

NBER WORKING PAPER SERIES

ORDER FLOW AND
EXCHANGE RATE DYNAMICS

Martin D. D. Evans
Richard K. Lyons

Working Paper 7317
<http://www.nber.org/papers/w7317>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1999

We thank the following for valuable comments: Menzie Chinn, Peter DeMarzo, Petra Geraats, Richard Meese, Michael Melvin, Andrew Rose, Mark Taranto, Ingrid Werner, and seminar participants at UC Berkeley and the 1999 NBER Summer Institute (IFM). Lyons thanks the National Science Foundation for financial assistance. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 1999 by Martin D. D. Evans and Richard K. Lyons. All rights reserved. Short sections of text, not to

exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Order Flow and Exchange Rate Dynamics

Martin D. D. Evans and Richard K. Lyons

NBER Working Paper No. 7317

August 1999

JEL No. F31, G15

ABSTRACT

Macroeconomic models of nominal exchange rates perform poorly. In sample, R^2 statistics as high as 10 percent are rare. Out of sample, these models are typically out-forecast by a naïve random walk. 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 exchange-rate changes produces R^2 statistics above 50 percent. Out of sample, our model produces significantly better short-horizon forecasts than a random walk. For the DM/\$ spot market as a whole, we find that \$1 billion of net dollar purchases increases the DM price of a dollar by about 1 pfennig.

Martin D. D. Evans
Department of Economics
Georgetown University
Washington, DC 20057
and NBER
evansm1@gunet.georgetown.edu

Richard K. Lyons
Haas School of Business
UC Berkeley
Berkeley, CA 94720-1900
and NBER
lyons@haas.berkeley.edu

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.”¹

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

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

Froot 1987, and Hau 1998, among others). On *a priori* grounds, 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 rate variability. This, too, remains 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.² 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.

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 for the respective currencies. Order flow, denoted by x , is the sum over time of signed trades between foreign exchange dealers worldwide.³

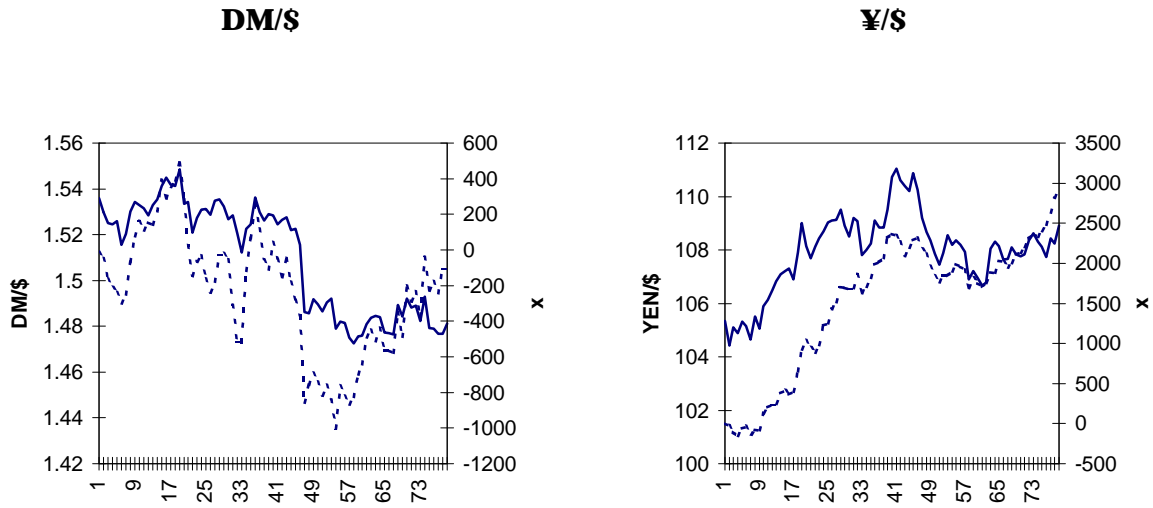
² Order flow is a measure of buying/selling pressure. It is the net of buyer-initiated orders and seller-initiated orders. In a dealer market, it is the dealers who absorb this order flow, and they are compensated for doing so. (In an auction market, limit orders absorb the flow of market orders.)

³ 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).

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. 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 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 heart of the market.

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 the shorter-horizon results 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 over 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, Bjornes and Rime 1998, Evans 1999, and Payne 1999.)⁴

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

⁴ The standard example of order flow that conveys non-public information is orders from central bank intervention. Probably more 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.)

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.⁵ 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.

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 Δp_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,

⁵ 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.

denoted here by the ellipsis.⁶ 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 ε_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.⁷ These models are all variations on the following specification:

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

Now Δp_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 order flow Δx , the change in net dealer positions (or inventory) ΔI , 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 Δx 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 Δp and Δx because order flow communicates non-public information, and once communicated, it is

⁶ 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.

⁷ 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.

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 Δp and Δx 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 Δp and ΔI 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:⁸

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

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 Δx . Our time-aggregated measure spans a much longer period than is addressed elsewhere within empirical microstructure.

⁸ 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.⁹ 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. 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 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).


⁹ Evidence that asset demand curves slope down is provided by Scholes (1972) and Shleifer (1986), among many others.

Specifics

Consider a pure exchange economy with T periods and two assets, one riskless, and one with a stochastic payoff representing foreign exchange. The time-T payoff on foreign exchange, denoted F, is composed of a series of increments, so that $F = \sum_{t=1}^T r_t$. The increments r_t are i.i.d. $\text{Normal}(0, \Sigma_r)$ and are observed before trading in each period. These realized increments represent the flow of publicly available macroeconomic information over time (e.g., changes in interest rates).

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 period is:

Daily Timing



The dealers and customers all have identical negative exponential utility defined over time-T wealth.

Trading Round 1

At the beginning of each period t, all market participants observe r_t , the period's increment to the payoff F. On the basis of this increment and other avail-

able information, each dealer simultaneously and independently quotes a scalar price to his customers at which he agrees to buy and sell any amount.¹⁰ 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 .¹¹ 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 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

¹⁰ 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.

¹¹ A natural extension of this specification is that customer orders reflect changing expectations of future r_t .

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_{i3} 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.

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 its expected return conditional on public information:

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

where the positive coefficient γ captures the aggregate risk-bearing capacity of the public, and Ω_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_{i1} , P_{i2} , and P_{i3} , and the dealer's interdealer trade T_{i2} (the latter being a component of Δx , the interdealer order flow). 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 r_t is common information at the beginning of round 1, order flow Δx_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 Δx_t .

Whether Δx 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:

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

with a constant coefficient α . 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:

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

where λ is a positive constant. That this price change includes the innovation in payoffs r_t one-for-one is unsurprising. The $\lambda \Delta x_t$ 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 $\lambda\Delta x = \lambda\Sigma_i T_{i2} = \lambda\alpha\Sigma_i c_{i1}$. The sum $\Sigma_i c_{i1}$ is this total portfolio shift from round 1. The public's total demand in round 3, c_3 , is not perfectly elastic, and λ insures that at the round-3 price $c_3 + \Sigma_i c_{i1} = 0$.

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 ΔP using the change in the interest differential as an independent variable, rather than the differential's level. 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.¹² 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.

¹² 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.”

4. Data

Our data set contains time-stamped, tick-by-tick data on actual transactions for the two largest spot markets—DM/\$ and ¥/\$—over a four-month period, May 1 to August 31, 1996. These data were collected from the Reuters Dealing 2000-1 system via an electronic feed customized for the purpose. Dealing 2000-1 is the most widely used electronic dealing system. According to Reuters, over 90 percent of the world's direct interdealer transactions take place through the system.¹³ All trades on this system take the form of bilateral electronic conversations. The conversation is initiated when a dealer uses the system to call another dealer to request a quote. Users are expected to provide a fast two-way quote with a tight spread, which is in turn dealt or declined quickly (i.e., within seconds). To settle disputes, Reuters keeps a temporary record of all bilateral conversations. This record is the source of our data. (Reuters was unable to provide the identity of the trading partners for confidentiality reasons.)

For every trade executed on D2000-1, our data set includes a time-stamped record of the transaction price and a bought/sold indicator. The bought/sold indicator allows us to sign trades for measuring order flow. This is a major advantage: we do not have to use the noisy algorithms used elsewhere in the literature for signing trades. A drawback is that it is not possible to identify the size of individual transactions.¹⁴ For model estimation, order flow Δx is therefore measured as the difference between the number of buyer-initiated trades and the number of seller-initiated trades.

Three 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

¹³ As noted in footnote 3, interdealer transactions account for about 75 percent of total trading in major spot markets. This 75 percent from interdealer trading breaks into two transaction types, direct and brokered. Direct trading accounts for about 60 percent of interdealer trade and brokered trading accounts for about 40 percent. For more detail on the Reuters Dealing 2000-1 System see Lyons (1995) and Evans (1997).

¹⁴ This drawback may not be acute. There is evidence that the size of trades has no information content beyond that contained in the number of transactions. See Jones, Kaul, and Lipson (1994).

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 market-wide transactions data are not observed by individual FX dealers as they trade. Though dealers have access to their own transaction records, they cannot observe others' transactions on the system. Our data therefore represent activity that, at the time, participants could only infer indirectly. This is one of those rare situations where the researcher has more information than market participants themselves (at least in this dimension).

Third, our data cover a relatively long time span (four months) in comparison with other micro data sets. 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 ¥/\$), Δp_t , is the log change in the purchase 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, Δx_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).

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 Δp_t is the change in the log spot rate (DM/\$ or ¥/\$) from the end of day $t-1$ to the end of day t , $\Delta(i_t - i_t^*)$ is the change in the overnight interest differential from day $t-1$ to day t (* denotes DM or ¥), and Δx_t is the order flow from the end of day $t-1$ to the end of day t (negative denotes net dollar sales).¹⁵

The coefficient β_2 on our portfolio shift variable Δx_t is correctly signed and significant, with t -statistics above 5 in both equations. To see that the sign is correct, recall from the model that net purchases of dollars—a positive Δx_t —should lead to a higher DM price of dollars. The traditional macro-fundamental—the interest differential—is correctly signed, but is only significant in the yen equation. (The sign should be positive because, in the sticky-price monetary model for example, an increase in the dollar interest rate i_t requires an immediate dollar appreciation—increase in DM/\$—to make room for the expected dollar depreciation required by uncovered interest parity.) The overall fit of the model is striking relative to

¹⁵ 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.

traditional macro models, with R^2 statistics of 64 percent and 45 percent for the DM and yen equations, respectively. In terms of diagnostics, the DM equation shows some evidence of heteroskedasticity, so we correct the standard errors in that equation using a heteroskedasticity-consistent covariance matrix (White correction).¹⁶

Table 1

In-sample fit of portfolio shifts model

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

	$\Delta(i_t - i_t^*)$	Δx_t	Diagnostics		
			R^2	Serial	Hetero
DM	0.52 (0.35)	2.10 (0.20)	0.64	0.78 0.41	0.08 0.02
Yen	2.48 (0.92)	2.90 (0.46)	0.45	0.50 0.37	0.96 0.71

The dependent variable Δp_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 Δx_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.

¹⁶ To check robustness, we examine several obvious variations on the model. For example, we include a constant in the regression, even though the model does not call for one. The constant is insignificant for both currencies. We also include the level of the interest differential in lieu of its change. This is also insignificant in both cases. Estimating the whole model in levels rather than changes produces a pattern similar to that in Table 1: order flow is highly significant, the interest differential is insignificant, and R^2 is 0.75 for the DM equation and 0.61 for the Yen equation. With this levels regressions, however, beyond the usual concerns about non-stationarity, there is also strong evidence of serial correlation and heteroskedasticity (both tests are significant at the 1 percent level for both currencies). Finally, recall that our price series is measured from purchase transactions. Results using 4 pm sale prices are identical. We address additional robustness issues in the next subsection.

The size of our order flow coefficient is consistent with past estimates based on single-dealer data. The coefficient of 2.1 in the DM equation implies that a day with 1000 more dollar purchases than sales induces an increase in the DM price by 2.1 percent. Given the average trade size in our sample of \$3.9 million, \$1 billion of net dollar purchases increases the DM price of a dollar by 0.54% ($= 2.1/3.9$). At a spot rate of 1.5 DM/\$, this implies that *\$1 billion of net dollar purchases increases the DM price of a dollar by 0.8 pfennig*. At the single-dealer level, Lyons (1995) finds that information asymmetry induces the dealer he tracks to increase price by $1/100^{\text{th}}$ of a pfennig (0.0001 DM) for every incoming buy order of \$10 million. That translates to 1 pfennig per \$1 billion. Though linearly extrapolating this estimate is certainly not an accurate description of single-dealer behavior, with multiple dealers it may be a good description of the market's aggregate elasticity.

The striking explanatory power of these regressions is almost wholly due to order flow Δx_t . Regressing Δp_t on $\Delta(i_t - i_t^*)$ alone, plus a constant, produces an R^2 statistic less than 1 percent in both equations, and coefficients on $\Delta(i_t - i_t^*)$ that are insignificant at the 5 percent level.¹⁷ That the interest differential regains significance once order flow is included, at least in the Yen equation, is consistent with omitted variable bias in the interest-rates-only specification. (The correlation between the two regressors Δx_t and $\Delta(i_t - i_t^*)$ is 0.02 for the DM and -0.27 for the Yen, though both are insignificant at the 5 percent level.)

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

¹⁷ 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."

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, Δx_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 .¹⁸ If a customer happens to buy dealer k 's position from him, then x_t remains at -2 . In this simple scenario, order flow measured only from trades between customers and dealers would have reverted to zero—the concluding customer trade offsets the initiating customer trade. 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. 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 Δx_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?

¹⁸ 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).

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.¹⁹ 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 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 β_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 β_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

¹⁹ 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.

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.²⁰ 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. (Put differently, we are not simply “regressing price on quantity,” as in the classic supply-demand identification problem.) Within microstructure theory more broadly, this direction of causality is well established: it is a feature of all the canonical models (Glosten and Milgrom 1985, Kyle 1985, Stoll 1978, Amihud and Mendelson 1980). Nevertheless, there are certainly alternative hypotheses under which causality is reversed.²¹ Let us consider the issue from a broader perspective.

²⁰ 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.

²¹ There is also the possibility that a common process simultaneously drives both order flow and price. But because our relation involves order flow rather than simply volume, this alternative is difficult to motivate on theoretical grounds.

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.

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 because price adjusts only after the news anticipated by order flow is commonly observed (Foster and Viswanathan 1990). Order flow might also precede price because price adjusts only after order flow itself is commonly observed (Lyons 1996). 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).²² Under the Feedback hypothesis, order flow lags price because of feedback trading. (Positive-feedback trading involves systematic buying in response to price increases, and selling in response to price decreases. Negative-feedback trading is the reverse.) Variations are distinguished by whether this feedback trading is rational (an optimal response 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.²³ Causality is less clear under the Anticipation hypothesis. The variation of the Anticipation hypothesis where price adjusts only after the news anticipated by order flow is observed is probably not relevant to foreign exchange (in contrast to equity markets, where insider order flow can anticipate a firm's earnings announcement). The other variation of the Anticipation

²² 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).

²³ This does not imply that price cannot influence order flow. Price does influence order flow in microstructure 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.

hypothesis, however, probably is relevant to the foreign exchange market. As noted in the Data section, order flow in this market is not common knowledge when realized. Consequently, there may be a lag in price adjustment (a lag which would not violate market efficiency). In this case, causality still runs from order flow to price, but the effects are delayed. In fact, we can test this variation of the Anticipation hypothesis 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.²⁴

Under the Feedback hypothesis, causality can clearly go in reverse, that is, from price to order flow.²⁵ This is the source of concern about simultaneity bias. If 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 Δx_t being explained, at least in part, by Δp_{t-1} . If our order-flow coefficient in Table 1 is picking up this daily-frequency positive feedback, then including lagged price change Δp_{t-1} in the Portfolio-Shifts regression should weaken, if not eliminate, the significance of order flow. 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.

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 lines of defense against this alternative interpretation of our results. The first is a set of three arguments why intraday

²⁴ 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 Δx 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.

²⁵ Note that the Feedback hypothesis does not imply that all causality runs in reverse. For example, the Feedback hypothesis does not rule out that feedback trading can affect prices.

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.)

Table 2

Portfolio shifts model: Alternative specifications

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

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

	$\Delta(i_t - i_t^*)$	Δx_t	Δx_{t-1}	Δp_{t-1}	Diagnostics		
					R²	Serial	Hetero
DM	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
Yen	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 Δp_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 Δx_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.

There are three reasons, a priori, why the hypothesis of intraday positive-feedback trading is unappealing in the context of our analysis. First, direct empirical evidence does not support it: the current literature provides no evidence of positive-feedback trading in the foreign exchange market. Second, even if systematic positive-feedback trading were present, it would be difficult to rationalize: intraday transaction studies find no evidence of the positive autocorrelation in price that would make positive-feedback an optimal response (Goodhart, Ito, and Payne 1996). Third, we view the fallback possibility of irrational positive-feedback trading as unlikely in the context of our analysis. Recall that the order flow we measure in this paper 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 Δx_t into two components:

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

where Δx_{t1} denotes exogenous order flow from portfolio shifts (a la our model), with variance equal to Σ_{x1} , and Δx_{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 ε_t represents common-knowledge (CK) news, and ε_t is iid with variance Σ_ε . 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 α will be a reduced form coefficient that depends on γ . Note that under these circumstances, equation (10) is a valid reduced-form equation that could be estimated by OLS if one had data on Δx_{t1} .

With data on Δx_t and Δp_t only, suppose we estimate

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

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

$$\begin{aligned} \beta &= \text{Cov}(\Delta p_t, \Delta x_t) / \text{Var}(\Delta x_t) \\ \delta &= \text{Var}(\Delta p_t) / \text{Var}(\Delta x_t) \end{aligned}$$

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:

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

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 ϕ 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 α and γ for given values of ϕ . The following table presents these implied values of α and γ .

Table 3

Bias Analysis

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

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

	$\phi = \Sigma_\varepsilon / \Sigma_{x1}$	α	γ
DM	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
Yen	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 α (order-flow-causes-price) and γ (price-causes-order-flow) implied by the sample moments and given values for the parameter ϕ . The parameter ϕ is the ratio of common-knowledge news to order-flow news.

Note that even for values of ϕ above 2, the feedback trading needed to generate our results is still negative. Note too that the parameter α —the order-flow-causes-price parameter—is not driven to zero until ϕ 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 far too extreme to be compelling.²⁶

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. 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

²⁶ In the end, even those readers who cling to the reverse-causality interpretation of our order-flow/price relation should find our results of interest—evidence of such strong positive feedback in this market would itself be a rather remarkable new finding.

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>
DM	<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)
Yen	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? The portfolio shifts in our model are exogenous and unrelated to macro fundamentals like interest rates. This is a straightforward way to introduce these shifts, but it is certainly not the only way. There are alternatives under which these shifts are driven by underlying determinants. The specific underlying determinants may be within the traditional set of macroeconomic variables (e.g., the “proxy-for-expectations” idea introduced in the introduction). Or they may be outside the traditional set, necessitating a broader definition of what we mean by “fundamentals.” Exploring these links to deeper determinants is a natural topic for future research. This will surely require a retreat back into intraday data.²⁷

²⁷ The role of macro announcements in determining order flow clearly 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).

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 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, these 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, Δx , 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.²⁸ This is a radical departure from traditional approaches to exchange rate determination.

Our Portfolio Shifts model is strikingly successful in accounting for realized rates. It accounts for more than 50 percent of daily changes in the DM/\$ rate, and more than 30 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 1 pfennig. This relation should be of particular interest to those working in the area of central bank intervention (though care should be exercised in mapping central bank orders to subsequent interdealer trades).

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. 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

²⁸ 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.

rate economics—that disturbing space between our successful modeling of very short and very long horizons.

Two issues raised by our measure of order flow deserve some final remarks. First, though our measure captures a substantial share of wholesale trading, it remains incomplete. As data sets covering customer-dealer trading and brokered interdealer 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.

Appendix: Model Solution

Each dealer determines quotes and speculative demand by maximizing a negative exponential utility function defined over terminal 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 τ 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:

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

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- τ 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 Ω_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:

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

where $P_{3,t-1}$ is the round-three quote from the previous period, and r_t is the public-information 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:

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

The constant λ 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 public-information innovation at the beginning of period t , r_t . (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:

$$(A3) \quad E[c_{i1} | \Omega_{Pi1}] + E[D_{i2}(P_{1,t}) | \Omega_{Pi1}] = 0$$

$$(A4) \quad E[c_{i1} | \Omega_{Pi1}] + E[D_{i2}(P_{2,t}) | \Omega_{Pi1}] = 0$$

$$(A5) \quad E[\sum_i c_{i1} | \Omega_{Pi3}] + E[C_3(P_{3,t}) | \Omega_{Pi3}] = 0$$

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 period's aggregate

portfolio shift. These equations pin down equilibrium price because any price except that which satisfies each would generate net excess demand in round-2 interdealer trading, which cannot be reconciled since dealers trade among themselves.

That $P_{1,t}=P_{2,t}=P_{3,t-1}+r_t$ 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 $P_{3,t}=P_{2,t}+\lambda\Delta x_t$ follows from the fact that Δx_t is a sufficient statistic for the period's aggregate portfolio shift $\Sigma_i c_{i1}$. 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, the round-3 price must satisfy:

$$c_3(P_{3,t}) = -\Sigma_i c_{i1}$$

Given the optimal rule for determining T_{i2} (which we establish below), we can write $\Sigma_i c_{i1}$ in terms of interdealer order flow Δx_t as:

$$\Sigma_i c_{i1} = (1/\alpha)\Delta x_t$$

and since the specification of c_3 in the text is:

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

this implies a market-clearing round-3 price of:

$$\begin{aligned} P_{3,t} &= E[P_{3,t+1} | \Omega_3] + (\alpha\gamma)^{-1}\Delta x_t \\ &= \sum_{i=1}^t (r_i + \lambda\Delta x_i) \end{aligned}$$

with $\lambda=(\alpha\gamma)^{-1}$, which is unambiguously positive. This sum is the expected payoff on the risky asset (the r_i terms), adjusted for a risk premium, which is determined by cumulative portfolio shifts (the Δx_i terms). This yields equation (5) in the text:

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

where ΔP_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.

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:

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

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

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

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

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

The expected value of the other dealers' trades in Δx 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:

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

The first-order condition of this problem is:

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

which implies a speculative demand of:

$$(A12) \quad D_{i2} = \left(\frac{1}{q\mathbf{s}^2 - 2\mathbf{l}} \right) c_{i1}$$

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

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

The second-order condition for a maximum, $(2\lambda - \theta\sigma^2) < 0$, 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.
- Bjornnes, 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, 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.
- 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), 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.
- 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.
- 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.
- 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.

</ref_section>