

NBER WORKING PAPER SERIES

FIXED VERSUS FLEXIBLE: LESSONS FROM EMS ORDER FLOW

William P. Killeen
Richard K. Lyons
Michael J. Moore

Working Paper 8491
<http://www.nber.org/papers/w8491>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2001

Killeen now works for Setanta Asset Management in Dublin. The original version of this paper was written while Killeen worked at BNP Paribas Asset Management in London. We thank the following for valuable comments: Kathryn Dominguez, Ken Froot, Carol Osler, Helene Rey, Andrew Rose, participants at the Fall 2000 NBER IFM meeting, and the lunch- groups at Berkeley and Queens. Killeen also wishes to thank Mr. Loic Meinnel, Global Head of Foreign Exchange, BNP Paribas for very useful discussions on FX markets during this project. Thanks too to Tina Kane. 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, BNP Paribas or Setanta Asset Management.

© 2001 by William P. Killeen, Richard K. Lyons and Michael J. Moore. 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.

Fixed versus Flexible: Lessons from EMS Order Flow
William P. Killeen, Richard K. Lyons and Michael J. Moore
NBER Working Paper No. 8491
September 2001
JEL No. F31, G12, G15

ABSTRACT

This paper addresses the puzzle of regime-dependent volatility in foreign exchange. We extend the literature in two ways. First, our microstructural model provides a qualitatively new explanation for the puzzle. Second, we test implications of our model using Europe's recent shift to rigidly fixed rates (EMS to EMU). In the model, shocks to order flow induce volatility under flexible rates because they have portfolio-balance effects on price, whereas under fixed rates the same shocks do not have portfolio-balance effects. These effects arise in one regime and not the other because the elasticity of speculative demand for foreign exchange is (endogenously) regime-dependent: low elasticity under flexible rates magnifies portfolio-balance effects; under credibly fixed rates, elasticity of speculative demand is infinite, eliminating portfolio-balance effects. New data on FF/DM transactions show that order flow had persistent effects on the exchange rate before EMU parities were announced. After announcement, determination of the FF/DM rate was decoupled from order flow, as predicted by the model.

William P. Killeen
Setanta Asset Management
College Park House
South Frederick Street
Dublin 2 Ireland

Richard K. Lyons
Haas School of Business, S545
U.C. Berkeley
Berkeley, CA 94720-1900
and NBER
Tel: 510-642-1059
Fax: 510-642-4700
Email: lyons@haas.berkeley.edu

Michael J. Moore
School of Mgmt and Economics
Queen's University, Belfast
Belfast BT7 1NN
Northern Ireland, UK

Fixed versus Flexible: Lessons from EMS Order Flow

1. Introduction

If there is a topic at the center of international macroeconomics, it is fixed versus flexible exchange rates. Whether teaching Mundell-Fleming, speaking about the Impossible Trinity, or writing about “excess” volatility, the fixed-versus-flexible debate is deeply relevant. At the same time, many of the issues in this debate remain unresolved. Important among them is the regime-volatility puzzle: similar macroeconomic environments produce much more exchange-rate volatility under flexible-rate regimes (e.g., Flood and Rose 1995).¹ This finding has prompted many authors to conclude that the critical determinants of flexible-rate volatility are not macroeconomic. It is unclear, however, what these non-macro determinants might be.

This paper addresses the regime-volatility puzzle from a new perspective. Our approach augments the traditional macro-asset approach with some price-determination microeconomics. In particular, we focus our theoretical and empirical analysis on the role of order flow. (Order flow is signed volume; seller-initiated trades are negative order flow and buyer-initiated trades are positive order flow.) In microeconomic models of asset-price determination, order flow plays an important causal role that arises because order flow conveys information. The type of information that order flow conveys includes any information that is relevant to the realization of uncertain demands, so long as that information is not common knowledge. (If common knowledge—as is generally assumed in macro exchange-rate models—then price adjusts without any need for an order-flow role.) For example, order flow may convey information about dispersed shifts in portfolio balance (e.g., shifts in hedging demands or risk tolerances), or about differential interpretation of common news. In a non-common-knowledge setting, order

¹ To some, lower volatility under fixed rates may seem obvious. But empirically, fixed rates do have a distribution over time (due to parity changes). And in most models, keeping the variance of this distribution below that under flexible rates requires keeping the variance of fundamentals below that under flexible rates. As an empirical matter, this has not been the case, per Flood and Rose (1995). Our explanation provides a source of volatility that operates only under flexible rates, but is not included among the fundamentals previously considered.

flow becomes the intermediate link between evolving information and price—a proximate cause of price movements (Evans and Lyons 1999).

Order flow is potentially relevant for a high-level debate like fixed-versus-flexible because its impact on exchange rates is persistent. Though our model makes this persistence explicit, let us provide some perspective. Note that if order flow conveys information, then its price impact *should* persist, at least under the identifying assumption used regularly in empirical work that information effects on price are permanent (e.g., French and Roll 1986, Hasbrouck 1991). That is not to suggest that order flow cannot also have transitory effects on price, e.g., from temporary indigestion effects (sometimes called inventory effects). But insofar as order flow communicates information—along the lines noted in the previous paragraph—then some portion of order flow’s effects on price will persist.

Order flow may be key to resolving the regime-volatility puzzle. One reason it has not been considered to date is because empirical work on this puzzle examines macro determinants, whereas order flow is generally viewed as non-macro.² If order flow is indeed a determinant, then it might explain why flexible regimes produce more volatility than macroeconomics predicts. There is now considerable evidence that the “if” part of that last sentence is met: order flow *is* a determinant (Lyons 1995, Evans 1999, Payne 1999, Evans and Lyons 1999, Rime 2000). Whether the relation found by these authors for flexible rates is affected by differences in exchange-rate regime is an open question, however, one that we address in this paper (both theoretically and empirically).

The main lesson from the theoretical portion of the paper is the following: exchange rates are more volatile under flexible rates because of order flow. Importantly, this is not because order flow is more volatile under flexible rates (indeed, its volatility is unchanged in our model across regimes). The intuition for order flow’s role is tied to the

² Though generally viewed as non-macro, order flow need not be divorced from macro. Consider models in microstructure finance, where order flow signals that some agents’ expectations about future fundamentals have changed. For exchange rates, one can write the price of foreign exchange, P_t , as a function of current and expected future macro fundamentals: $P_t = g(f_t, f_{t+1}^e)$. Order flow may be the way that price-setters learn about changes in f_{t+1}^e . Under this hypothesis, order flow is a proximate cause of price changes, but macro (macro expectations) remains the underlying cause. Note too that this hypothesis offers a potential explanation for the empirical results of Meese and Rogoff (1983): if the macro variables that order flow forecasts are largely beyond the one-year horizon, then the empirical link between exchange rates and *current* macro variables will be loose. That Meese-Rogoff-style empirical results are more positive at horizons beyond one year is consistent with this “anticipation” hypothesis.

elasticity of speculative demand. Under flexible rates, the elasticity of speculative demand is (endogenously) low: volatility causes rational speculators to trade less aggressively. Their reduced willingness to take the other side of shocks to order flow makes room for portfolio-balance effects on price. The size of those portfolio-balance effects is determined by the size of the observed order-flow shocks. This is the price-relevant information that order-flow conveys. Under credibly fixed rates, the elasticity of speculative demand is infinite (return volatility shrinks to zero), which precludes portfolio-balance effects, thereby eliminating order flow's information role. Consequently, order flow as a return factor is shut down.

To test our explanation for regime-dependent volatility empirically, we exploit a natural experiment for why order flow can induce volatility under flexible rates, but not under fixed rates. The natural experiment is the switch from flexible (wide band) to rigidly fixed rates in the transition from the European Monetary System (EMS) to the European Monetary Union (EMU).³ Starting in January 1999, the euro-country currencies have been rigidly fixed to one another—as close to a perfectly credible fixed-rate regime as one might hope to observe. Before January 1999, however—particularly before May 1998—there was still uncertainty about which internal parities would be chosen and about the timing of interest-rate harmonization in the May-to-December period.

Figure 1 provides an initial, suggestive illustration of our results. It shows the relationship between the FF/DM exchange rate in 1998 and cumulative order flow. (These are interdealer orders; see section 3 for details.) In panel A, order flow is measured as the number of buyer-initiated trades minus the number of seller-initiated trades (consistent with Evans and Lyons 1999). In panel B, order flow is measured in DM amounts rather than numbers of trades. The vertical line is 4 May 1998. This was the first trading day after the announcement of the internal parities of the euro-participating currencies.⁴ The positive relationship between the two series before that date is clear: the

³ Though it allowed some flexibility, the EMS was not a free float. That said, the transition to EMU was indisputably a transition toward exchange-rate fixity. Low variability of the FF/DM rate in the EMS portion of our sample (relative to major flexible rates such as Yen/\$) does not undermine the validity of our tests. (Variability in that portion of our sample was certainly high enough to be significant economically for market participants, given the low transaction costs.) Extending the model of the next section to environments of imperfectly credible fixed rates is a natural direction for further research.

⁴ On the weekend of May 2/3, 1998, the Council of EU finance ministers (ECOFIN) decided which countries would become founder members of the euro. They also decided on the internal parities for the euro-zone currencies.

correlation is about 0.7. (This accords with the strong positive relationship under flexible rates found for other currencies and samples, e.g., by Evans and Lyons 1999 and Rime 2000). After 4 May, however, there is a sharp unwinding of long DM positions with no corresponding movement in the exchange rate. In fact, there is a negative correlation during the second period. Though total variation in the exchange rate is small (roughly 20 times the median bid-offer spread), the effect of order flow appears to have changed from one of clear impact—as has been found in other studies for flexible rates—to one of no impact.⁵ (Visually it appears the relationship is loosening in April, but the statistical evidence of a break we provide below does not emerge until early May.) The model we develop in the following section provides a framework for addressing why these order-flow effects might disappear, and how their disappearance helps to resolve the regime-volatility puzzle.

Our paper is the first in the literature to address fixed versus flexible rates using the concept of order flow (despite the concept's long history within the field of finance). The two most closely related bodies of work on fixed-versus-flexible rates include the balance-of-payments flow approach and a more recent literature that introduces non-rational traders to account for high flexible-regime volatility. Work on the balance-of-payments flow approach dates back to Robinson (1949) and Machlup (1949). (See also the survey in Rosenberg 1996.) In those models, exchange rates are determined from balance-of-payments flows, e.g., imports and exports. Balance-of-payments flows depend, in turn, on the exchange-rate regime. Empirically, however, this approach has not borne fruit: balance-of-payments flows are unsuccessful in accounting for exchange-rate movements (see, e.g., Frankel and Rose 1996).⁶ The second related body of work introduces non-rational traders to account for high flexible-regime volatility, for example, Hau (1998) and Jeanne and Rose (1999). In both of these papers, the mechanism that induces volatility under flexible rates is entry of traders in the foreign-exchange market, which increases the level of noise. The mechanism in our (rational) model is different. In

⁵ A test of whether the variance is equal across the two sub-periods is rejected at the 5 percent level.

⁶ These negative empirical results are less applicable to the noise-trade models of Osler (1998) and Carlson and Osler (2000) because the “current account” flows in those models can be interpreted more broadly (e.g., as flows from hedging demand), which need not manifest as identifiable balance-of-payments flows. Importantly for our approach here, order flows and balance-of-payment flows are not one to one; see, e.g., the discussion in Lyons (2001), chapter 9.

our case it is the endogenous unwillingness of speculators to take the other side of order-flow shocks that produces higher volatility.

The remainder of this paper is organized as follows. Section 2 presents our model and the analytical results that provide a new explanation for the regime-dependent volatility puzzle. Section 3 describes our data. Section 4 presents empirical tests of our model's implications. Section 5 concludes.

2. Model

The trading model developed in this section serves three important purposes. First, it expresses market dynamics in terms of measurable variables, most notably order flow.⁷ The type of order flow that is currently measurable (the type shown in Figure 1) is interdealer order flow, which necessitates a specification of interdealer trading behavior and how that behavior relates to underlying demands in the economy. Second, the model provides a clear demonstration of the type of dispersed information that order flow can convey, and how that information is subsequently impounded in price. Third, the model shows why under flexible rates order flow (cumulative) and price share a long run, cointegrating relationship. The properties of this relationship and how it is affected by the shift to fixed rates provide a set of testable implications that we examine empirically in section 4.

The model includes two trading regimes: a flexible-rate regime followed by a fixed-rate regime. The shift from flexible to fixed rates is a random event that arrives with constant probability p at the end of each trading day (after all trading).⁸ Once the regime has shifted to fixed rates it remains there indefinitely. Though a regime shift is not included in the model of Evans and Lyons (1999), our specification of trading within each day is identical to that earlier specification, so our exposition below is fullest in areas where the models differ.

⁷ This is not a property of more abstract approaches to modeling trading, for example, rational expectations models (such as Grossman and Stiglitz 1980 and Diamond and Verrecchia 1981), which are not directly estimable. (In rational expectations models, trades cannot be identified as a buy or a sell because counterparties are symmetric—neither side is the initiator.)

⁸ This formulation has two important advantages. First, the effective horizon over which foreign exchange is priced in the flexible-rate regime remains constant. Second, the parameter p provides a compact means of describing regime shifts as far or near. As an empirical matter, particularly in the context of the EMS-EMU transition, this specification serves as a convenient abstraction from reality.

Consider an infinitely lived, pure-exchange economy with two assets, one riskless and one with stochastic payoffs (foreign exchange). Each day t , foreign exchange earns a payoff R_t , which is composed of a series of increments, so that

$$R_t = \sum_{\tau=1}^t \Delta R_\tau \quad (1)$$

The increment ΔR_t is observed publicly on day t before trading. These realized increments represent innovations over time in public macroeconomic information (e.g., changes in interest rates). Under the flexible-rate regime, the ΔR_t increments are i.i.d. $\text{Normal}(0, \sigma_R^2)$. On the first morning of the fixed rate regime, the central bank (credibly) commits to pegging the exchange rate at the previous day's closing price and maintains $\Delta R_t=0$ thereafter.

The foreign exchange market is organized as a dealership market with N dealers, indexed by i , and a continuum of non-dealer customers (the public). The mass of customers on $[0,1]$ is large (in a convergence sense) relative to the N dealers. (This assumption will drive the model's overnight risk-sharing features.) Dealers and customers all have identical negative exponential utility defined over periodic wealth, with coefficient of absolute risk aversion θ .

The timing of trading within each day is summarized in Figure 2. Within each day there are three trading rounds. 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 throughout the economy).

Each day begins with public observation of the payoff R_t . Then each dealer quotes a scalar price to his customers at which he agrees to buy and sell any amount (quoting is simultaneous and independent).⁹ We denote this round 1 price of dealer i on day t as P_{it}^1 . Each dealer then receives a customer-order realization C_{it}^1 that is executed at his quoted

⁹ Introducing a bid-offer spread (or price schedule) in round one to endogenize the number of dealers is a straightforward—but distracting—extension.

price P_{it}^1 . Let $C_{it}^1 < 0$ denote net customer selling (dealer i buying). The individual C_{it}^1 's are distributed

$$C_{it}^1 \sim \text{Normal}(0, \mathbf{s}_C^2)$$

They are uncorrelated across dealers and uncorrelated with the payoff R_t at all leads and lags. These orders represent exogenous portfolio shifts of the non-dealer public. Their realizations are not publicly observable, and they arrive every day, regardless of regime (with an unchanged distribution). For the analysis below, it is useful to define the aggregate public demand in round 1 as:

$$C_t^1 = \sum_{i=1}^N C_{it}^1$$

In round 2, each dealer quotes a scalar price to other dealers at which he agrees to buy and sell any amount (simultaneous and independent). These interdealer quotes are observable and available to all dealers. Each dealer then trades on other dealers' quotes (simultaneous and independent). Orders at a given price are split evenly across dealers quoting that price. Let T_{it} denote the (net) interdealer trade initiated by dealer i in round 2 (we denote T_{it} as negative for dealer- i net selling). At the close of round 2, all dealers observe the interdealer order flow X_t from that period:

$$X_t = \sum_{i=1}^N T_{it} . \tag{2}$$

In round 3 of each day, 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 quotes a scalar price P_{it}^3 at which he agrees to buy and sell any amount (simultaneous and independent). These quotes are observable and available to the public. We assume that aggregate public demand for the risky asset in round-3, denoted C_t^3 , is less than infinitely elastic. With the earlier assumptions, this allows Evans and Lyons to write public demand as a linear function of expected return:

$$C_t^3 = \mathbf{g} E \left[\Delta P_{t+1} + R_{t+1} \mid \Omega_t^3 \right] \quad (3)$$

where

$$\mathbf{g} \propto \text{Var}^{-1} \left[\Delta P_{t+1} \mid \Omega_t^3 \right]$$

Thus, the positive coefficient γ captures the elasticity of public demand—the public’s aggregate willingness to absorb exchange rate risk. The information in Ω_t^3 is that available to the public at the time of trading in round three of day t (which includes all past R_t and X_t).

Equilibrium

The equilibrium relation between interdealer order flow and price adjustment follows from results established for the simultaneous-trade model of Lyons (1997). Consider the determination of prices. Propositions 1 and 2 of that paper show that each trading round, all dealers quote a common price (this is necessary to prevent arbitrage). It follows that this price is conditioned on common information only. Though R_t is common information on day t at the beginning of round 1, order flow X_t is not observed until the end of round 2. The price for round-3 trading, P_t^3 , reflects the information in both R_t and X_t . The payoff information in R_t is straightforward.

The information in X_t is as follows. In equilibrium, each dealer’s interdealer trade, T_{it} , is proportional to the customer order flow C_{it}^1 he receives in round-one. This implies that when dealers observe X_t at the end of round 2 (equation 2), they can infer the aggregate portfolio shift on the part of the public in round 1, C_t^1 . Dealers also know that, for a risk-averse public to re-absorb this portfolio shift in round 3, price must adjust—a portfolio-balance effect. In particular, price must adjust in round 3 so that $C_t^3 + C_t^1 = 0$, where C_t^3 is given by equation (3).

The resulting price level at the end of day t can be written as:

$$P_t = \begin{cases} I_1 \sum_{t=1}^t \Delta R_t + I_2 \sum_{t=1}^t X_t & \text{under flexible rates (t} \leq \mathbf{T}) \\ I_1 \sum_{t=1}^T \Delta R_t + I_2 \sum_{t=1}^T X_t + I_3 \sum_{t=T+1}^t X_t & \text{under fixed rates (t} > \mathbf{T}) \end{cases} \quad (4)$$

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

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

$$\Delta P_t = I_1 \Delta R_t + I_2 X_t \quad (5)$$

where I_1 and I_2 are positive constants.¹⁰ The portfolio-balance effects from order flow enter through I_2 , which depends inversely on γ —the elasticity of public demand in equation (3)—and also on the variances \mathbf{s}_R^2 and \mathbf{s}_C^2 (see Evans and Lyons 1999 for details).¹¹ The parameter I_2 is commonly referred to as a price impact parameter (it governs the price impact of order flow).

The connection between equation (4) and traditional measures of macro fundamentals deserves some attention. For traditional measures, it is useful to distinguish between “narrow” fundamentals and “broad” fundamentals. Under the monetary macro approach, the set of variables considered fundamental (e.g., money supplies, interest rates, and output levels) do not include variables that affect equilibrium risk premia

¹⁰ Note that we have not yet added an error term to either equation (4) or equation (5). The cointegration model we estimate in section 4 adds a stationary (but not necessarily white noise) error term to equation (4). When differenced, this adds a non-invertible moving average term to equation (5), which represents the error-correction mechanism we estimate.

¹¹ The probability p of the regime shift adds a parameter to the Evans-Lyons solution that has no qualitative impact on the coefficients of interest here, namely λ_2 and λ_3 . This model can also be used to generate multiple equilibria. Introducing multiple equilibria obscures the essential portfolio-balance logic, so we do not pursue this direction.

because monetary models do not admit risk premia. Fundamental variables under this approach constitute the set of narrow fundamentals. In contrast, the portfolio-balance macro approach does admit risk premia, so variables affecting these premia become fundamental drivers of exchange rates under this approach (e.g., changes in hedging demands, risk tolerances, or asset supplies). When added to the set of narrow fundamentals, these variables affecting risk-premia define the set of broad fundamentals. In terms of equation (4), the set of narrow fundamentals includes the payoff terms, ΔR_t , but does not include the portfolio-balance terms, X_t . The set of broad fundamentals includes all the terms (in keeping with the idea that equilibrium determination of risk premia is fundamental to asset pricing). As we show next, a change in regime does not change which variables are included within broad fundamentals, but it does alter the price response to variables in that set (specifically, the price response to a given amount of order flow—the coefficient on X_t in equation 4).

Differences Across Trading Regimes

Understanding the effects of the different trading regimes—and the changing role of order flow—comes from the effect of the exchange-rate regime on equations (2) and (3). Specifically, the parameter γ , which represents the elasticity of public demand, is regime-dependent. This comes from the regime-dependence of $\text{Var}[\Delta P_{t+1} | \Omega_t^3]$ (γ being proportional to the inverse of this variance, per equation 3). The elimination of portfolio-balance effects under fixed rates reduces this variance (shown below), implying that:

$$g_{\text{flexible}} < g_{\text{fixed}} \tag{6}$$

Public demand is therefore more elastic in the (credible) fixed-rate regime than the flexible-rate regime. The implication for the price impact parameters I_2 and I_3 in equation (4)—henceforth I_{flexible} and I_{fixed} respectively—is the following:

$$I_{\text{flexible}} > I_{\text{fixed}} \tag{7}$$

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

$$I_{\text{fixed}} = 0 \quad (8)$$

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

Intuition for Cointegration in Equation (4)

Consider P_{T+1} , the price at the close of the first day of the fixed-rate regime. Foreign exchange is a riskless asset at this point, with return variance equal to zero. A return variance of zero implies that the elasticity of the public's speculative demand is infinite, and the price impact parameter λ_3 in equation (4) equals zero. This yields a price at the close of trading (round 3) on day T+1 of:

$$P_{T+1} = I_1 \sum_{t=1}^T \Delta R_t + I_2 \sum_{t=1}^T X_t$$

The summation over the payoff increment ΔR_t does not include an increment for day T+1 because the central bank maintains ΔR_t at zero in the fixed regime. Though X_{T+1} is not equal to zero, this has no effect on price because $\lambda_3=0$, as noted. This logic holds throughout the fixed-rate regime.

To understand why order flow has price impact under flexible rates, consider the price at the close of trading on the final flexible-rate day, day T. According to equation (4):

$$P_T = I_1 \sum_{t=1}^T R_t + I_2 \sum_{t=1}^T X_t$$

At the close of day T , the payoff R_{T+1} is uncertain because ΔR_{T+1} is uncertain. (Recall that the shift in regime is random and not yet known at the close of trading on day T . Note too that forward-looking variables do not enter this expression due to our simple specification of ΔR_t and C_t^1 as independently distributed across time with mean zero.) The second term is the portfolio balance term. To understand this term, note that in the Evans-Lyons flexible-rate model, each dealer's trading rule is:

$$T_{it} = \alpha C_{it}^1$$

where α is a positive coefficient. This implies that:

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

Therefore, we can write:

$$\sum_{t=1}^T X_t \propto \sum_{t=1}^T C_t^1$$

The sum of the portfolio shifts C_t^1 over time represents changes in “effective” asset supply: exogenous shifts out of foreign exchange (by some agents for non-speculative purposes) are an increase in the net supply that the speculating public must re-absorb. (We couch this in terms of supply to connect with traditional portfolio-balance intuition; see footnote 12.) The total increase in net supply is the sum of past portfolio shifts out of foreign exchange,

$$\text{Increase in net supply} = - \sum_{t=1}^T C_t^1$$

where the minus sign arises because an aggregate shift out of foreign exchange corresponds to C_t^1 being negative.

As is standard in portfolio balance models, increases in supply lower price, and decreases in supply raise price. This is why a positive cumulative X_t in equation (4) raises price: if cumulative X_t is positive, this implies that cumulative C_t^1 is also positive, which is a decrease in net supply, requiring an increase in price to clear the market. X_t is the variable that conveys this information about the decrease in net supply (C_t^1 is unobservable). The round-three price on day T, P_T , depends on the sum of the X_t because each additional decrease in supply C_t^1 requires an incremental increase in price. As payoff uncertainty shrinks to zero (as in the fixed-rate regime), the arrival of new X_t no longer induces portfolio balance effects.¹²

In our model, a credible fixed-rate regime is one in which the private sector, not the central bank, absorbs innovations in order flow. This theoretical point contrasts with extant models where central bank reserves are used to absorb innovations in order flow, but at a cost (e.g., selling reserves at domestic-currency prices that are too low; see, e.g., Guembel and Sussman 2001). Empirically, as we turn attention to testing implications of our model below, our understanding is that there was little intervention by the national central banks or the ECB in the period from May to December, 1998 (per conversations with ECB and national central bank officials—hard data are not available). The Bretton Woods era, too, provides many periods in which exchange-rate volatility was quite low, yet central banks were intervening very little (relative to the size of the market).

3. Data

Our data set includes the daily value of purchases and sales in the FF/DM market for twelve months, January to December 1998. Electronic Broking Services (EBS), the

¹² Here is a simple one-period example that illustrates the basic economics of the model. An uncertain payoff R is realized at time 1 and the market-clearing gap $E[R]-P_0$ will be a function of the risky asset's net supply. In traditional portfolio balance models, demand D is a function of relative returns, and supply S is time-varying. That is: $D(E[R]-P_0)=\tilde{S}$ where the tilde denotes random variation. Our model looks different. In our model (gross) supply is fixed. But what we are calling "net supply" is moving over time, due to demand shifts that are unrelated to $E[R]-P_0$. These demand shifts are the realizations of C_t^1 (which one could model explicitly as arising from hedging demands, liquidity demands, or changing risk tolerances). Conceptually, our model looks more like: $D(E[R]-P_0, \tilde{C})=\bar{S}$ where \bar{S} denotes fixed gross supply and \tilde{C} denotes shifts in net supply, that is, shifts in demand unrelated to $E[R]-P_0$. In this one-period example, the higher the $t=0$ realization of \tilde{C} , the lower the net supply to be absorbed by the rest of the public, and the higher the market-clearing price P_0 (to achieve stock equilibrium). In a sense, our multi-period model is akin to a single-period model in which net supply is "shocked" multiple times before trading takes place, each shock having its own incremental impact on price.

leading foreign-exchange broker, provided the data. Each trading day (weekday) covers the twenty four-hour period starting from midnight GMT.

A brief overview of the market structure may be useful for understanding the data and evidence.¹³ There are three types of trades in the forex market: customer-dealer trades, direct interdealer trades, and brokered interdealer trades. Customers are non-financial firms and non-dealers in financial firms (e.g., corporate treasurers, hedge funds, mutual funds, pension funds, proprietary trading desks, etc.). Dealers are marketmakers employed in banks worldwide (the largest 10 dealers account for more than half of the volume in major currencies). At the time of our sample, these three trade types accounted for roughly equal shares of total volume in major markets—one-third each.

Our data come from the third trade type: brokered interdealer trading. There are two main interdealer broking systems, EBS and Reuters Dealing 2000-2. EBS has a much larger market share than Dealing 2000-2 in most currency pairs that do not involve sterling (including FF/DM). We estimate that our EBS sample amounts to about 18% of all trading in spot FF/DM in 1998. This estimate is based on comparing our EBS volume data with data provided by the BIS (1999) for April 1998—one of the months in our sample. The 18% estimate is based on the following calculations. For the FF/DM rate, the BIS-reported total daily volume for April 1998 was \$10.1 billion (which includes spot, forwards, and swaps).¹⁴ Spot transactions account for 80% of this volume or \$8.1 billion.¹⁵ In our EBS data, average daily spot volume in FF/DM for April 1998 amounted to \$1.5 billion. Our EBS volume of \$1.5 billion divided by the BIS total of \$8.1 billion produces our 18% estimate.

Three features of our data set are noteworthy. First, it spans a considerably longer period than previous data on order flow. For example, Evans and Lyons (1999) use four months of data. Second, the brokered interdealer trading it covers is the most rapidly growing category of trade (see BIS, 2001), and anecdotal evidence suggests that EBS is

¹³ For more detail see Lyons (2001) and the EBS website at www.ebsp.com. Trading in other European cross rates on EBS was not thick enough to estimate our model on a panel of currencies.

¹⁴ According to BIS (1999), the FF/DM currency pair accounted for 0.7% of world forex volume in April 1998, which in total amounted to about \$1.5 trillion per day.

¹⁵ The 80% estimate comes from the following. According to BIS (1999), spot transactions amounted to 75% of trade between the DM and other EMS currencies in April 1998. In April 1995, this figure was 86% for the FF/DM market. Unfortunately, no information on this percentage is provided specifically for the FF/DM market for April 1998; 80% is derived as an interpolation.

also increasing its market share. Third, our data include daily order flow measured in terms of DM value. The Evans (1997) data, also used by Evans and Lyons (1999), include only order flow measured as the number of buys minus the number of sells (e.g., a sale of any DM amount is measured as -1).¹⁶ This allows us to use a measure of cumulative order flow that matches our model, namely:

$$\sum_{t=1}^t X_t = \sum_{t=1}^t (B_t - S_t)$$

where B_τ and S_τ are the DM values of buyer-initiated and seller-initiated orders on day τ , respectively.

The rest of the data is measured as follows. The spot exchange rate is measured as the French franc price of a DM and is sampled daily at the close of business in London.¹⁷ The interest differential is calculated from the daily Paris (PIBOR) and Frankfurt (FIBOR) monthly interbank offered rates.¹⁸ These rates are also sampled at the London close. Both exchange rate and interest rate data are from Datastream.

Institutional Details on EBS

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

¹⁶ Per figure 1, our data set allows us to construct both of these two measures. They behave quite similarly. For example, the correlation between the two order-flow measures in the flexible-rate portion of our sample (January to April) is 0.98.

¹⁷ We know of no source for transaction prices sampled at midnight GMT. Even if a source were available, it would be relatively noisy because the market is quite thin around that time, with wide bid-offer spreads (which generate measurement error due to transaction prices bouncing from bid to ask). Most all trading in FF/DM is carried out before the London close (there is relatively little trading in this currency pair in New York). Finally, it is not feasible for us to re-measure our daily order flow as of the London close because we do not have the flow data on an intraday basis.

¹⁸ We use monthly interest rates because of the noise in the daily variation of overnight rates. Our results are not affected in any substantive way by this choice.

There are two ways that dealers can trade currency on EBS. Dealers can either post prices (i.e., submit “limit orders”), which does not insure execution, or dealers can “hit” prices (i.e., submit “market orders”), which does insure execution. To construct a measure of order flow, trades are measured as positive or negative depending on the direction of the market order that initiates each transaction. For example, a market order to sell DM 10 million that is executed against a posted limit order would generate order flow of –10 million Deutschemarks.

When a dealer submits a limit order, she is displaying to other dealers an intention to buy or sell a given cash amount at a specified price.¹⁹ Bid prices (limit order buys) and offer prices (limit order sells) are submitted with the hope of being executed against the market order of another dealer—the “initiator” of the trade. To be a bit more precise, not all initiating orders arrive in the form of market orders. Sometimes, a dealer will submit a limit-order buy that is equal to or higher than the current best offer (or will submit a limit-order sell that is equal to or lower than the current best bid). When this happens, the incoming limit order is treated as if it were a market order, and is executed against the best opposing limit order immediately. In these cases, the incoming limit order is the initiating side of the trade.

4. Results

The analytical results in section 2 offer five testable hypotheses that we collect here as a guide for our empirical analysis:

Hypothesis 1: Under flexible rates, the level of the exchange rate, cumulative interdealer order flow, and cumulative public payoff information are individually nonstationary and jointly cointegrated.

Hypothesis 2: Cumulative interdealer order flow remains nonstationary after the shift from flexible to fixed rates, but the level of the exchange rate and cumulative public payoff information become stationary.

Hypothesis 3: A structural break in the cointegrating relationship occurs when the regime shifts from flexible to fixed rates.

¹⁹ EBS has a pre-screened credit facility whereby dealers can only see prices for trades that would not violate their counterparty credit limits, thereby eliminating the potential for failed deals because of credit issues.

Hypothesis 4: Under flexible rates, error correction in the cointegrating relationship occurs through exchange rate adjustment, not order flow adjustment (weak exogeneity of order flow).

Hypothesis 5: Under flexible rates, there is no Granger causality from the exchange rate to order flow (strong exogeneity of order flow).

Hypotheses 1-3 summarize the section-two discussion of equation (4). Hypotheses 4 and 5 follow from the model's specification of round-one public order flow as exogenous. This particular assumption of the model is a strong one; our cointegration framework provides a natural way to test its implications.

The empirical analysis proceeds in two stages. First, we address hypotheses 1-3 by testing for unit roots, cointegration, and structural breaks. This first stage also examines the related issue of coefficient size within the cointegrating relationship. The second stage addresses hypotheses 4 and 5 by estimating the appropriate error-correction model and (separately) testing for Granger causality.

4.1 Stage 1: Testing Hypotheses 1 to 3

Let us begin by repeating equation (4) from the model, which establishes the relationship between the level of the exchange rate P_t , a variable summarizing public information about payoffs ($\Sigma\Delta R_t$), and accumulated order flow (ΣX_t).

$$P_t = \begin{cases} I_1 \sum_{t=1}^t \Delta R_t + I_2 \sum_{t=1}^t X_t & \text{under flexible rates (t} \leq \text{T)} \\ I_1 \sum_{t=1}^T \Delta R_t + I_2 \sum_{t=1}^T X_t + I_3 \sum_{t=T+1}^t X_t & \text{under fixed rates (t} > \text{T)} \end{cases} \quad (4)$$

Like Evans and Lyons (1999), we use the interest differential (PIBOR-FIBOR) as our

measure of cumulative public information about foreign-exchange payoffs.²⁰

Stationarity

The first step is to test for the stationarity of all variables in the two sub-periods of 1998, January 1 to May 1 and May 4 to December 31. Table 1 shows the results of six Dickey-Fuller tests, one for each of the three variables in each of the two sub-periods. Consistent with hypothesis 1, in the first four months of 1998 all variables appear non-stationary: the unit-root null cannot be rejected.²¹ Though predicted by our model, cumulative order flow being non-stationary is not obvious; indeed, a common intuition is that market clearing would produce cumulative order flow that rapidly reverts to zero.

Consistent with hypothesis 2, for the remaining eight months the exchange rate and interest differential appear stationary (unit-root null is rejected), whereas cumulative order flow remains non-stationary. The combination in the latter period of a stationary exchange rate and non-stationary cumulative order flow is consistent with a price impact parameter λ_3 of zero, as predicted by the model. It remains to be determined whether equation (4) actually holds for the January 1 to May 1 period, i.e., whether the variables are cointegrated (which we return to below).

The bottom panel of Table 1 speaks to hypothesis 2 by implementing a univariate test for a structural break in the spot rate process. The two tests shown, due to Banerjee, Lumsdaine, and Stock (1992), are described in more detail in the table legend. The null for both of these tests is that spot-rate process is nonstationary over the whole sample period with no structural breaks in the constant or trend. In both cases, the null is rejected. We return to a direct test of stability in the cointegrating relationship below, following the cointegration results. (See also the appendix for results from implementing the Rigobon 1999 test for parameter stability in settings with heteroskedasticity, endogeneity, and omitted variables.)

²⁰ At the daily frequency, the interest differential is arguably the best public-information measure of changing macroeconomic conditions. That said, it does not constitute a well-specified macro model of expected payoffs. We return to the measurement error this entails below.

²¹ Though the possibility always remains that these variables are not non-stationary, these tests indicate that they are at the very least close to being so (in a statistical sense), in which case proceeding with our cointegration analysis as if they are non-stationary is the safe route (e.g., in terms of consistency).

Cointegration

Because the log exchange rate, cumulative order flow, and the PIBOR-FIBOR interest rate differential are all non-stationary in the flexible-rate period from January to May 1, equation (4) only holds if they are cointegrated. We use two different procedures to test for cointegration: the Granger-Engle ADF test and the Johanson test. Both begin with a baseline model that includes the three variables of our model of section two, a constant, and a trend. We determine the lag length for the Johansen vector autoregression using Sims' likelihood ratio tests; a lag length of four allows for all significant dynamic effects (results available from authors on request).

Hypothesis 1 is borne out: evidence for cointegration (rejecting the null of no-cointegration) over the flexible-rate period is strong. Table 2 presents results for the Granger-Engle ADF test. In this test, the log spot rate is regressed against cumulative order flow, the interest differential, a constant, and a trend. (We use Phillips-Hansen fully modified estimation. This is a semi-parametric technique that gives more accurate point estimates in small samples by diminishing second-order asymptotic bias. It also yields consistent standard errors. The residuals from the regression are then tested for stationarity using conventional Augmented Dickey-Fuller tests.) The null of non-cointegration is rejected at the 5 percent level, per the Dickey-Fuller tau statistics at the bottom. The top of the table shows the estimated coefficients (constant not reported). Note that the interest differential is insignificant. We do not take this insignificance to mean that macro does not matter (though it is in keeping with the results of Meese and Rogoff 1983 and the long empirical literature that followed that paper). It may be due to the measurement error (or mis-specification) inherent in our use of the interest differential for the variable R_t from the model. The cointegration we find between the exchange rate and cumulative order flow demonstrates that the model is indeed able to account for a steady state relationship. Having found significance, henceforth we focus on the bivariate cointegrating relationship between the exchange rate and cumulative order flow.²²

²² The second of our cointegration tests—using the Johanson procedure—tests the null hypothesis of no-cointegration in the bivariate model against the alternative of a single cointegrating vector. That test is rejected at the 5 percent level (not reported). The Johanson procedure also suggests that there is only one cointegrating vector: A Johanson test of the null of one cointegrating vector against the alternative of two is not rejected at the 5 percent level. In independent work, Bjonnes and Rime (2001) also find evidence of cointegration between cumulative interdealer order flow and the level

Next we estimate the magnitude of the coefficients in the bivariate cointegrating relationship (flexible-rate period). These are reported in Table 3. Of particular importance is the coefficient on cumulative order flow.²³ Our use of the log spot rate makes this coefficient easy to interpret: an increase in cumulative order flow of DM 1 billion (i.e., net DM purchases) increases the French franc price of a DM by 3 basis points. This is much smaller than the contemporaneous impact of order flow estimated by Evans and Lyons (1999) for the DM/\$ market (roughly 50 basis points per \$1 billion).²⁴ Three factors are likely to account for this. First, one must correct for the fact that the DM price of a dollar at the time was roughly 1.5 DM/\$. This scales the Evans-Lyons coefficient down by about a third. Second, Evans and Lyons are measuring the impact effect, whereas we are measuring the long-run impact (i.e., the persistent impact).²⁵ To the extent any of the impact effect is transitory, this will account for part of the difference. Finally, and perhaps most important in the context of our model, one would expect the sensitivity of price to order flow to be lower in the FF/DM market precisely because the EMS target zones were not a free float (and therefore the elasticity of absorptive private demand is higher).

We turn now to evaluating hypothesis 3: a structural break should occur in the cointegrating relationship when the regime shifts from flexible to fixed rates. Results from two different tests appear in Table 4. The first, the F-test (see Hanson 1992), is akin to a non-stationary analog of the familiar Chow test in stationary settings. Importantly, it is valid for settings in which the date of the structural break is known in advance. Using May 4, 1998, as a known break date, the F-test rejects the null of no structural break over the year at the 5 percent level (column two), consistent with hypothesis 3. For robustness, we also tested for a structural break in the cointegrating relationship over the full year using a test that does not rely on knowing the break date in advance. That test, the Mean-

of the spot exchange rate (they examine the DM/\$ and Norwegian Krona/DM markets). The relationship is not the focus of their paper, however; it is addressed in a closing subsection.

²³ A reassuring feature is the similarity between the modified least squares estimates in the trivariate model of Table 2 and the maximum likelihood bivariate estimates of Table 3.

²⁴ Our use of the logged spot rate in estimation (rather than unlogged) makes our work directly comparable to earlier work. It also squares with the empirical distribution of exchange rates, which is (approximately) Lognormal.

²⁵ One might be tempted to check robustness of our result that order flow effects persist by regressing the *level* of the spot rate on lags of unaccumulated order flow (i.e., past daily flows). The notion being that more distant order flow lags might be negative. Econometrically, however, this is an unbalanced regression, mixing non-stationary and stationary variables, thereby undermining inference. Cointegration tests are the proper way to resolve this issue.

F, also rejects the null of no structural break at the 5 percent level (column two). The Mean-F test can also be applied to specifically to the flexible-rate portion of our sample. There is no evidence of structural instability in the bivariate relationship over the flexible period from January through May 1 (p-value >20%). (For visual evidence that order flow's price impact had dropped off by May 4, see the appendix.)

4.2 Stage 2: Testing Hypotheses 4 and 5

From the Granger representation theorem (Engle and Granger 1987), we know that a cointegrated system has a vector error-correction representation. Tables 5 and 6 explore this and its implications for whether order flow can be considered exogenous, as predicted by hypotheses 4 and 5.

Table 5 presents evidence that order flow over the flexible-rate period is indeed weakly exogenous, as predicted by hypothesis 4. Recall that weak exogeneity of order flow means that error correction in the cointegrating relationship occurs through exchange rate adjustment, not order flow adjustment. Or, econometrically, it means that the error-correction term is significant in the exchange rate equation but not in the order flow equation. Table 5 shows that the null of weak exogeneity of order flow cannot be rejected at conventional levels (p-value 10 percent). In contrast, the null of weak exogeneity of the exchange rate is strongly rejected (p-value 1 percent).²⁶ It appears that the burden of adjustment to long-run equilibrium falls exclusively on the exchange rate.

Table 6 presents evidence that order flow is, in fact, strongly exogenous, as predicted by hypothesis 5. The key result is that for the test labeled F(8,67), noted in the table legend with "h". This is a test for any feedback to order flow from lagged values of either the exchange rate or the interest differential. Because exclusion of these variables from the general (unconstrained) model cannot be rejected, there is no evidence of Granger causality running from these variables to order flow (thus, there is no evidence of feedback trading). This combination of weak exogeneity and absence of Granger

²⁶ This test is based on the full information maximum likelihood approach of the Johansen (1992) procedure, which takes account of possible cross equation dependencies (as opposed to testing the significance of the error-correction term equation by equation).

causality implies that cumulative order is strongly exogenous.²⁷

The error-correction representation allows us to answer an important question: How rapidly does this system return to its long-run equilibrium? (Though important, this question is purely empirical in that our model makes no prediction.) The answer to this question comes directly from the estimate of the system's error-correction term. That estimate is -0.237 (reported in appendix Table A2), implying that about one-quarter of departures from long-run equilibrium is dissipated each day.²⁸

This result is also helpful for judging whether four months of data is sufficient to produce reliable analysis of cointegration. If one-quarter of any departure is dissipated each day, the half-life of departures is less than three days. Four months of data is enough to cover about 40 of these half-lives, quite a lot in the context of estimating cointegrating relationships. For comparison, adjustment back to the cointegrating relationship of Purchasing Power Parity (PPP) has a half-life around 5 years. One would need 200 years of data on a single exchange rate to estimate PPP error correction with as many half-lives in the sample. At the same time, we recognize that cointegration may not be literally true; if not, one is left with the result that order flow effects on price are very persistent, but not truly permanent.

5. Conclusions

Previous theoretical work offers three approaches to resolving the regime-dependent volatility puzzle: flexible-rate regimes induce either additional policy shocks (e.g., from greater policy autonomy), additional noise (e.g., entry of noise traders), or additional equilibria. Our explanation of the puzzle does not rely on any of these. Instead, flexible rates produce additional price volatility because the market's willingness to absorb unchanged shocks to order flow is reduced. (Though we chose not to do so, these shocks could be explicitly modeled as arising from, for example, changing risk preferences, changing endowment-risk hedging, or changing liquidity demands.) Under

²⁷ It is unlikely, but possible, that this lack of Granger causality is due to the six-hour mismatch in the timing of our order flow and exchange rate data. Remember, however, that little of the order flow in FF/DM occurs between 6pm and midnight GMT. For a test of an even stronger form of statistical exogeneity—strict exogeneity—see the appendix.

²⁸ For those less familiar with cointegration models, note that this result does not imply that order flow's effect on price is transitory: this error-correction estimate applies to *departures* from the long-run relationship, not to the long-run relationship itself. Note too that the error-correction term is significant in all the specifications we examined in Table A2 for the exchange rate equation, with little variation in its magnitude.

flexible rates, these shocks produce portfolio-balance effects on price because elasticity of speculative demand is (endogenously) low under flexible rates. Under credibly fixed rates, the elasticity of speculative demand is infinite, eliminating portfolio-balance effects.

Testable implications of our explanation for the puzzle are borne out in the data. We use a unique data set on FF/DM order flow in 1998 to show that before the rigid parity-rates were announced, cumulative order flow and the spot rate were cointegrated, as our model predicts. (This result emerges despite the FF/DM rate varying considerably less than major flexible rates such as DM/\$.) Thus, at least some of the effects of order flow on the exchange rate appear to be permanent. This is contrary to received wisdom: many people believe that order flow has only transitory “indigestion” effects on price, but this is not the case.²⁹ After the conversion rates for the euro-participating currencies were announced, the FF/DM rate was decoupled from order flow. The model we develop predicts this as well.

We also address the degree to which order flow can be considered exogenous, as our model assumes. Our findings are supportive in this regard. We find that order flow is (at least) strongly exogenous. This has two key implications. First, the burden of adjustment to long-run equilibrium falls on the exchange rate, not on order flow. Second, there is no Granger causality running from the exchange rate back to order flow (i.e., feedback trading does not appear to be present).

It is common to characterize fixed-rate regimes in terms of the central bank’s willingness to trade domestic currency at a predetermined price; i.e., it is the central bank that absorbs the order flow. In our model, a credible fixed-rate regime is one in which the private sector, not the central bank, absorbs innovations in order flow. As a practical matter, if the central bank needs to intervene, the fixed exchange rate regime is already in difficulty because the private sector’s demand is no longer perfectly elastic, and the portfolio-balance channel is operative.

Though we model a regime shift from flexible to fixed rates, it may be useful to revisit models of currency crises (i.e., shifts from fixed to flexible) with order flow’s role in mind. For example, our model directs attention to the variance of order flow shocks

(C_t^1 in the model) and to the evolving elasticity of order-flow absorption by the private sector (the γ coefficient in the model). Interest differentials and other empirical proxies for credibility and expected devaluation can be recast in terms of these other, now increasingly measurable variables. In this setting, policymakers take their cues on necessary adjustment of interest rates and reserves from the private order flows they observe in the market (not from a macro model). Order flows become the vehicle through which dispersed market information about credibility and fundamentals is manifested. At the same time, order flow is by no means a noiseless signal, leading (potentially) to learning dynamics quite different from those in macro collapse models (for work in this direction, see, e.g., Carrera 1999, Calvo 1999, and Corsetti, Pesanti, and Roubini 2001).

²⁹ Though received wisdom, it should also be noted that transitory price effects of significant size are difficult to reconcile with market efficiency anyway.


References

- Banerjee, A., R. Lumsdaine, and J. Stock, 1992, Recursive and sequential tests of the unit-root and trend-break hypotheses: Theory and international evidence, *Journal of Business and Economic Statistics*, 10: 271-287.
- BIS (Bank for International Settlements), 1999, *Central bank survey of foreign exchange market activity in April 1998*, publication of the Monetary and Economics Department, BIS, May (available at www.bis.org).
- BIS (Bank for International Settlements), 2001, *BIS 71st Annual Report*, June (available at www.bis.org).
- Bjonnes, G., and D. Rime, 2001, FX trading live! Dealer behavior and trading systems in foreign exchange markets, typescript, Norwegian School of Management, University of Oslo, March (www.uio.no/~dagfinri).
- Calvo, G., 1999, "Contagion in Emerging Markets: When *Wall Street* Is a Carrier," University of Maryland working paper.
- Carlson, J., and C. Osler, 2000, Rational speculators and exchange rate volatility, *European Economic Review*, 44: 231-253.
- Carrera, J., 1999, Speculative attacks to currency target zones: A market microstructure approach, *Journal of Empirical Finance*, 6, 555-582.
- Corsetti, G., P. Pesenti, and N. Roubini, 2001, The role of large players in currency crises, NBER Working Paper 8303, May.
- Devereux, M., and C. Engel, 1999, The optimal choice of exchange-rate regime: Price-setting rules and internationalized production, NBER Working Paper 6992.
- Diamond, D. and R. Verrecchia 1981. "Information Aggregation in a Noisy Rational Expectations Economy," *Journal of Financial Economics*, 9: 221-235.
- Engle, R., and C. Granger, 1987, Cointegration and error correction: Representation, estimation and testing, *Econometrica*, 55: 251-276.
- Evans, Martin, 2001, FX trading and exchange rate dynamics, NBER Working Paper 8116, February.
- Evans, M., and R. Lyons, 1999, Order flow and exchange rate dynamics, typescript, U.C. Berkeley (available at www.haas.berkeley.edu/~lyons/wp.html), forthcoming in the *Journal of Political Economy*.
- Evans, M., and R. Lyons, 2001, Portfolio balance, price impact, and secret intervention, NBER Working Paper 8356, July.
- Flood, R., and A. Rose, 1995, Fixing exchange rates: A virtual quest for fundamentals, *Journal of Monetary Economics*, 36: 3-37.
- French, K., and R. Roll, 1986, Stock return variance: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 99-117.
- Froot, K., and M. Obstfeld, 1991, Exchange Rate Dynamics Under Stochastic Regime Shifts: A Unified Approach, *Journal of International Economics*, 31: 203-229.
- Grossman, S., and J. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.
- Guembel, A., and O. Sussman, 2001, Optimal exchange rates: A market-microstructure approach, typescript, Said Business School, Oxford University.

- Hansen, B., 1992, Tests for parameter instability in regressions with I(1) processes, *Journal of Business and Economic Statistics*, 10: 321-335.
- Hasbrouck, J., 1991, Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.
- Hau, H., 1998, Competitive entry and endogenous risk in the foreign exchange market, *Review of Financial Studies*, 11: 757-788.
- Hau, H., W. Killeen, and M. Moore, 2000, The euro as an international currency: Explaining the first evidence, CEPR Discussion Paper No. 2510, July.
- Jeanne, O. and A. Rose, 1999, Noise trading and exchange rate regimes, NBER Working Paper #7104, April, forthcoming in the *Quarterly Journal of Economics*.
- Johansen, S., 1992, Cointegration in partial systems and the efficiency of single equation analysis, *Journal of Econometrics*, 52: 389-402.
- Krugman, P., and M. Miller, 1993, Why have a target zone? *Carnegie-Rochester Conference Series on Public Policy*, 38: 279-314.
- Lyons, R. 1995. "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 39: 321-351.
- Lyons, R., 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics*, 42: 275-298.
- Lyons, R., 2001, *The Microstructure Approach to Exchange Rates*, MIT Press: Cambridge (chapters at www.haas.berkeley.edu/~lyons), November.
- Machlup, F., 1949, The theory of foreign exchanges. Reprinted in *Readings in the Theory of International Trade*, edited by H. Ellis and L. Metzler, Philadelphia: Blakiston.
- 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.
- Rime, D., 2000, Private or public information in foreign exchange markets? An empirical analysis, typescript, Norwegian School of Management, University of Oslo, March (www.uio.no/~dagfinri).
- Rigobon, R., 1999, On the measurement of the international propagation of shocks, NBER Working Paper 7354, September 1999.
- Robinson, J., 1949, The foreign exchanges. Reprinted in *Readings in the Theory of International Trade*, edited by H. Ellis and L. Metzler, Philadelphia: Blakiston.
- Rosenberg, M., 1996, *Currency Forecasting: A Guide to Fundamental and Technical Models of Exchange Rate Determination*. Chicago: Irwin Professional Publishing.


Figure 1

Panel A: Log FF/DM and Cumulative Order flow (No. of DM Buy Trades minus No. of DM Sell Trades)



Source: EBS, Datastream

Panel B: Log FF/DM and Cumulative Order flow (DM Bought minus DM Sold)



Source: EBS, Datastream

Figure 2

Three Trading Rounds Within Each Day




Table 1
Hypotheses 1 and 2: Unit Root Tests (Dickey-Fuller)

Sample Period	Tau Statistic ^a	DF Lag ^b
<u>1 Jan 1998 - 1 May 1998 (85 obs)</u>		
Log FF/DM	-2.06	1
Interest Rate Differential (PIBOR -FIBOR)	-0.11	0
Cumulative Order Flow ^c	-1.23	0
<u>4 May 1998 - 31 December 1998 (174 obs)</u>		
Log FF/DM	-3.02	0
Interest Rate Differential (PIBOR -FIBOR)	-3.00	0
Cumulative Order Flow	-1.91	0

Hypothesis 2: Structural Break Tests on Log FF/DM (Banerjee et al. 1992)^d

Sequential Minimum ADF Unit Root Test	Maximum F-Test (p-value)^e 1492.6 (0.000)
Recursive Minimum ADF Unit Root Test	Minimum ADF Tau^f -4.41

a: DF 5% Critical Value for Unit Root Test with Constant: -2.84

b: Lag length in first differences in DF Equation needed to obtain white noise error

c: Cumulative Order Flow is the cumulated index of Daily **Cash** DM buys net of DM sells

d: Tests conducted under zero lags and no trend

e: Critical values for the sequential F-test are 16.72 (10%), 19.01 (5%) and 21.31 (2.5%).

f: Critical values for the recursive unit root tests are -3.24 (10%), -3.61 (5%) and -3.91 (2.5%),

The top portion of the table tests the null hypothesis of a unit root in each of the three variables in our model over the flexible -rate and fixed-rate sub-samples (January 1 to May 1 and May 4 to December 31, 1998, respectively). The unit-root null cannot be rejected for any of the three variables in the flexible-rate sub-sample. In the fixed-rate sub-sample, the unit-root null for cumulative order flow is still not rejected, but the null is rejected for the log spot rate and interest differential.

The bottom portion of table tests the hypothesis of a structural break in the log spot rate process more directly using two tests proposed by Banerjee et al. (1992). The sequential test estimates the process both before and after each observation in the full sample. The null of the associated F test—one for each observation—is that the process is nonstationary with no structural breaks in the constant or trend. The largest of these F statistics is compared to a tabulated critical value. The rejection of this no-breaks null using the sequential test is overwhelming.

The Recursive test shown in the bottom portion of the table also has a null of a unit root and no structural breaks in the constant or trend throughout the sample. In this case, one calculates a Dickey-Fuller-style unit root test for every sample size on an expanding basis. The smallest value (large and negative) of the test statistic is then compared against a critical value, obtained from Monte Carlo simulations. In this case the test rejects the no-break null, with the test statistic of -4.41 being well below the 5% critical value of -3.61.

Table 2**Hypothesis 1: Cointegration Test (Phillips-Hansen Estimation)^a**

1 Jan 1998 - 1 May 1998

Dependent Variable: Log FF/DM^b	Coefficient (St. Err.)	T-Statistic
Cumulative Order Flow	0.00243 (0.0003)	3.43
Interest Rate Differential (PIBOR-FIBOR)	0.1923 (0.15)	0.81
Trend	0.1762 (0.149)	5.21
Trivariate Dickey-Fuller Tau	-4.19	
Bivariate Dickey-Fuller Tau (excludes the interest rate differential)	-4.25 ^c	

a: Regression of Log FF/DM on a constant, a trend, and cumulative order flow

b: Log FF/DM is multiplied by 10,000, so coefficient units are basis points per million Deutschemarks

c: Cointegrating ADF 5% critical value with two variables, constant and Trend: -3.78

Table 3

Cointegrating Vector from Johansen Procedure

1 January 1998 - 1 May 1998

	Coefficient	T-Statistic ^a
Cumulative Net Order Flow	0.003	2.81
Trend	0.145 ^b	2.69

a: Tests for exclusion of each variable

b: The constant is unrestricted and is not reported. The trend is restricted to be within the cointegration vector.

Table 4**Hypothesis 3: Hansen Stability Tests**

Test Description	Jan-Dec 1998		Jan-April 1998	
	Test-Statistic	p-value	Test-statistic	p-value
F-test ^a	8.1	<5%	not applicable	not applicable
Mean-F ^b	6.3	4.7%	1.817	>20%

a: Test conditional on known break date (4 May 1998)

b: Test for an unspecified structural break

Table 5**Hypothesis 4: Testing for Weak Exogeneity of Order Flow**

1 January 1998 - 1 May 1998

Likelihood Test ^a	Chisq(1) ^b
Diff. FF/DM	6.95 (0.01)
Diff. Cum. Order Flow	2.80 (0.10)

a: Tests the null hypothesis that the variable is weakly exogenous

b: The p-value is in parentheses

Table 6

Hypothesis 5: Testing for Strong Exogeneity of Order Flow

1 Jan 1998 - 1 May 1998

Equation 1: Diff. Cum. Order Flow	Coeff. (Excl. Lags)	T-statistic (Excl. Lags)																																																
Dependent Variable: Diff. Cum. Order Flow																																																		
Constant	6.240	0.233																																																
Diff. Cum. Order Flow(-1)	0.258	2.41																																																
Diagnostics																																																		
<table border="0" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th colspan="3" style="text-align: left; border-right: 1px solid black;">General Order Flow Returns Model Specification^a</th> <th colspan="3" style="text-align: right;">Final Model Specification^b</th> </tr> </thead> <tbody> <tr> <td style="width: 33%;">R² (adjusted)</td> <td style="width: 33%;">4.0%</td> <td style="width: 33%; border-right: 1px solid black;"></td> <td style="width: 33%;">R² (adjusted)</td> <td style="width: 33%;">5.5%</td> <td style="width: 33%;"></td> </tr> <tr> <td>Q (36) p-value^c</td> <td>0.67</td> <td style="border-right: 1px solid black;"></td> <td>Q (36) p-value</td> <td>0.36</td> <td></td> </tr> <tr> <td></td> <td style="text-align: center;">F-Test^d</td> <td style="text-align: center;">p-value</td> <td></td> <td style="text-align: center;">F-Test</td> <td style="text-align: center;">p-value</td> </tr> <tr> <td>F(4,67)^e</td> <td style="text-align: center;">1.90</td> <td style="text-align: center;">0.12</td> <td style="border-right: 1px solid black;">F(11,67)ⁱ</td> <td style="text-align: center;">0.865</td> <td style="text-align: center;">0.570</td> </tr> <tr> <td>F(4,67)^f</td> <td style="text-align: center;">0.55</td> <td style="text-align: center;">0.69</td> <td style="border-right: 1px solid black;">F(1,67)^j</td> <td style="text-align: center;">5.37</td> <td style="text-align: center;">0.023</td> </tr> <tr> <td>F(4,67)^g</td> <td style="text-align: center;">1.43</td> <td style="text-align: center;">0.23</td> <td style="border-right: 1px solid black;"></td> <td></td> <td></td> </tr> <tr> <td>F(8,67)^h</td> <td style="text-align: center;">1.08</td> <td style="text-align: center;">0.35</td> <td style="border-right: 1px solid black;"></td> <td></td> <td></td> </tr> </tbody> </table>			General Order Flow Returns Model Specification ^a			Final Model Specification ^b			R ² (adjusted)	4.0%		R ² (adjusted)	5.5%		Q (36) p-value ^c	0.67		Q (36) p-value	0.36			F-Test^d	p-value		F-Test	p-value	F(4,67) ^e	1.90	0.12	F(11,67) ⁱ	0.865	0.570	F(4,67) ^f	0.55	0.69	F(1,67) ^j	5.37	0.023	F(4,67) ^g	1.43	0.23				F(8,67)^h	1.08	0.35			
General Order Flow Returns Model Specification ^a			Final Model Specification ^b																																															
R ² (adjusted)	4.0%		R ² (adjusted)	5.5%																																														
Q (36) p-value ^c	0.67		Q (36) p-value	0.36																																														
	F-Test^d	p-value		F-Test	p-value																																													
F(4,67) ^e	1.90	0.12	F(11,67) ⁱ	0.865	0.570																																													
F(4,67) ^f	0.55	0.69	F(1,67) ^j	5.37	0.023																																													
F(4,67) ^g	1.43	0.23																																																
F(8,67)^h	1.08	0.35																																																

a : General model includes constant, Diff.Cum Order Flow(-1 to -4), Diff. Pibor-Fibor(-1 to -4) and Diff. FF/DM(-1 to -4)

b: Final model includes constant and Diff. Cum. Order Flow(-1)

c: This tests for residual autocorrelation up to 36th order

d: **All F tests relate to exclusions from the general model**

e: Test for the exclusion of Diff. Pibor-Fibor (-1 to -4)

f: Test for the exclusion of Diff. Log FF/DM (-1 to 4)

g: Test for the exclusion of Diff. Cum. Order Flow (-1 to -4)

h: Test for exclusion of Diff. FF/DM(-1 to -4) and Diff. Pibor-Fibor (-1 to -4)

i: Test for exclusion of all variables **except** Diff. Cum. Order Flow(-1)

j: Test for exclusion of Diff. Cum. Order Flow(-1)

Appendix : Related Tests and Robustness

Testing Structural Shifts versus Changing Shock Variance (Rigobon 1999)

Rigobon (1999) proposes a powerful and flexible test for parameter stability in settings with possible heteroskedasticity, endogeneity, and omitted variables. (Previous tests are unable to handle all three simultaneously.) His DCC³⁰ test is based on the assumption that the data can be divided into two sub-samples: a tranquil and a turbulent period, with a known break. The null hypothesis is that heteroskedasticity is caused by a shift in the variance of just one of the shocks affecting the system, which can include structural shocks or innovations in unobservable omitted variables. The alternative hypothesis is that the heteroskedasticity is caused by parameter change. The test is simple to implement: the variance-covariance matrix is calculated for each sub-sample, the two estimated matrices are subtracted, and the determinant of their difference is calculated. Under the null, the determinant is zero. Large absolute values provide evidence of parameter instability.

Though Rigobon's test is designed for systems that exhibit change from tranquility to turbulence (e.g., contagious transmission of a financial crisis from one country to another), there is no reason why the test cannot be applied to systems that move in the opposite direction, such as ours. The null, in our case, is that the volatility of the FF/DM rate simply declined without any structural change at the beginning May 1998. The alternative is that fundamental parameter change occurred, which is what we contend (that the cointegrating relationship between the nominal exchange rate and cumulative order flow disintegrated).

Table A1 presents results for the DCC test in our sample.³¹ One of the maintained hypotheses for the test is that the decline in volatility is due to a decline in a single structural variance. The longer the second sub-sample, the more implausible that

³⁰ Acronym for *determinant* of the *change* in *covariance* matrix.

³¹ The distribution of the DCC test is non-standard and dependent on nuisance parameters. The standard deviations and p-values in the table are obtained by bootstrapping. A *Gauss* program to implement this is available from rigobon@mit.edu. For the three sub-samples studied in this paper, the distribution for the DCC statistic is available from the authors on request.

assumption is (lowering the power of the test). On the other hand, the shorter the second sub-sample, the less precise are the estimates of the covariance matrix, making the size of the test unreliable. To balance these considerations, we look at three different lengths for the second sub-sample. For the longest period—May to December 1998—parameter stability is rejected. For the shortest period—May 1998 only—parameter stability cannot be rejected. Both of these results are suspicious for the reasons given above. Most convincing is the result for the intermediate period, May and June 1998. This test rejects parameter stability at the 5% level, providing further support for treating the period before May 2/3 as parametrically different from the remainder of 1998.³²

Table A1

Second Sub-sample	DCC Test Value	p-value %
May 1998	-6327	22.7
May and June 1998	-12166	3.5
May to December 1998	-139816	0.0


Visual Evidence that Order Flow's Impact on Price Drops Off

Though not appropriate statistically (given the cointegrating relationship), it is nevertheless instructive to consider the time pattern of coefficients in a regression of exchange rate changes on order flows. The following Figure A1 provides some visual evidence from a rolling regression of FF/DM changes (50 observation window) on current and past (up to lag 5) order flows. The solid line plots the sum of those order-flow coefficients that are statistically significant in the January to May 1 sample (the

³² The DCC result does not permit us to conclude that the relationship between FF/DM and order flow changed from a cointegrating relationship to a non-cointegrating one, nor that May 2/3 is *the* date of change. A version of this test has not yet been developed for the case of an unknown break (nor can it be specific about the precise nature of the parameter instability).

coefficients are summed at each window-sample snap shot). The plot begins in mid-March to accommodate the 50 (trading day) rolling period. The plot indicates a clear drop off of order flow impact. The timing of the drop off is also suggestive of market anticipation of the May2/3 announcement (subject to the statistical caveat noted above).

Figure A1



Testing Strict Exogeneity of Order Flow

Given the evidence in section 4 that order flow is strongly exogenous, for completeness we test for strict exogeneity (a stronger concept than weak and strong exogeneity). Although our model also predicts strict exogeneity, the economic content of this prediction is less substantive than the notions of weak and strong exogeneity that underlie hypotheses 4 and 5 (which is why we address it in the appendix). One can test for strict exogeneity (given strong exogeneity) by testing the significance of contemporaneous order flow (i.e., differenced cumulative order flow) in the error-correction equation for the exchange rate. Including contemporaneous order flow amounts to a test for contemporaneous correlation across the two innovations in the two-equation vector error-correction representation.³³ If order flow is strictly exogenous, the coefficient on contemporaneous order flow should be insignificant.

From the result corresponding to note “e” in Table A2 (in bold), order flow does appear to be strictly exogenous: a test of excluding this variable cannot be rejected at conventional levels (p-value 34 percent). The top of the table presents coefficient estimates for the final error-correction model, which includes only those variables that enter significantly. The remainder of the table presents various diagnostics relevant to the two models (the general model, which includes all variables, and the final model with only the significant variables).

³³ If order flow is weakly exogenous, then our vector error-correction representation can be reduced to a single equation, and the coefficient on the contemporaneous variable takes the sign of the correlation between the two innovations of the system (Johansen 1992).

Table A2

Testing Strict Exogeneity

1 Jan 1998 - 1 May 1998

Equation 2: Diff. Log FF/DM	Coeff. (Excl. Lags)	T-statistic (Excl. Lags)
Dependent Variable: Diff. Log FF/DM		
Constant	0.210	1.45
ECM(-1)	-0.237	-3.10
Diff. Cum. Order Flow(-1)	0.002	2.74
Diff. Cum. Order Flow(-3)	-0.001	-2.36
Diff. Cum. Order Flow(-4)	0.001	2.66

Diagnostics

General FF/DM Returns Model Specification ^a				Final Model Specification ^b		
R ² (adjusted)	20.7%			R ² (adjusted)	24.5%	
Q (36) p-value ^c	0.97			Q (36) p-value	0.97	
	F-Test^d	p-value			F-Test	p-value
F(1,65)^e	0.94	0.34		F(11,65) ^j	0.637	0.778
F(1,65) ^f	6.15	0.02		F(4,65) ^k	5.88	0.000
F(4,65) ^g	0.06	0.99				
F(4,65) ^h	3.69	0.01				
F(4,65) ⁱ	0.72	0.58				

a :General Model: constant, Diff. Cum Order Flow(0 to -4), Diff. Pibor-Fibor(-1 to -4), Diff. FF/DM(-1 to -4) and ECM(-1)

b: Final model includes constant, ECM(-1) and Diff. Cum. Order Flow(-1,-3, -4); coefficient values reported at top of table

c: This tests for residual autocorrelation up to 36th order

d: **All F tests relate to exclusions from the general model**

e: Test for the exclusion of contemporaneous Diff. Cum. Order Flow

f: Test for the exclusion of ECM(-1)

g: Test for the exclusion of Diff. Log FF/DM (-1 to -4)

h: Test for the exclusion of Diff. Cum. Order Flow (-1 to -4)

i: Test for exclusion of Pibor-Fibor (-1 to -4)

j: Test for the exclusion of **all variables except** ECM(-1) and Diff. Cum. Order Flow (-1, -3, -4)

k: Test for the exclusion of ECM(-1) and Diff. Cum. Order Flow (-1,-3,-4)