subsequent prices.

It is well known that if a random variable  $W$  is distributed  $N(\mu, \sigma^2)$  and the utility function  $U(W) = -\exp(-\theta W)$ , then:

$$(A6) \quad E[U(W)] = -\exp[-\theta(\mu - \theta\sigma^2/2)]$$

Maximizing  $E[U(W)]$  is therefore equivalent to maximizing  $(\mu - \theta\sigma^2/2)$ . This result allows us to write the dealers speculative-demand problem as:

$$(A7) \quad \text{Max}_{D_{i2}} \left( E[P_3 | \Omega_{T_{i2}}] - P_2 \right) - D_{i2}^2 (\theta/2)\sigma^2$$

where the information set  $\Omega_{T_{i2}}$  is defined above, and  $\sigma^2$  denotes the conditional variance of  $E[P_3 | \Omega_{T_{i2}}] - P_2$ . Now, from Proposition 2, we can write:

$$(A8) \quad E[P_3 | \Omega_{T_{i2}}] - P_2 = E[\beta_2 \Delta x | \Omega_{T_{i2}}]$$

And from the definitions of  $\Omega_{T_{i2}}$  and  $\Delta x$  we know that:

$$(A9) \quad E[\beta_2 \Delta x | \Omega_{T_{i2}}] = \beta_2 T_{i2}$$

The expected value of the other dealers' trades in  $\Delta x$  is 0 under our specification because (i) customer trades are mean-zero and independent across dealers and (ii) there is no information in the model other than customer trades to motivate speculative demand. This fact also implies that dealer  $i$ 's trade in round 2,  $T_{i2}$  from equation (A2), is equal to:

$$T_{i2} = D_{i2} + c_{i1}$$

Therefore, we can write the dealer's problem as:

$$(A10) \quad \text{Max}_{D_{i2}} D_{i2} \beta_2 (D_{i2} + c_{i1}) - D_{i2}^2 (\theta/2)\sigma^2$$

The first-order condition of this problem is:

$$(A11) \quad 2\beta_2 D_{i2} + c_{i1} - \theta\sigma^2 D_{i2} = 0$$

which implies a speculative demand of:

$$(A12) \quad D_{i2} = \left( \frac{1}{\theta\sigma^2 - 2\beta_2} \right) c_{i1}$$

This demand function and the fact that  $T_{i2} = D_{i2} + c_{i1}$  imply:

$$(A13) \quad T_{i2} = \left( \frac{1}{q\mathbf{s}^2 - 2\mathbf{b}_2} + 1 \right) c_{i1} \equiv \mathbf{a} c_{i1}$$

The second-order condition for a maximum,  $(2\beta_2 - \theta\sigma^2) < 0$ , insures that  $\alpha > 0$ .

## References

- Amihud, Y., and H. Mendelson, 1980, Dealership market: Marketmaking with inventory, *Journal of Financial Economics*, 8: 31-53.
- Andersen, T., and T. Bollerslev, 1998, Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, 53: 219-266.
- Bagwell, L., 1992, Dutch auction repurchases: An analysis of shareholder heterogeneity, *Journal of Finance*, 47: 71-105.
- Bjornes, G., and D. Rime, 1998, FX trading ... live: Impact of new trading environments, typescript, Norwegian School of Management, December.
- Blanchard, O., 1979, Speculative bubbles, crashes, and rational expectations, *Economics Letters*, 14: 387-389.
- Burnham, J., 1991, Current structure and recent developments in foreign exchange markets, in S. Khoury (ed.), *Recent Developments in International Banking and Finance* (North Holland, Amsterdam), 123-153.
- Cai, J., Y. Cheung, R. Lee, and M. Melvin, 1999, "Once-in-a-generation" yen volatility in 1998: Fundamentals, intervention, or order flow? typescript, Arizona State University, July.
- Cheung, Y., and M. Chinn, 1998, Macroeconomic implications of the beliefs and behavior of foreign exchange traders, typescript, U.C. Santa Cruz, November.
- Cheung, Y., and C. Wong, 1998, A survey of market practitioners' views on exchange rate dynamics, *Journal of International Economics*, forthcoming.
- Chinn, M., 1991, Some linear and non-linear thoughts on exchange rates, *Journal of International Money and Finance*, 10: 214-230.
- Chinn, M., and R. Meese, 1994, Banking on currency forecasts, *Journal of International Economics*, 38: 161-178.
- Covrig, V., and M. Melvin, 1998, Asymmetric information and price discovery in the FX market: Does Tokyo know more about the yen?, typescript, Arizona State University.
- Degennaro, R., and R. Shrieves. 1997, Public information releases, private information arrival, and volatility in the foreign exchange market, *Journal of Empirical Finance*, 4: 295-315.
- DeLong, B., A. Shleifer, L. Summers, and R. Waldmann, 1990, Positive feedback investment strategies, and destabilizing rational speculation, *Journal of Finance*, 45, 379-396.
- Dominguez, K., 1986, Are foreign exchange forecasts rational? New evidence from survey data, *Economic Letters*, 21: 277-281.
- Dornbusch, R., 1976, Expectations and exchange rate dynamics, *Journal of Political Economy*, 84: 1161-1176.
- Dornbusch, R., 1982, Equilibrium and disequilibrium exchange rates, *Zeitschrift fur Wirtschafts und Sozialwissenschaften*, 102: 573-799; reprinted in R. Dornbusch, *Dollars, Debts and Deficits*, MIT Press: Cambridge, MA.
- Easley, D., and M. O'Hara, 1992, Time and the process of security price adjustment, *Journal of Finance*, 47: 577-605.

- Evans, G., 1986, A test for speculative bubbles in the sterling-dollar exchange rate, *American Economic Review*, 76: 621-636.
- Evans, M., 1997, The microstructure of foreign exchange dynamics, Georgetown University typescript, November, presented at the NBER Market Microstructure Conference, December 4, 1997.
- Evans, M., 1999, What are the origins of foreign exchange movements?, Georgetown University typescript, March.
- Fair, R., 1970, Estimation of simultaneous equations models with lagged endogenous variables and first order serially correlated errors," *Econometrica*, 38: 507-516.
- Flood, M., 1994, Market structure and inefficiency in the foreign exchange market," *Journal of International Money and Finance*, 13: 131-158.
- Flood, R., and R. Hodrick, On testing for speculative bubbles, *Journal of Economic Perspectives*, 4: 85-101.
- Flood, R., and A. Rose, 1995, Fixing exchange rates: A virtual quest for fundamentals, *Journal of Monetary Economics* 36, 3-37.
- Flood, R., and M. Taylor, 1996, Exchange rate economics: What's wrong with the conventional macro approach?, in Frankel, J., G. Galli, and A. Giovannini, (eds.) *The Microstructure of Foreign Exchange Markets* (Chicago: The University of Chicago Press), 261-294.
- Foster, D., and S. Viswanathan, 1990, A theory of interday variations in volumes, variances, and trading costs in securities markets, *Review of Financial Studies*, 3: 593-624.
- Frankel, J., and K. Froot, 1987, Using survey data to test standard propositions regarding exchange rate expectations, *American Economic Review*, 77: 133-153.
- Frankel, J., G. Galli, and A. Giovannini, 1996, Introduction, in *The Microstructure of Foreign Exchange Markets* (Chicago: The University of Chicago Press), 1-15.
- Frankel, J., and A. Rose, 1995, Empirical research on nominal exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science: Amsterdam, 1689-1729.
- Frenkel, J., 1976, A monetary approach to the exchange rate: Doctrinal aspects and empirical evidence, *Scandinavian Journal of Economics*, 78: 200-224.
- Glosten, L., and P. Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *Journal of Financial Economics*, 14: 71-100.
- Goldberg, L., and R. Tenorio, 1997, Strategic trading in a two-sided foreign exchange auction, *Journal of International Economics*, 42: 299-326.
- Goodhart, C., 1988, The foreign exchange market: A random walk with a dragging anchor, *Economica*, 55: 437-460.
- Goodhart, C., T. Ito, and R. Payne, 1996, One day in June 1993: A study of the working of the Reuters 2000-2 electronic foreign exchange trading system. In *The Microstructure of Foreign Exchange Markets*, eds. J. Frankel, G. Galli and A. Giovannini, pp. 107-179. University of Chicago Press, Chicago, IL.
- Goodhart, C., and M. O'Hara, 1997, High frequency data in financial markets: Issues and applications, *Journal of Empirical Finance*, 4: 73-114.

- Grinblatt, M., S. Titman, and R. Wermers, 1995, Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior, *American Economic Review*, 85, 1088-1105.
- Hansch, O., N. Naik, and S. Viswanathan 1998. "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange," *Journal of Finance*, 53: 1623-1656.
- Hau, H., 1998, Competitive entry and endogenous risk in the foreign exchange market, *Review of Financial Studies*, 11: 757-788.
- Ho, T., and H. Stoll, 1981, Optimal dealer pricing under transactions and return uncertainty, *Journal of Financial Economics*, 9, 47-73.
- Hsieh, D., and A. Kleidon, 1996, Bid-ask spreads in foreign exchange markets: Implications for models of asymmetric information, in Jeffrey Frankel et al. eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago) 41-65.
- Isard, P., 1995, *Exchange Rate Economics*, Cambridge University Press: Cambridge, UK.
- Ito, T., R. Lyons, and M. Melvin, 1998, Is there private information in the FX market? The Tokyo experiment, *Journal of Finance* 53, 1111-1130.
- Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance*, 48, 65-92.
- Jones, C., G. Kaul, and M. Lipson, 1994, Transactions, volume, and volatility, *Review of Financial Studies*, 7: 631-651.
- Kaul, A., M. Vikas, and M. Randall, 2000, Demand curves for stock do slope down: New evidence from an index weights adjustment, *Journal of Finance*, 55: 893-912.
- Kouri, P., 1976, The exchange rate and the balance of payments in the short run and in the long run: A monetary approach, *Scandinavian Journal of Economics*, 78, 280-304.
- Kouri, P., and M. Porter, 1974, International capital flows and portfolio equilibrium, *Journal of Political Economy*, 82, 443-467.
- Kyle, A., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.
- Lyons, R., 1995, Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321-351.
- Lyons, R., 1996, Optimal transparency in a dealer market with an application to foreign exchange, *Journal of Financial Intermediation*, 5, 225-254.
- Lyons, R., 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275-298.
- Lyons, R., 1998, Profits and position control: A week of FX dealing, *Journal of International Money and Finance* 17, 97-115.
- Lyons, R., 1999, *The Microstructure Approach to Exchange Rates*, MIT Press (chapters at [www.haas.berkeley.edu/~lyons/NewBook.html](http://www.haas.berkeley.edu/~lyons/NewBook.html)), forthcoming.
- Mark, N., 1995, Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review*, 85: 201-218.
- Meese, R., 1986, Testing for bubbles in exchange markets, *Journal of Political Economy*, 94: 345-373.
- Meese, R., 1990, Currency fluctuations in the post-Bretton Woods era, *Journal of Economic Perspectives*, 4: 117-134.

- Meese, R., and K. Rogoff, 1983a, Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.
- Meese, R., and K. Rogoff, 1983b, The out-of-sample failure of empirical exchange rate models, in J. Frenkel (ed.), *Exchange Rate and International Macroeconomics*, University of Chicago Press: Chicago.
- Meese, R., and K. Rogoff, 1988, Was it real? The exchange rate-interest differential relation of the modern floating-rate period," *Journal of Finance*, 43: 933-948.
- Mussa, M., 1976, The exchange rate, the balance of payments, and monetary and fiscal policy under a regime of controlled floating, *Scandinavian Journal of Economics*, 78: 229-248.
- Naranjo, A., and M. Nimalendran, 1999, Government intervention and adverse selection costs in foreign exchange markets, *Review of Financial Studies*, forthcoming.
- Obstfeld, M., and K. Rogoff, 1995, Exchange rate dynamics redux. *Journal of Political Economy*, 103: 624-660.
- Osler, C., 1998, Short-term speculators and the puzzling behavior of exchange rates, *Journal of International Economics*, 45: 37-57.
- 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, 1589-1614.
- Reiss, P., and I. Werner, 1998, Does risk sharing motivate interdealer trading? *Journal of Finance*, 53: 1657-1704.
- 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](http://www.uio.no/~dagfinri)).
- Roll, R., 1988,  $R^2$ , *Journal of Finance*, 43: 541-566.
- Romer, D., 1993. Rational asset-price movements without news, *American Economic Review*, 83: 1112-1130.
- Scholes, M., 1972, The market for securities: Substitution versus price pressure and the effect of information on share price, *Journal of Business*, 45: 179-211.
- Shleifer, A., 1986, Do demand curves for stocks slope down? *Journal of Finance*, 41: 579-590.
- Stoll, H., 1978, The supply of dealer services in securities markets, *Journal of Finance*, 33, 1133-1151.
- Taylor, M., 1995, The economics of exchange rates, *Journal of Economic Literature* 33, 13-47.
- Viswanathan, S., and J. Wang, 1998, Why is interdealer trade so pervasive in financial markets? Working paper, Duke University.
- Vitale, P., 1998, Two months in the life of several gilt-edged market makers on the London Stock Exchange, *Journal of International Financial Markets, Institutions, & Money*, 8: 301-326.
- Vogler, K., 1997, Risk allocation and interdealer trading, *European Economic Review* 41, 417-441.

Wei, S., and J. Kim, 1997, The big players in the foreign exchange market: Do they trade on information or noise?, NBER Working Paper 6256, November.  
Yao, J. 1998. Market making in the interbank foreign exchange market, New York University Salomon Center Working Paper #S-98-3.